Formalizing Interdisciplinary Collaboration in the CogSci Community

Lauren Oey* (loey@ucsd.edu)
Isabella DeStefano* (idestefa@ucsd.edu)
Erik Brockbank (ebrocbba@ucsd.edu)
Edward Vul (evul@ucsd.edu)

University of California, San Diego, Department of Psychology
9500 Gilman Dr., La Jolla, CA 92093 USA

Abstract

Is cognitive science interdisciplinary or multidisciplinary? We contribute to this debate by examining the authorship structure and topic similarity of contributions to the Cognitive Science Society from 2000 to 2019. We compare findings from CogSci to abstracts from the Vision Science Society over the same time frame. Our analysis focuses on graph theoretic features of the co-authorship network—edge density, transitivity, and maximum subgraph size—as well as clustering within the topic space of CogSci contributions. We also combine structural and semantic information with an analysis of homophily. We validate this approach by predicting new collaborations in this year’s CogSci proceedings. Our results suggest that cognitive science has become increasingly interdisciplinary in the last 19 years. More broadly, we argue that a formal quantitative approach which combines structural co-authorship information and semantic topic analysis provides inroads to questions about the level of interdisciplinary collaboration in the cognitive science community.

Keywords: co-authorship networks; topic modeling; interdisciplinaryity; multidisciplinaryity; scientometrics

Introduction

Since its foundation, the Cognitive Science Society sought to unify various disciplines of study under one interdisciplinary research field. Recently, criticism of the success of this mission has sparked a debate about whether cognitive science is fundamentally multidisciplinary rather than interdisciplinary (Núñez et al., 2019; Gray, 2019). The distinction between these community structures is subtle, making any claims favoring one or the other difficult to evaluate. Broadly, a research community might be considered to be more *multidisciplinary* if collaborations happen mostly within small groups and there is greater topical isolation of each group from the rest. On the other hand, a more *interdisciplinary* research community will show fewer isolated groups structurally and less separation of research interests across groups.

But how do we measure interdisciplinaryity in a way that captures meaningful differences within diverse communities? Currently, there is no consensus on a single measure that best aligns with this abstract concept. Previous studies quantified interdisciplinaryity by considering the journals as tags for different disciplines. Some of these studies have examined the distribution of journals cited (Goldstone & Leydesdorff, 2006; Porter, Cohen, Roessner, & Perreault, 2007; Núñez et al., 2019), the citation networks (Rafols & Meyer, 2010), and the journals that authors previously published in (Bergmann, Dale, Sattari, Heit, & Bhat, 2017). But this earlier research

aiming to quantify interdisciplinaryity was primarily targeted at the categorization of disciplines. These measures suffer from inconsistencies across classification systems, leading to variable conclusions (Wagner et al., 2011). Others have used departmental affiliation and educational background (Núñez et al., 2019; Schunn, Crowley, & Okada, 1998), but research interests often shift over the course of a lifetime which makes the affiliation label a transient indicator (Porter et al., 2007).

In the present work, we address the challenges of defining and measuring interdisciplinaryity through a combination of co-authorship network features, topic analysis, and assessment of graph homophily that unifies both structure and content of publications. We validate our measures using full papers from the Cognitive Science Society proceedings between 2000 and 2019 and abstracts from the Vision Science Society (only abstracts are submitted) over a similar time frame (2001 to 2019).

First, the degree to which a community is interdisciplinary or multidisciplinary may in large part be revealed by who collaborates with whom. Scientific collaboration can be represented as an undirected graph, in which nodes correspond to individual authors and edges between nodes indicate whether any two authors co-authored a paper together (Newman, 2001, 2004; Barabási et al., 2002). Co-authorships within a community containing multiple areas of study can range from highly integrated to highly modular, and the structure of the resulting co-authorship network will reflect this spectrum of possibilities.

Second, while the collaboration structure of a community no doubt reveals something about the modularity of interdisciplinary work that occurs within it, the ways in which research interests combine must play a role as well. To better understand how the *content* of collaborations informs the interdisciplinaryity of the field, we use a topic model (Griffiths & Steyvers, 2004) to extract high level patterns in cognitive science research over the last 19 years. Topic models have been used in previous research to capture trends in the published work within a discipline, including within cognitive science (Cohen Priva & Austerweil, 2015; Rothe, Rich, & Zhi-Wei, 2018). Studies specifically addressing interdisciplinaryity have used topic models to complement pre-defined discipline tagging (Nichols, 2014). In the present work, we apply clustering algorithms to the topics that authors study, addressing the separability of the interests and methods of researchers in

Figure 1: The co-authorship network of CogSci in 2000 and 2019 and the network of VSS in 2001 and 2019.

the field. More distinct clusters in topic space imply greater division between disciplines.

Finally, we propose a unified approach that draws on both the structure and the content of collaboration within the cognitive science community. We analyze the *homophily* of the co-authorship network with respect to the topic similarity of individual authors (McPherson, Smith-Lovin, & Cook, 2001). Homophily within a network measures the degree to which similar nodes (for any given similarity function) are likely to be connected. In a highly interdisciplinary community, we might expect people to be more likely to collaborate with those who have more distant interests from their own, while in a multidisciplinary community, such collaboration is less likely. Thus homophily provides a formal but intuitive metric for the amount of interdisciplinary collaboration in an authorship network.

Together, our analyses measure (1) interconnectedness in the co-authorship network structure, (2) clusters in the author topic space, and (3) homophily, unifying the co-authorship network and topic space. Not only do these metrics quantitatively illustrate how authorship within cognitive science has changed over time, but we also believe these measures may provide a meaningful contribution to the multidisciplinary-interdisciplinary debate across science¹.

Data

We retrieved 11,553 full text PDFs (with 12,203 unique authors) from the published *Proceedings of the Annual Meeting of the Cognitive Science Society* from 2000 to 2019².

¹All code used in this analysis can be found at: https://github.com/isabelladestefano/formalizing_interdisciplinary_collaboration

This data is primarily full text conference proceedings papers but also includes submitted abstracts. In addition, we retrieved 22,504 Vision Science Society Annual Meeting abstracts (with 23,842 unique authors) published in the *Journal of Vision* from 2001 to 2019³. Both data sets were processed to extract unique authors, publication year, and the full text of each paper or abstract.

Co-Authorship Network

Using the publication data collected from CogSci and VSS proceedings, we generated a co-authorship network for each year of the conferences with nodes representing authors and edges representing co-authored publications by pairs of authors in that year’s proceedings. The graphs were unweighted, i.e., edges represented whether two authors published together *at all* in a given year. We analyze three graph-theoretical measures which, when applied to the collaboration networks, provide insight into the level of interdisciplinarity within these conference communities: edge density, transitivity, and maximum subgraph size.

Edge density refers to the proportion of edges within the network relative to the theoretical maximum. Here, the theoretical maximum is determined by the number of edges possible given the total number of publications in that year. For every paper, there exists a fully connected subgraph of the paper’s authors with $n(n-1)/2$ edges, where n is the number of authors on that paper. Thus, the full set of N papers, and their associated number of co-authors, sets a theoretical maximum number of edges at $\sum_{i=1}^N \frac{n_i(n_i-1)}{2}$. We define edge density for a given year by normalizing the observed number of edges by this theoretical maximum (eq. 1).

$$\text{edge density} = \frac{|E(G)|}{\sum_{i=1}^N \frac{n_i(n_i-1)}{2}} \quad (1)$$

where $|E(G)|$ is the total number of edges in the co-authorship network G for that year, N is the total number of papers published in that year, and n_i is the number of authors on any given paper i . Our edge density metric measures the degree of repeated collaboration between any two authors, as a proportion of the amount of possible collaboration: a higher edge density indicates a higher rate of unique co-authorships. In an interdisciplinary community, we expect a *higher edge density*, indicating that authors tend to publish with a broad set of collaborators.

The edge density metric is shown in Figure 2. The edge density for both CogSci and VSS appears relatively stable over the range considered. Critically, we note that the edge density for VSS is significantly lower than CogSci ($\beta = -0.05$, $p < 0.001$) and, perhaps more importantly, the CogSci edge

²(a) 2000-2014 papers are hosted at <https://escholarship.org/uc/cognitivesciencesociety/>, retrieved 9 December 2018; (b) 2010-2019 papers are hosted at <https://mindmodeling.org/cogsciYEAR/>, retrieved 9 December 2018 and CogSci 2019 retrieved 3 December 2019

³All abstracts hosted at <https://jov.arvojournals.org/>, retrieved 6-8 January 2020

Figure 2: (a) Edge density, or the proportion of edges in the graph to the theoretical maximum given the number of papers and authors per paper. (b) Transitivity, or the proportion of authors whose co-authors also publish together. (c) The maximum subgraph size, or how many authors are in the largest island relative to the full graph.

density measure is relatively close to the theoretical maximum for this measure. This suggests that on average, CogSci authors publish with many unique authors.

Transitivity measures the probability of a node’s adjacent nodes also being connected by an edge, i.e., closed triads. Also referred to as the clustering coefficient, transitivity approximates the commonality of local clustering in the graph, such that higher transitivity indicates more clustering. Thus, we would expect an interdisciplinary community to have *lower transitivity*—authors publish with authors across group boundaries.

The transitivity for CogSci appears to decrease over time whereas the transitivity of VSS remains low over the range considered. Indeed the slope of a regression against year is significantly negative ($\hat{\beta} = -0.012, p < 0.001$), suggesting that the transitivity of the CogSci network is decreasing meaningfully. This could be influenced by a number of factors, including the possibility that authors have published more papers in the proceedings over time. Nonetheless, the decreasing transitivity suggests that collaborations are often between a more diverse set of individuals: that is, CogSci has become less “clique-y”.

The *size of the maximum subgraph* specifies the proportion of nodes in the graph that are connected to the largest island. A network with a large island relative to the overall size of the graph indicates that many authors are connected to many other authors through their co-authors’ and their co-authors’ co-authors’ collaborations. We would expect an interdisciplinary community to have a *large maximum subgraph*, reflecting the tendency of a large subset of the field to be connected in the same collaboration network.

Across both VSS and CogSci, the maximum subgraph appears to grow over the analyzed time period. Broadly, this suggests that the network of authors within the CogSci community has become increasingly interconnected: the positive slope of this increase in the CogSci data is significant ($\hat{\beta} = 0.014, p < 0.001$).

Figure 3: Frequent words from selected topics: these examples illustrate the level of granularity that the topic model is able to extract from the CogSci texts with 100 topics.

Topic Space

To extract the research topics studied by the cognitive science community, we used the `stm` package in R (Roberts, Stewart, & Tingley, 2014) to fit a topic model to the full text of the papers from the CogSci and VSS proceedings. `stm` provides functions for cleaning the data by removing punctuation, stopwords, and numbers, then lemmatizing the remaining text. Finally, we fit a Latent Dirichlet Allocation (LDA) topic model to the full text documents (Blei, Ng, & Jordan, 2003; Griffiths & Steyvers, 2004). In the model fitting process we specified 100 topics, which yielded niche yet enduring topics and methods, e.g. theory formation (Gopnik & Sobel, 2000), rational analysis (Chater & Oaksford, 1999), and connectionism (Rumelhart & McClelland, 1986). See Figure 3 for several examples of high probability words belonging to particular topics fit by the model. The topic model estimates a distribution over the 100 topics for each paper (or abstract); author locations in topic space were computed to be the overall distribution of their topics across all papers they had published in a given year. To alleviate unusually

Figure 4: K-means cluster analysis ($k = 5$) on topic space of authors, mapped onto 2 dimensions via MDS. The cluster maps show the clustering of topics studied by authors in the earliest (left) and most recent year (right) for both CogSci (blue) and VSS (orange). The line graph shows the ratio of within- to between-cluster sums of squares for each year. CogSci is becoming less clustered over time.

high spikes within topic distributions resulting from authors that publish only one paper, we smoothed the distributions by regularizing individual authors' topic distributions in a given year to the overall topic distribution for each year.

To understand how *integrated* the topics were year over year, we first applied multidimensional scaling (MDS) to the authors' distributions across the 100 topics to reduce the space to two dimensions, which is easier to visualize. We computed clusters on the scaled topic space of authors via k-means clustering (we used $k = 5$ which seemed to balance resolution of salient clusters and consistency across years). If authors are more clustered in topic space, that reflects less connectivity between disciplines and suggests a multidisciplinary community. To measure the separability of clustering across years, we computed the ratio of the within-cluster sum of squares to the between-cluster sum of squares based on the k-means centroids. A higher ratio reflects greater dispersion within clusters compared to between clusters, indicating that the clusters are not very separated—in other words, that authors are less siloed in disciplinary enclaves, as would be the case in a more interdisciplinary field.

The central plot in Figure 4 shows the ratio of the total within-cluster sum of squares to the between-cluster sum of squares for CogSci and VSS between 2000 and 2019. While VSS appears relatively stable (a regression on the data during this range is in fact negative: $\hat{\beta} = -0.006$, $p = 0.004$), the CogSci data has increased dramatically during this time ($\hat{\beta} = 0.014$, $p < 0.001$). Our results suggest that clusters in topic space have become less separable over time. The left and right sides of Figure 4 are the author clusters for the earliest and most recent years of the CogSci and VSS data sets. The increase in topic overlap (decreased separability) in the set of CogSci authors is apparent in the two plots while topic consolidation in VSS does appear more nominal.

Homophily

In the previous sections, we argue that structural measures of collaboration and general trends in topic space are both useful in trying to quantify interdisciplinarity. However, interdisciplinarity is not only about community structure and topic distributions alone, but about the distribution of topics studied *within* the co-authorship structure. In this vein, we combine the previous methods to better capture the interdisciplinarity of cognitive science. The most natural way to combine network structure and topic space is through the concept of *homophily*: do people who study the same topics publish with people who have similar topic interests or different interests? Since most of our analysis involves the evolution of the CogSci community, we want a method that will tell us if (1) homophily exists within the community and (2) To what extent does homophily predict new collaborations and persisting collaboration network structure. Our measure of homophily in the CogSci co-authorship network allows us to address both of these questions.

Here, we frame homophily in the CogSci network as a link prediction problem: how does the topic similarity of authors contribute to the likelihood that they would publish a paper together the following year? If the topic similarity between two authors is a strong predictor that they will publish together, then this suggests a degree of homophily in the authorship network: authors are more likely to publish with other authors who have similar research interests. In a multidisciplinary field, we expect a greater degree of homophily in the network: proximity between authors likely reflects a great deal of topic similarity, so topic similarity should be a strong predictor of whether authors will publish together in future years. In contrast, a more interdisciplinary field should see disproportionately more collaboration among authors that are more distant in topic space, yielding a weaker relationship

between topic similarity and future collaborations.

We measured similarity in topic space between two authors in a given year as the cosine similarity: the cosine of the angle θ between two authors' 100 topic vectors fitted by the topic model.

$$\cos(\theta) = \frac{a \cdot b}{||a|| ||b||} \quad (2)$$

The cosine similarity between each pair of authors was log-transformed to eliminate skewedness. Given the topic similarity between any two authors, the question is whether this is a strong predictor of future collaboration between those authors. Homophily predicts that two authors are more likely to publish together next year if they have similar topic distributions this year. Some amount of homophily is inevitable, assuming coherent and stable research interests; however, excess homophily may lead to insular research clusters, i.e., a more multidisciplinary network. On the other hand, in a more interdisciplinary community, homophily should be lower and topic similarity should be a weaker predictor of future collaborations.

Using the topic similarity between authors, we fit a logistic regression to the co-authorships during each year with the similarity between authors from the previous year as a predictor. To control for the auto-correlation in network structure (i.e. authors who published together one year are likely to publish together in the following year) we used a binary variable representing whether two authors had a "prior collaboration" the previous year as a covariate in our model.

We first look at whether topic similarity is a significant predictor of whether authors will publish together in a given year. Then, as a measure of homophily, we compare the strength of the similarity coefficient for CogSci and VSS regressions. After controlling for the covariate of a "prior collaboration", similarity was found to be a significant predictor of a new collaboration for both CogSci and VSS. For both CogSci and VSS the slope on similarity was positive ($\hat{\beta} = 2.248$, $p < 0.001$ and $\hat{\beta} = 2.598$, $p < 0.001$ respectively) suggesting that there is at least some level of homophily within both of these communities. A Wald-Z-test between the regression parameters of VSS and CogSci model revealed that the slope on similarity was significantly greater for VSS than for CogSci ($z = 5.486$, $p < 0.001$). This suggests that the VSS community may be more homophilous than the CogSci community.

To assess how ingrained the amount of homophily is, we looked at the interaction between "prior collaboration" and topic similarity between two authors. If the interaction term is positive, that means that topic similarity has a greater impact among prior co-authors, indicating that authors selectively publish again only with particularly like-minded collaborators. This would suggest an aversion to interdisciplinary collaboration within the network. If the interaction term between these two variables is negative, that means that topic homophily is lower among prior co-authors than among individuals who have not published together the previous year. If it is not only negative, but also larger than the positive main

effect of similarity, that would indicate that authors preferentially publish again with prior collaborators who are *less* similar to them than their other collaborators.

We do find that the slope of the interaction is negative for both CogSci ($\hat{\beta} = -3.862$, $p < 0.001$) and VSS ($\hat{\beta} = -3.960$, $p < 0.001$). As suggested above, since the slope on similarity is positive and the slope on the interaction between similarity and "prior collaboration" is both negative and larger in magnitude than the main effect, once two authors have collaborated they will be more likely to collaborate again in the following year if their topic interests are dissimilar. So while new co-authorship edges arise from topic homophily, the persisting collaborations tend to cross topics. This process suggests that the core network structure is fundamentally interdisciplinary since the collaborations that are likely to continue from year to year are between sets of authors with diverse interests.

To ensure the strength of all the predictors in our model, we compare the full model described above—which predicts new collaborations on the basis of prior publication, topic similarity, and their interaction—to lesioned models with only prior publication and only main effects. The full model outperformed both lesioned models (topic similarity: *deviance* = 1150, $p < 0.001$; interaction: *deviance* = 518, $p < 0.001$), suggesting that topic similarity and the interaction between topic similarity and prior publication improve predictions of novel collaborations.

Predicting 2020 Collaborations

Using the dynamic homophily regression with prior publication and topic similarity and their interaction as predictors and training data from CogSci 2000 to 2019, we generate predictions about who co-authors together in CogSci 2020. A subset of these predictions are shown in Figure 5a.

To evaluate the model's effectiveness, we compare the model's predictions to holdout data: the full set of collaborations from CogSci 2020⁴. Figure 5b shows a Receiver Operator Characteristic (ROC) curve (Swets, 1988). This curve is generated for *new collaborations only* (when there was no prior collaboration the previous year), since a model that predicts 2020 co-authorships might obtain high accuracy simply by assuming that authors who previously collaborated together will do so again. Instead, we are interested in the model's effectiveness predicting collaborations on the basis of the authors' topic similarity. We use the area under the curve (AUC) to evaluate how well our model predicted new publications; an AUC of 0.5 indicates chance performance and an AUC of 1 indicates perfect classification accuracy. We found our model had an $AUC = 0.689$, which indicates our model is well above chance when making predictions about new collaborations⁵. At the optimal threshold, the model's

⁴We thank the organizers of CogSci 2020 for providing us with the co-authorship data for the purposes of this analysis.

⁵Based on *all* model predictions (including authors with and without prior collaboration) we obtain an $AUC = 0.869$.

Figure 5: a) Prediction of co-authorships in 2020 for the 50 most eigencentral authors of 2019. The lighter the tile, the more likely our model predicts two authors will publish together. b) ROC curve created by using different thresholds on the probability of new publication to make binary predictions. We evaluate only cases where authors did not publish together in the previous year. The dotted line shows where an ROC curve would fall for a model making predictions at chance.

predictions have specificity (i.e., true negatives) of 0.802 and sensitivity (i.e., true positives) of 0.504.

Discussion

It has been argued that science is becoming more interdisciplinary across a broad range of research areas (Porter & Rafols, 2009). However, a recent debate in the cognitive science community raises questions whether the diverse fields that contribute to cognitive science pursue integrated research or are better described as multidisciplinary (Núñez et al., 2019; Gray, 2019). We argue that this discussion—and broader investigations into the interdisciplinary nature of research—is complemented by a formal, bottom-up treatment of the collaboration structure and content within the field. Using the full text and author data from 19 years of published proceedings of the Cognitive Science Society, we analyze the evolution of the co-authorship network and assess changes in topic space year over year. Since these methods are novel in their application and have not yet been externally validated, we provide a second example of how these same analyses apply to the full set of abstracts published in the Vision Science Society over a similar time period.

This bottom-up approach yields converging support for the claim that the Cognitive Science Society has become more interdisciplinary over the past two decades. First, the cognitive science co-authorship network shows that the field is becoming (structurally) less clustered and more interconnected, as evidenced by the decreasing transitivity of co-authorships and increasing maximum subgraph size. Second, co-authorship edge density, though more stable over time, is consistently

higher for CogSci than VSS, suggesting that CogSci authors tend to publish with more unique authors. Third, beyond the structure of collaboration networks in CogSci, we find that the clustering of authors by topic within the CogSci proceedings has become less separable over time. We argue that this provides some evidence that distinctions among disciplines may be shrinking. Finally, by combining co-authorship network and topic information, our analysis suggests that new collaborations reflect some degree of homophily—as they are predicted by prior collaborations and topic similarity—but this tendency to publish with similar authors is reversed for prior collaborators. These results suggest that although new co-authorships are driven by similar topics, persisting CogSci co-authorships are more likely to reflect interdisciplinary collaborations.

The use of topic modeling, characteristics of the co-authorship network, and the combination of the two offers a novel set of measures for understanding interdisciplinarity in a given field. The strength of these measures, apart from their formality, is the degree to which they are sensitive to the data in the research itself. Rather than pre-specifying the unique disciplines or fields within the community, we let graph clusters and topic separability speak to the connectedness of the research being done. This may allow for broader application across a range of distinct fields. However, such an approach may also suffer from its lack of structure.

First, the success of the approach depends on the richness of the data used. In the present investigation, we compare the full text of CogSci proceedings to abstracts published in VSS proceedings. One concern is that abstracts have a higher pro-

portion of general introductory material, and a smaller proportion of more idiosyncratic details. This would yield less divergence in topic space than a corpus of full length articles, thus potentially undermining a comparison of the two for measures of interdisciplinarity in topic space. Future work using data driven approaches should aim to use maximally similar data sets as the basis for any comparison across fields.

Second, results that rely on graph theoretic and topic space measures like the ones proposed here, which lack predefined sub-fields or topics, can be difficult to interpret. Although VSS serves as a pragmatic comparison point, future work should test our measures against theoretically motivated baselines, e.g., simulated interdisciplinary and multidisciplinary network structures and topic spaces. We believe that the possibility of combining top-down analysis techniques used in prior literature (Núñez et al., 2019) with network connectivity and topic similarity methods explored in the present work will allow for a more nuanced depiction of interdisciplinary collaboration. Building on our current results, future investigation into the interdisciplinary nature of research in cognitive science and beyond might draw not only on top-down measures of collaboration across disciplines, but on data driven measures of community structure and topic change.

Acknowledgments

We thank Rafael Núñez, Carson Miller Rigoli, Michael Allen, and Richard Gao for helpful discussion. We also thank the CogSci 2020 organizers for graciously sharing an early list of the authors published in this year's proceedings. This material is based upon work supported by the NSF Graduate Research Fellowship under Grant No. DGE-1650112.

References

- Barabási, A. L., Jeong, H., Néda, Z., Ravasz, E., Schubert, A., & Vicsek, T. (2002). Evolution of the social network of scientific collaborations. *Physica A*, 311(3-4), 590–614.
- Bergmann, T., Dale, R., Sattari, N., Heit, E., & Bhat, H. S. (2017). The interdisciplinarity of collaborations in cognitive science. *Cognitive Science*, 41(5), 1412–1418.
- Blei, D. M., Ng, A. Y., & Jordan, M. I. (2003). Latent dirichlet allocation. *Journal of machine Learning research*, 3, 993–1022.
- Chater, N., & Oaksford, M. (1999). Ten years of the rational analysis of cognition. *Trends in Cognitive Sciences*, 3(2), 57–65.
- Cohen Priva, U., & Austerweil, J. L. (2015). Analyzing the history of *Cognition* using topic models. *Cognition*, 135, 4–9.
- Goldstone, R. L., & Leydesdorff, L. (2006). The import and export of *Cognitive Science*. *Cognitive Science*, 30(6), 983–993.
- Gopnik, A., & Sobel, D. M. (2000). Detectingblickets: How young children use information about novel causal powers in categorization and induction. *Child Development*, 71(5), 1205–1222.
- Gray, W. D. (2019). Welcome to cognitive science: The once and future multidisciplinary society. *Topics in Cognitive Science*, 11(4), 838–844.
- Griffiths, T. L., & Steyvers, M. (2004). Finding scientific topics. *Proceedings of the National Academy of Sciences*, 101(1), 5228–5235.
- McPherson, M., Smith-Lovin, L., & Cook, J. M. (2001). Birds of a feather: Homophily in social networks. *Annual Review of Sociology*, 27(1), 415–444.
- Newman, M. E. J. (2001). The structure of scientific collaboration networks. *Proceedings of the National Academy of Sciences*, 98(2), 404–409.
- Newman, M. E. J. (2004). Coauthorship networks and patterns of scientific collaboration. *Proceedings of the National Academy of Sciences*, 101(suppl. 1), 5200–5205.
- Nichols, L. G. (2014). A topic model approach to measuring interdisciplinarity at the national science foundation. *Scientometrics*, 100(3), 741–754.
- Núñez, R., Allen, M., Gao, R., Miller Rigoli, C., Relaford-Doyle, J., & Semenuks, A. (2019). What happened to cognitive science? *Nature Human Behaviour*, 3(8), 782–791.
- Porter, A. L., Cohen, A. S., Roessner, D., & Perreault, M. (2007). Measuring researcher interdisciplinarity. *Scientometrics*, 72(1), 117–147.
- Porter, A. L., & Rafols, I. (2009). Is science becoming more interdisciplinary? measuring and mapping six research fields over time. *Scientometrics*, 81(3), 719–745.
- Rafols, I., & Meyer, M. (2010). Diversity and network coherence as indicators of interdisciplinarity: Case studies in bionanoscience. *Scientometrics*, 82(2), 263–287.
- Roberts, M. E., Stewart, B. M., & Tingley, D. (2014). stm: R package for structural topic models. *Journal of Statistical Software*, 10(2), 1–40.
- Rothe, A., Rich, A. S., & Zhi-Wei, L. (2018). Topics and trends in cognitive science (2000–2017). In *Proceedings of the 40th annual conference of the cognitive science society* (pp. 979–984). Austin, TX.
- Rumelhart, D. E., & McClelland, J. L. (1986). *Parallel distributed processing: Explorations in the microstructure of cognition, volume 1: Foundations*. Cambridge, MA: MIT Press.
- Schunn, C. D., Crowley, K., & Okada, T. (1998). The growth of multidisciplinary in the cognitive science society. *Cognitive Science*, 22(1), 107–130.
- Swets, J. A. (1988). Measuring the accuracy of diagnostic systems. *Science*, 240(4857), 1285–1293.
- Wagner, C. S., Roessner, J. D., Bobb, K., Klein, J. T., Boyack, K. W., Keyton, J., ... Börner, K. (2011). Approaches to understanding and measuring interdisciplinary scientific research (idr): A review of the literature. *Journal of Informetrics*, 165, 14–26.