



# Predicting the future success of scientific publications through social network and semantic analysis

Andrea Fronzetti Colladon<sup>1</sup> · Ciriaco Andrea D'Angelo<sup>2</sup> · Peter A. Gloor<sup>3</sup>

Received: 26 October 2019 / Published online: 29 April 2020  
© Akadémiai Kiadó, Budapest, Hungary 2020

## Abstract

Citations acknowledge the impact a scientific publication has on subsequent work. At the same time, deciding how and when to cite a paper, is also heavily influenced by social factors. In this work, we conduct an empirical analysis based on a dataset of 2010–2012 global publications in chemical engineering. We use social network analysis and text mining to measure publication attributes and understand which variables can better help predicting their future success. Controlling for intrinsic quality of a publication and for the number of authors in the byline, we are able to predict scholarly impact of a paper in terms of citations received 6 years after publication with almost 80% accuracy. Results suggest that, all other things being equal, it is better to co-publish with rotating co-authors and write the papers' abstract using more positive words, and a more complex, thus more informative, language. Publications that result from the collaboration of different social groups also attract more citations.

**Keywords** Social network analysis · Text mining · Social capital · Abstract · Citability · Scholarly impact

## Introduction

Measuring the value of a scientific publication is extremely complex but also crucial for many decisions related to research management and science policy. Scientific publications encoding new knowledge have different values, depending on their impact on future

---

✉ Ciriaco Andrea D'Angelo  
dangelo@dii.uniroma2.it

Andrea Fronzetti Colladon  
andrea.fronzeticolladon@unipg.it

Peter A. Gloor  
pgloor@mit.edu

<sup>1</sup> Department of Engineering, University of Perugia, Via G. Duranti, 93 – 06125 Perugia, Italy

<sup>2</sup> Department of Engineering and Management, University of Rome “Tor Vergata”, Via del Politecnico, 1 – 00133 Rome, Italy

<sup>3</sup> MIT Center for Collective Intelligence, 245 First Street, Cambridge, MA 02142, USA

scientific advancements and ultimately on social and economic development. As a proxy for such impact, bibliometricians adopt citation-based indicators. The choice of using citation indicators as a proxy for the impact of scientific production is based on assumptions deriving from sociology of science. In a narrative review of studies on the citing behavior of scientists, Bornmann and Daniel (2008), analyze the motivations that push scientists to cite the work of others. The findings show that “citing behavior is not motivated solely by the wish to acknowledge intellectual and cognitive influences of colleague scientists, since the individual studies reveal also other, in part non-scientific, factors that play a part in the decision to cite”. Nevertheless, “there is evidence that the different motivations of citers are not so different or randomly given to such an extent that the phenomenon of citation would lose its role as a reliable measure of impact”. In particular, previous literature proposes two different theories of citing behavior: the normative theory and the social constructivist view. The first, based on the work of Robert Merton (1957), affirms that scientists, through the citation of a scientific work, recognize a credit towards a colleague whose results they have used. In this case, the citation represents an intellectual or cognitive influence on their scientific work. The social constructivist view on citing behavior is based instead on constructivist theory in sociology of science (Knorr-Cetina 1981; Latour and Woolgar 1986). This approach contests the assumptions at the basis of normative theory and thus weakens the validity of evaluative citation analysis. Constructivists argue that “scientific knowledge is socially constructed through the manipulation of political and financial resources and the use of rhetorical devices” (Knorr-Cetina 1991), with the direct consequence that citations are not linked in direct and consequential manner to the scientific contents of the cited article. The bibliometric approach is based instead on the assumption that this link is strong and direct, meaning that citational analysis can be the principal instrument for evaluating the impact of scientific production.

We agree with this assumption and the Mertonian normative concept of what citations signify, although there might be exceptions (uncitedness, negative citations, fraudulent cross-citations, etc.). Although both theories have their merits, it still remains to understand: (1) which of them better explains the citability of a scientific work and; (2) whether there are other determinants of citability not covered by these two theories. On this last issue, we assume that there are “hidden honest signals” underlying the cognitive and intellectual process that produces a paper, that draw the attention of readers, influencing its citability, beyond its intrinsic quality and the social capital of its authors. Before “citing” a paper, a scholar needs to read it. Therefore maybe some semantic features related to the content of a paper (and, consequently, to its cognitive/intellectual appeal) might explain its readability and accessibility, and therefore its subsequent citability. When analyzing literature on a given topic, scientists generally rely on websites of journals’ publishers, on bibliometric platforms (WoS, Scopus), or on science social media (Mendeley, Academia, Researchgate, Google Scholar). Before downloading and reading the full text, they analyze the abstracts resulting from a specific search query. We wonder if some features of a publication’s abstract might affect its readability and, therefore, its citability. It is known that it matters “what” you publish and “with whom”: now, we want to investigate whether the “how” also counts, meaning “how an author sells” (in the abstract) the outcomes of her/his research to prospective readers and, as a consequence, to prospective citers.

In this work, we propose an empirical analysis based on a dataset of 2010–2012 worldwide publications in chemical engineering, indexed in SCOPUS. In particular, we compare publication metrics at the time of publishing, with scholarly impact of the paper 6 years after publication. Controlling for intrinsic quality of a publication, proxied by the impact factor of the journal and by the number of authors in the byline, we aim at

understanding the importance of three sets of predictive variables: structural social network metrics, dynamic changes in network position of authors, and complexity and sentiment of abstracts. The first two sets of variables complement the cardinality of the byline, in proxying the “social capital” of authors. The third set tries to catch cognitive/intellectual appeal of a paper based on semantic features of its summary offered to the reader.

Using machine learning, we are able to predict the impact of a paper in terms of numbers of citations 6 years after publication with 79% accuracy. We found that it is better to co-publish with many well-connected authors, write the abstract using more positive words, and employ a more complex, thus more informative, language.

The next section offers a picture of previous literature on different issues related to our analysis; the “[Data collection and methodology](#)” section illustrates methodological issues of the work, i.e. data collection and variables of the inferential model; the “[Results](#)” section presents results of the analysis; the “[Discussion and conclusions](#)” section closes the work discussing results and proposing concluding remarks.

## Literature review

Our paper aims at analyzing the predictability of long-term citations received by a publication observing the social capital of the authors in the byline and the features of its abstract. A summary of the main contributions of these two literature streams will be presented below.

### Social capital, research collaboration and impact of co-authored publications

The scientific environment is no different from other human activities by requiring to work in cooperation, because the individual scientist cannot possess all the competencies and resources needed for the resolution of the problem she/he is working on. Three concomitant factors help explain the remarkable increase of collaboration among scientists, research groups, and institutions, witnessed during the last decades: (1) the increasing complexity and cost of research to solve global societal problems, mostly interdisciplinary in nature (Bennett and Gadlin 2012; Persson et al. 1997); (2) the general reduction in travel costs, as well as the diffusion of inexpensive new communication technologies, in particular the Internet, which has greatly reduced the qualitative divide between distant and face-to-face communication (Hoekman et al. 2010; Olson and Olson 2000); (3) the existence of incentive systems towards collaborative research (Defazio et al. 2009). These factors have a systemic impact: at the level of individuals, they encourage scientists to increase their own “social capital”, defined as the whole of the resources obtainable through one’s social network (Jha and Welch 2010). Such resources include both the social network itself and those that are accessible via the network. For Nahapiet and Ghoshal (1998), social capital is a concept involving three dimensions: structural, cognitive and relational. The structural dimension concerns the general degree of connection and density of the network structure. The cognitive dimension concerns the sharing of knowledge between the actors of the network; the relational dimension concerns the quality of interpersonal relations in terms of trust, respect, friendship, etc. The relational dimension is the one that most influences the availability and use of resources in a social network of researchers (Burt 1995).

In the context of research systems, social capital is integral to the more encompassing concept of scientific and technical human capital (S&T human capital; STHC). Social

capital and STHC are highly interdependent. Each enables growth of the other. To be able to grow their social capital, scientists have to develop some basis of STHC over the course of their early career, in order to catch the interest of other colleagues (Bozeman et al. 2001; Dietz 2000; Murray 2005). It is no accident that scientists with tenure and the largest research projects tend to have larger, more heterogeneous and cosmopolitan, collaboration networks. They expand their networks beyond home institutions (Bozeman et al. 2001; Bozeman and Corley 2004) and countries (Melkers and Kiopa 2010). As social capital increases, the potential intensity and quality of research collaboration increases in parallel with growth in STHC. Scientists use their social networks for multiple purposes, including the identification and selection of collaborators (Beaver 2004; Katz and Martin 1997; Maglaughlin and Sonnenwald 2005). According to Wagner et al. (2015), future stars consciously build collaboration networks with other future stars well before they become famous. Sekara et al. (2018) have identified a “chaperone effect” where senior highly cited researchers help junior researchers in their team to establish themselves in a field and acquire senior status themselves. On the other hand, analyzing the scientific impact of a platform’s programming community that produces digital scientific innovations, Brunswicker et al. (2017) state that being surrounded by star performers can be harmful.

The impulse to undertake research collaboration studies has been supported by the development of specific bibliometric tools, which permit the measurement of different dimensions that characterize the phenomenon. In the literature, bibliometrics and the analysis of co-authorships have become the standard ways of observing research collaborations and measuring social capital. It should also be noted that co-authorships should be handled with care as a source of evidence for true scientific collaboration: this assumption has been questioned by many bibliometricians (Kim and Diesner 2015; Laudel 2002; Lundberg et al. 2006; Melin and Persson 1996). As Katz and Martin (1997) stated, some forms of collaboration do not generate co-authored articles and some co-authored articles do not reflect actual collaboration. However, in contradiction to the limitations noted above, this approach offers notable advantages both in terms of sample size (and consequent power of analysis) and of cost-effectiveness.

Social Network Analysis (SNA) is frequently used in the evaluation of the scientists’ social capital. The diffusion of collaboration studies based on SNA was particularly stimulated by Melin and Persson (1996), whose seminal study outlined procedures for the construction and analysis of co-authorship networks. In the literature on research collaboration, indicators of centrality have often been used in attempts to validate hypotheses related to social capital theory (Jha and Welch 2010; Nahapiet and Ghoshal 1998) and the contextual development of human and social capital (Bozeman et al. 2001; Bozeman and Corley 2004). A subject of great attention has been the so-called mechanisms of preferential attachment, meaning that when a scientist begins publishing, s/he will tend to collaborate with other scientists having a higher level of degree centrality (Barabási et al. 2002; Li et al. 2007; Perc 2010). In this manner, the cumulative advantage of the most popular scientists increases, in line with the *Matthew Effect* (Merton 1968), and the role of the hub within the network continues to strengthen. Previous research suggested the existence of a tight relationship between the number of authors in the byline and the long-term citation impact of publications (Abramo and D’Angelo 2015; Bornmann et al. 2014; Franceschet and Costantini 2010; Larivière et al. 2015; Matveeva and Poldin 2016; Waltman and van Eck 2015).

A few studies have focused on how centrality indicators of authors interact and affect citations for publications. In general, these studies claim that a papers’ citations are related to the node attributes of their authors in the collaboration network. The only exception was

presented by Wang (2014) who, exploring the Matthew effect, found no impact of authors' networking and prestige on solo-authored papers' citations. By contrast, working on a sample of more than 30 thousand authors in Google Scholar, Matveeva and Poldin (2016) discovered a positive relationship between scholars' citation counts and authors' centrality. Using wind-energy paper data collected from WoS, Guan et al. (2017) found that the structural holes of authors have positive but non-significant effects on a paper's citations, while the authors' centrality has an inverted U effect.

Lastly, Li et al. (2013) defined six specific indicators of co-authorship network characteristics according to the social capital theory and provided several strategies for leveraging social capital, meant to support scholars who want to enhance their research impact.

As better detailed in the “**Study variables**” section, for measuring the “social capital” of authors, along with the cardinality of the byline, we propose both structural social network metrics and the analysis of dynamic changes in network position of authors. To the best of our knowledge, this last set of variables represents a novelty compared to previous studies on the same topic.

### **The influence of textual content on the citability of publications**

Many authors have investigated the impact of factors other than intrinsic quality and authors' social capital on publication citations. Bornmann et al. (2014) showed that the number of cited references, and the number of pages are useful covariates in the prediction of long-term citation impact. Others have tested the effect of the presence of a country's name in the title (Abramo et al. 2016; Jacques and Sebire 2010; Nair and Gibbert 2016; Paiva et al. 2012) or of the ordering of authors in the byline (Abramo and D'Angelo 2017; Huang 2015; Ong et al. 2018; Shevlin and Davies 1997). Other studies have concentrated on the importance of the article title because, as Haggan (2004, p. 293) reasons, “the title plays an important role as the first point of contact between writer and potential reader and may decide whether or not the paper is read”. We point out a set of works on the relation between the structure of the title and citation rates (Habibzadeh and Yadollahie 2010; Jacques and Sebire 2010; Jamali and Nikzad 2011; Subotic and Mukherjee 2014). Falahati et al. (2015) conducted a morphological analysis of titles, to study the link between citability and title length/number of punctuation marks. The results of the analysis, made on a sample of 650 articles published in the journal *Scientometrics* over the years 2009–2011, show that: (1) title length and article citations are not correlated; (2) the number of punctuation marks does not serve as a reliable predictor of citations. Habibzadeh and Yadollahie (2010) studied the correlation between the length of an article title and the number of citations, for the area of the medical sciences. Longer titles seem to be associated with higher citation rates, with a larger effect for articles published in journals with a high impact factor. Using a sample including all the articles published in six PLOS journals, Jamali and Nikzad (2011) investigated the influence of the type of article title on the number of citations and downloads that an article receives. They observed that: (1) “question” articles tend to be downloaded more often, but cited less compared to others; (2) articles with longer titles are downloaded less than those with shorter titles; (3) titles with colons tend to be longer, and therefore receive less downloads and citations. Rostami et al. (2014) studied the association between some features of titles relative to the number of citations, examining the articles of the 2007 volume of *Addictive Behavior*: their results indicate that the type of title, as well as the number of keywords different from the words in the title, can contribute to predicting the number of citations. Uddin and Khan (2016) showed that

author selected keywords have a positive impact on the long-term citation count. van Wesel et al. (2014) focused their attention on what they call “superficial factors” influencing citations, including the number of words in title, number of pages, number of references, but also sentences in the abstract and readability in general. In fact, if the title plays an important role as a “touch point” for attracting the reader towards the manuscript, the abstract should do so even more by “advertising” its content and encouraging the full reading of the paper. According to Plavén-Sigray et al. (2017), the abstracts reflect the overall writing style of entire articles and “the readability of scientific texts is decreasing over time” and this should worry scientists and the wider public, as they impact both the reproducibility and accessibility of research findings. As for the influence of the abstract on the citability of a publication, Weinberger et al. (2015) found that shorter abstracts (fewer words and fewer sentences) consistently lead to fewer citations, with short sentences being beneficial only in Mathematics and Physics. Similarly, using more (rather than fewer) adjectives and adverbs is beneficial. Different conclusions are reached by Letchford et al. (2016) who found that journals publishing papers with shorter abstracts and containing more frequently used words receive on average slightly more citations per paper. Lastly, Freeling et al. (2019) suggested that increases in clarity, narrative structure, and creativity in the abstract of a paper could translate to a boost in citations it receives.

As better detailed in the “Study variables” section, in order to assess the possible dependence of citations accrued by a publication, by the cognitive/intellectual appeal of its content, we consider semantic features of the abstract and, specifically, its length, sentiment, complexity, diversity, and commonness. In terms of sentiment, our approach is partially explorative, as only few studies addressed the topic of extraction of opinions from scientific literature so far. In general, we would expect an objective, factual-based, communication style used in scientific abstracts—i.e. a more technical language than the one appearing on news, reviews or narrative texts (Athar 2011; Justeson and Katz 1995). However, some studies showed that technical terms can convey sentiment as well, and that “sentiment carrying science-specific terms exist and are relatively frequent” (Athar 2011 p.82; Athar and Teufel 2012; Athar 2014).

## Data collection and methodology

Our dataset is made of publications indexed in Scopus in 2010–2012 and hosted by sources tagged as “Chemical engineering” with respect to the ASJC (All Science Journal Classification) schema.<sup>1</sup> The choice of Scopus as bibliometric source is due to a powerful feature available on this repository, the author name disambiguation system<sup>2</sup>: for each publication SCOPUS provides not only the authors’ list but also a list of unique codes associated with each author. Kawashima and Tomizawa (2015) estimated the accuracy of the author identification in Scopus and found a recall and precision for Japanese researchers of about 98% and 99% respectively, which makes us particularly confident in terms of accuracy of the social networks that we will analyze.

The choice of the 3-year time window maximizes the tradeoff between computational effort and the robustness of the analysis (Wallace et al. 2012); in fact, scientific

<sup>1</sup> See [https://service.elsevier.com/app/answers/detail/a\\_id/15181/supporthub/scopus/](https://service.elsevier.com/app/answers/detail/a_id/15181/supporthub/scopus/) for details. Last accessed on March 19, 2020.

<sup>2</sup> <https://www.scopus.com/freelookup/form/author.uri>. Last accessed on March 19, 2020.

**Table 1** Bibliometric dataset

	Year	2010	2011	2012	Total
	Unique authors	199,497	224,462	241,205	498,598
	Publications	68,599	76,514	78,445	223,558
	Solo author paper	6.8%	5.9%	4.5%	5.7%
No. of authors	Average	4.22	4.33	4.49	4.35
	Max	202	125	37	202
	St. Dev	2.39	2.35	2.30	2.35
Cites	Average	32.6	29.5	26.1	29.3
	Max	5815	6759	6126	6759
	St. Dev	78.7	67.4	55.5	67.4

production is subject to uncertainty due to: (1) personal events, (2) patterns in research projects; (3) editorial and indexing processes (Luwel and Moed 1998; Trivedi 1993), (4) accidental facts and errors in bibliometric repositories (Karlsson et al. 2015). According to Abramo, D’Angelo, and Cicero (2012) a three-year publication period is appropriate for filtering randomness and assessing research performance and collaboration.

The focus on a specific field poses, on the one hand, problems of possible generalizability of results but, on the other hand, is necessary for a smaller-scale analysis as we are doing here, because all the variables at stake are field specific: the intensity of publication and citation, collaboration patterns, structural features of social networks, etc.

For the construction of the dataset we directly queried SCOPUS through the advanced search box, which returned almost 298,000 records. Given the aim of our analysis, it was necessary to eliminate about 74,000 of these results lacking impact metrics of the hosting source or abstracts. We focused in our analysis on research articles published on scientific journals—excluding reviews, conference papers, book chapters and other document types, such as letters, which appeared much less frequently. The final dataset was made of 223,558 publications, indexed in 657 unique sources. For each publication in the 2012 dataset we counted citations on January 1st, 2019, meaning that the citation window is 6 years. If we exclude the so-called “sleeping beauties”, a term coined by van Raan (2004) for indicating papers whose importance is not recognized for several years after publication, this is an adequate citation window for predicting long term impact of publications (Abramo et al. 2011), especially in chemical engineering, a subject category characterized by significant “immediacy”, i.e. high speed in reaching the peak of citations. As for the impact of the hosting source we use the Scimago Journal Ranking-SJR, 2012 edition (Guerrero-Bote and Moya-Anegón 2012).

As shown in Table 1, in this period, we register an increase in both the average number of co-authors per publication (from 4.22 in 2010 to 4.49 in 2012) and the share of “collaborative” publications (the share of solo-author papers drops from 6.8% in 2010 to 4.5% in 2012). These figures are fully in line with previous literature indicating a worldwide increase in scientific collaborations (Milojevi 2014), attested both by a rapid decline of the share of single-authored publications (Uddin et al. 2012), and by a significant increase in the average number of authors per publication (Larivière et al. 2015).

## Study variables

As described in the previous sections, our intent is to evaluate the importance of authors' social capital and semantic structure of abstracts, in predicting scientific success of papers, measured in terms of citations received 6 years after publication.

In doing so, we must control for the number of authors in the byline and for the impact factor of the hosting source. Journal impact metrics are generally aggregated measures of the impact of hosted articles: high impact articles are published in high impact journals and viceversa (Leimu and Koricheva 2005; Mingers and Xu 2010). Of course, there are evident exceptions and bibliometricians suggest not to use impact factors for measuring the quality and impact of individual publications (Marx and Bornmann 2013; Moed and van Leeuwen 1996; Petersen et al. 2019; Weingart 2005). However, here we must control for the intrinsic quality of a paper without having any other information available than the impact of the hosting journal (in our case the SJR).

As for the social capital of authors, the publication data retrieved from Scopus allowed us the construction of two social networks: the first, which we call *author network*, linking authors who collaborated in the writing of one or more papers; the second, which we call *publication network*, linking publications which share one or more authors. Both networks correspond to undirected graphs, where we indicate with  $n$  the number of nodes and  $m$  the total number of edges. In the author network, nodes represent scholars and there is an edge between two nodes if the corresponding scholars wrote at least one paper together; edges are weighted according to the number of co-authored papers. We use this network to evaluate the social capital of authors and their co-publication patterns. In the publication network, on the other hand, nodes represent publications, connected by edges weighted by the number of authors they share. Therefore, if paper A shares three authors with paper B, there will be a link connecting nodes A and B of weight equal to three. This second network tracks the social position of a publication, given the relationships maintained by its authors. Considering the above-mentioned graphs, we were able to calculate well-known centrality metrics, in order to study the network position of each publication and of its authors.

*Degree centrality* It corresponds to the number of direct links of a network node, weighted by summing the weights of its adjacent arcs (Freeman 1979; Wasserman and Faust 1994). In the author network, it represents the total strength of the direct connections a node has. In the publication network, it counts how many times the authors of a paper are shared with other papers in the network.

*Betweenness centrality* This very well-known centrality metric measures how many times a node lies in-between the shortest network paths that connect the other nodes (Freeman 1979; Wasserman and Faust 1994). Nodes with high betweenness centrality often serve as indirect connection between other pairs of nodes, thus having high brokerage power (Borgatti et al. 2013). Betweenness of node  $i$  can be calculated according to the following formula (Wasserman and Faust 1994):

$$B(i) = \sum_{j < k} \frac{g_{jk}(i)}{g_{jk}}$$

where  $g_{jk}$  is the number of shortest network paths linking the generic pair of nodes  $j$  and  $k$ , and  $g_{jk}(i)$  is the number of that paths that include node  $i$ . The formula can be normalized dividing it by its maximum  $(n - 1)(n - 2)/2$ .

*Closeness centrality* It measures the embeddedness of a node in the social network. The higher the closeness of a node, the shorter the network paths that connect it to its peers. To



put it in other words, closeness is measured as the reciprocal of the sum of the length of the shortest paths between the node and all other nodes in the graph (Freeman 1979; Wasserman and Faust 1994):

$$C(i) = \frac{1}{\sum_{j=1}^n d_{ij}}$$

where  $d_{ij}$  is the length of the shortest path connecting nodes  $i$  and  $j$ . Closeness can be normalized, multiplying its value by  $(n - 1)$ , which is its maximum and reflects the case of node  $i$  being adjacent to all other nodes.

*Constraint (Structural Holes)* It measures the value of network constraint, for each node (either author or publication), as presented in the work of Burt (1995). The idea behind this metric is that nodes which can mediate across unconnected peers are less constrained by their ego-network, thus also having higher social capital (Burt 2004). For instance consider an example with three nodes, A, B and C, where A is linked to B and C, but a link between these last two is missing. That missing link is called “structural hole” and gives social advantage to A that could mediate interactions between B and C, thus being less “constrained” by its ego-network. This is something A could not do, if B and C were directly connected.

*Rotating leadership* It counts the number of oscillations in betweenness centrality an author has in the network, considering subsequent publication years, i.e. if the author’s betweenness centrality changes significantly from one year to the other, reaching local maxima or minima (Allen et al. 2016; Kidane and Gloor 2007). Rotating leaders are authors which frequently change their network position, not remaining statically central or peripheral. This metric largely proved its potential in past research, which showed, for example, that rotating styles can favor both online community growth (Antonacci et al. 2017) and startups’ innovative performance (Allen et al. 2016).

The first four SNA metrics are calculated for both the *author network* and the *publication network*. The *Rotating Leadership* relates to the *author network* only, so that we have a total of nine metrics.

Analyzing the abstract of each publication, we derived metrics of text mining and semantic analysis, to see which variables related to publication content affect its future scholarly impact. Prior to the calculation of these metrics, we processed abstracts in order to remove those words which give little contribution to the text, such as the words “the” or “and”, also known as stop-words. Moreover, we removed word affixes to reduce each word to its stem—a procedure known as stemming, which was carried out using the NLTK package and the Python programming language (Perkins 2014). After this preprocessing phase, we proceeded in calculating:

*Abstract length*, i.e. the number of text characters in the abstract.

*Sentiment* It measures the positivity or negativity of the language used in a paper abstract, by means of the VADER rule based model for sentiment analysis (Hutto and Gilbert 2014), included in the NLTK python package. Values range from -1 to 1, where positive values represent a positive average sentiment and negative values correspond to the expression of negative feelings. Even if not context-specific, the VADER lexicon showed a good performance in past research (e.g., Hutto and Gilbert 2014; Newman and Joyner 2018).

*Complexity* Lexical complexity of an abstract is measured by looking at the standard deviation of the frequency distribution of words used in the text. This metric—successfully used in past research (e.g., Fronzetti Colladon and Vagaggini 2017; Gloor et al. 2017a,

b)—originates from the idea that there is a number of common words which will occur more often in a text, but when more complex ideas are presented different words will appear, thus increasing the variance of the word frequency distribution. Higher scores indicate higher complexity.

*Lexical diversity* Is measured as the ratio of different unique word stems to the total number of words used in an abstract (Malvern et al. 2004).

*Commonness* This metrics examines the uniqueness of words used in each abstract, based on their overall frequency in all text documents. In a first step, the overall frequency of each word is computed (excluding stop-words and after stemming), considering all abstracts. Subsequently, frequencies are averaged for all words of a single abstract, to assess its commonness. If words used are common to all other abstracts then commonness will be high. Conversely, distinctive abstracts will use words that appear less frequently.

We also tested other variants for complexity, lexical diversity and commonness metrics. One approach was to measure complexity as the likelihood distribution of words within an abstract, i.e. the probability of each word to appear in the text based on the term frequency/inverse document frequency (TF-IDF) information retrieval metric (Brönnimann 2014). However, different metrics did not lead to better results.

In the end we have:

- two control variables: the SJR of the hosting journal and the number of co-authors of the publication;
- nine variables related to social capital of its authors, i.e. their social network position and oscillations ( $X_1$ – $X_9$ ), and;
- five variables related to article content, measured by the semantic analysis of its abstract ( $X_{10}$ – $X_{14}$ ).

Table 2 shows the main descriptive statistics for all the above variables. Note that the networks were built considering all publications in the dataset (2010–2012). To properly assess authors' collaboration patterns but in order not to use future information, predictions were carried out only for 2012 publications,<sup>3</sup> excluding those with incomplete data (for the byline, abstract, citation count, or SJR).

## Results

Table 3 shows the correlations of the variables at stake. Since they are often not normally distributed and the relationships among them not necessarily linear, we used a nonparametric approach, i.e. the Spearman's rank correlation (Spearman 1904).

As the table shows, many of our predictors significantly and positively correlate with the number of citations accrued by publications after six years. Journal ranking is the one with the strongest correlation. In addition, the number of authors and their position in the author network seem to play an important role: citations are higher for those papers whose authors are more central in terms of direct connections (degree centrality) and betweenness centrality. It could be that more connected authors can leverage their social capital to diffuse

<sup>3</sup> - This prevent the need for normalizing citation count, since all publication used for prediction are of the same year and subject field.

**Table 2** Descriptive statistics of variables used in the analysis

Variable code	Variable name	Unit of analysis	M	SD
Y	Citations	Publications 2012	20.65	32.092
SJR	SJR	Publications 2012	1.630	1.261
No. of authors	Number of Authors	Publications 2012	4.400	2.184
X <sub>1</sub>	Degree—publication network	Publications 2010–2012	18.760	24.091
X <sub>2</sub>	Constraint—publication network	Publications 2010–2012	0.395	0.298
X <sub>3</sub>	Closeness—publication network	Publications 2010–2012	0.233	0.291
X <sub>4</sub>	Betweenness—publication network	Publications 2010–2012	1.417 × 10 <sup>-5</sup>	5.050 × 10 <sup>-5</sup>
X <sub>5</sub>	Degree—author network	Publications 2010–2012	19.619	21.627
X <sub>6</sub>	Constraint—author network	Publications 2010–2012	0.483	0.238
X <sub>7</sub>	Closeness—author network	Publications 2010–2012	0.291	0.321
X <sub>8</sub>	Betweenness—author network	Publications 2010–2012	7.393 × 10 <sup>-5</sup>	2.652 × 10 <sup>-4</sup>
X <sub>9</sub>	Rotating Leadership	Publications 2010–2012	2.130	2.083
X <sub>10</sub>	Abstract Length	Publications 2012	1135.950	431.963
X <sub>11</sub>	Sentiment	Publications 2012	0.516	0.525
X <sub>12</sub>	Complexity	Publications 2012	0.869	0.395
X <sub>13</sub>	Diversity	Publications 2012	0.745	0.099
X <sub>14</sub>	Commonness	Publications 2012	23,567.090	5610.448

their research and get more citations. Rotating leadership is also positively correlated with citations, supporting the idea that a bigger network dynamism of scholars is rewarded with more citations. Similarly, all network metrics related to the centrality of papers in the publication network significantly correlate with citations received. It could be that being highly cited is not just a matter of journal ranking, but also depends on the level of embeddedness in the two social networks we study. Consistently network constraint correlates negatively both for the author and the publication network, suggesting that when ego networks are more open, with more structural holes, there can be advantages of mediation across different social groups. Authors that have the power to link unconnected peers could be more effective in diffusing their ideas and research (Burt 2004). Similarly, papers that enable the collaboration of unconnected social groups could attract citations from a larger audience. On the other hand, metrics extracted from the analysis of paper abstracts seem to play a minor role; among them, abstract length is the one with the highest correlation. Of course these are just exploratory speculations, as correlation only reveals associations, without taking into account the combined effects of variables. For this reason, we extended the analysis with the intent of building a more comprehensive forecasting model that allows the identification of future highly cited papers—in particular those that, six years after publication, receive a number of citations high enough to be in the uppermost quartile.

We trained a parallel tree boosting machine learning model, namely XGBoost (Chen and Guestrin 2016), whose results are presented in Table 4. The model has been trained on 75% of observations and its performance has been subsequently evaluated considering the remaining 25% of data (out of sample). This process of random sampling without replacement of the training set and forecasting (on the remaining test set) has been repeated 300 times, i.e. we used Monte-Carlo cross validation (Dubitzky et al. 2007). We also evaluated the forecast performance of other algorithms, such as random forests (Breiman 2001), without getting to better results. Similarly, we tested other possible selections of highly

**Table 3** Spearman's correlation coefficients for variables used in the analysis

Y	SJR	No. of authors	X <sub>1</sub>	X <sub>2</sub>	X <sub>3</sub>	X <sub>4</sub>	X <sub>5</sub>	X <sub>6</sub>	X <sub>7</sub>	X <sub>8</sub>	X <sub>9</sub>	X <sub>10</sub>	X <sub>11</sub>	X <sub>12</sub>	X <sub>13</sub>
Y	1.000														
SJR	.574**	1.000													
No. of authors	.222**	.128**	1.000												
X <sub>1</sub>	.287**	.233**	.398**	1.000											
X <sub>2</sub>	-.248**	-.207**	-.337**	-.835**	1.000										
X <sub>3</sub>	.108**	.108**	.148**	.367**	-.142**	1.000									
X <sub>4</sub>	.253**	.206**	.462**	.582**	-.644**	.200**	1.000								
X <sub>5</sub>	.308**	.244**	.451**	.945**	-.822**	.315**	.564**	1.000							
X <sub>6</sub>	-.305**	-.246**	-.517**	-.764**	.736**	-.271**	-.735**	-.820**	1.000						
X <sub>7</sub>	-.019**	-.008	.020**	-.102**	-.026**	.485**	-.078**	-.061**	.161**	1.000					
X <sub>8</sub>	.303**	.250**	.320**	.859**	-.815**	.249**	.635**	.879**	-.782**	-.184**	1.000				
X <sub>9</sub>	.282**	.209**	.536**	.752**	-.603**	.268**	.615**	.737**	-.728**	-.126**	.660**	1.000			
X <sub>10</sub>	.146**	.011**	.090**	-.0005	0.007	-.042**	.030**	.016**	-.023**	-.011**	.017**	.019**	1.000		
X <sub>11</sub>	.050**	-.019**	.028**	.051**	-.045**	.010*	.029**	.048**	-.039**	-.011**	.048**	.252**	.046**	1.000	
X <sub>12</sub>	0.006	-.084**	.040**	0.004	.012**	-.013**	-.019**	.013**	.020**	.016**	0.003	.008*	.591**	.134**	1.000
X <sub>13</sub>	.031**	.098**	-.013**	0.002	-.021**	.020**	.031**	-.0001	-.042**	-.009*	0.004	0.006	-.546**	-.120**	-.917**
X <sub>14</sub>	-.020**	-.012**	-.076**	.106**	-.080**	.010*	.038**	.077**	0.005	-.048**	.106**	.057**	-.042**	.109**	-.108**

Statistical significance: \*\*\* $p < 0.001$ ; \*\* $p < 0.05$ ; \* $p < 0.1$

Y = Citations, X<sub>1</sub> = Degree—publication network, X<sub>2</sub> = Constraint—publication network, X<sub>3</sub> = Closeness—publication network, X<sub>4</sub> = Betweenness—publication network, X<sub>5</sub> = Degree—author network, X<sub>6</sub> = Constraint—author network, X<sub>7</sub> = Closeness—author network, X<sub>8</sub> = Betweenness—author network, X<sub>9</sub> = Rotating Leadership, X<sub>10</sub> = Abstract Length, X<sub>11</sub> = Sentiment, X<sub>12</sub> = Complexity, X<sub>13</sub> = Diversity, X<sub>14</sub> = Commonness

**Table 4** Feature importances

Variable	Mean absolute SHAP values	SD absolute SHAP values
SJR	1.421	0.025
No. of Authors	0.626	0.049
X <sub>1</sub> , Degree—publication network	0.070	0.018
X <sub>2</sub> , Constraint—publication network	0.053	0.012
X <sub>3</sub> , Closeness—publication network	0.047	0.011
X <sub>4</sub> , Betweenness—publication network	0.309	0.028
X <sub>5</sub> , Degree—author network	0.042	0.009
X <sub>6</sub> , Constraint—author network	0.047	0.009
X <sub>7</sub> , Closeness—author network	0.064	0.013
X <sub>8</sub> , Betweenness—author network	0.068	0.012
X <sub>9</sub> , Rotating Leadership	0.808	0.048
X <sub>10</sub> , Abstract length	0.169	0.012
X <sub>11</sub> , Sentiment	0.079	0.010
X <sub>12</sub> , Complexity	0.097	0.013
X <sub>13</sub> , Diversity	0.103	0.014
X <sub>14</sub> , Commonness	0.063	0.007

cited papers—for example considering the upper quintile instead of quartile—and obtained results similar to those we present here.

Accuracy of predictions was quite good and stable across 300 random repetitions, with the model returning, on average, correct answers in 79.2% of cases, with an average score of 0.41 for the Cohen’s Kappa and of 0.70 for the Area Under the ROC-curve. These results seem quite promising when compared with those reported by Abramo et al. (2019a, b) on a dataset of publications submitted to the first Italian research assessment exercise (VTR 2006), exclusively based on peer review. Contrasting the peer review rating with long term citation scores, the authors obtained a 75% agreement and a Cohen’s k equal to 0.172.

It is also important to notice that our main goal was not to obtain a 100% accurate model; more than finding the perfect forecast, we were interested in identifying variables that could be more relevant when predicting citations. Accordingly, Table 4 shows the importance of each predictor, calculated as the average of its absolute SHAP values (Lundberg and Lee 2017): the higher the score reported in the table, the more important the predictor. SHAP stands for SHapley Additive exPlanations and is a well-known evaluation approach, applicable to the output of different machine learning models. This method showed better consistency than previous approaches (Lundberg and Lee 2017) and proved to be particularly appropriate for tree ensembles (Lundberg et al. 2020). These last analyses were carried out using the Python programming language, specifically the packages SHAP (Lundberg and Lee 2017) and XGboost (Chen and Guestrin 2016).

Consistent with the results of the correlation analysis, we find that journal ranking is the most important predictor of highly cited papers, followed by rotating leadership, the number of authors and betweenness centrality in the publication network. It seems that social capital plays a role in terms of authors’ direct connections with peers, who could read and cite their papers. Keeping a dynamic position is also important. In addition, papers which result from the collaboration of different social groups also get more citations. Lastly,

writing longer, more informative abstracts seems to contribute a little to the improvement of model performance. The other variables, on the other hand, contribute little to our model predictions.

Journal ranking is by far the most important feature to forecast future citations and scholarly impact. Indeed, we notice that our sample comprises about 8000 papers that are both highly cited and published in top journals. However, a smaller number of papers, about 500, has the peculiar characteristic of being highly cited even if published in journals that have very low rankings (bottom 25% of the SJR distribution). How is that possible? We explored the differences between these two sets of papers through the t-tests presented in Fig. 1.

Apart from commonness, all the variables are significantly different. Successful papers published in low SJR journals seem to present more positive results (higher sentiment in the abstract) and new ideas (higher complexity), and have longer abstracts (even if this could be influenced by journal policies). Both these papers and their authors are closer to the network core (closeness is higher). Surprisingly, the number of authors and their connections—as well as betweenness centrality and rotating leadership—are lower with respect to papers in the top citations quartile, published in top journals. It seems that focused network embeddedness is the major driver of success for this set of papers (low SJR, high citations). It is not just a matter of being close to the network core, but also being part of a compact group with few structural holes. We speculate that in these cases unity is strength.

## Discussion and conclusions

In our research, we examined several characteristics of scientific papers which help predict their scholarly impact 6 years after publication. Results of our parallel tree-boosting machine learning model confirm findings of previous research, which indicate that journal impact factor and number of authors have a significant and positive effect on citations (Abramo and D'Angelo 2015; Bornmann et al. 2014; Leimu and Koricheva 2005; Mingers and Xu 2010; Waltman and van Eck 2015). We used these metrics as control variables and combined them with measures of social network and semantic analysis, which allowed the identification of highly cited papers with 79.2% accuracy. We found that authors' social capital has a role in attracting citations, thus publishing papers with well-connected authors can be an advantage. However, this effect is relatively small if compared with authors' rotating leadership, i.e. the ability to frequently change position in the collaboration network, moving back and forth from center to periphery. Indeed, authors' rotating leadership (change in betweenness centrality) emerged as one of the most important predictors of highly cited papers: it is not just a matter of authors' brokerage power, i.e. the ability to bridge connections across different social groups; authors' ability to activate bridging collaborations and subsequently leave space to others, without keeping dominant or static positions, was the third most important predictor. This is consistent with previous research showing that rotating leaders foster community growth and participation (Antonacci et al. 2017) and that dynamic social styles can favor innovation and knowledge sharing (Allen et al. 2016; Davis and Eisenhardt 2011). Accordingly, our study extends the research on the forecasting of scholarly impact, giving evidence to the contribution of new metrics of social network analysis, such as rotating leadership. In particular, we analyzed two social networks over a period of three years: the first, linking authors based on their scientific




Fig. 1 Characteristics of highly cited papers published in low ranked journals. (*T* tests, \*\*\* $p < .001$ ; \* $p < .05$ )

collaborations; the second, considering the social position of scientific papers based on their shared authors. The analysis of this second network revealed another important factor of publication success: scientific papers that resulted from the collaboration of different social groups—whose betweenness centrality was therefore higher—were more frequently ranked among the highly cited papers.

Predictors related to the semantic analysis of paper abstracts exhibited a lower, yet significant, importance. In particular, longer and more informative abstracts, whose texts have

a higher lexical diversity, seem to attract more citations. In this regard, our findings are aligned with research showing that shorter abstracts lead to fewer citations (Weinberger et al. 2015) and contrast with the study of Letchford et al. (2016) which proves the opposite. Our results also support the idea that abstracts that are more creative and diversified can attract more citations, as discussed by Freeling et al. (2019).

As a last step of analysis, we examined those papers which represented an exception to the idea that journal ranking plays a major role in attracting citations. In particular, we found about 500 articles which were published in low SJR journals but were highly cited. We compared them with highly cited papers published in top journals. Distinctive characteristics of successful low-SJR papers are that they present more positive results—abstract sentiment is higher on average—and have longer and more complex abstract texts, thus probably being even more informative than regular highly cited papers. Authors of these papers are close to the network core (high closeness); however, their rotating leadership is surprisingly lower than the one of authors of highly cited papers published in top journals. These publications also rarely involve scholars of different social groups. It seems that successful papers published in low-ranked journals mostly benefit from focused network embeddedness of their authors. Being part of a closed group with few structural holes, and being close to the network core, seem much more important than bridging social ties.

Our work not only extends research on the forecasting of paper citations, but also contributes to the identification of new metrics derived from social network and semantic analysis. The study has several limitations and the results of our analysis do only give limited insights about causality—which should be examined in future research. Is it that well-connected authors will get more citations in the future, or is it that highly cited papers will lead to more centrality for authors? One would assume that both statements are true.

Compared to past studies (Abramo et al. 2019a, b; Bornmann et al. 2014; Bruns and Stern 2016; Levitt and Thelwall 2011; Stegehuis et al. 2015; Wang et al. 2013), we present a model that considers the combined effects of a high number of predictors, i.e. scientific paper features. Future research could use our model and predictors to examine citations dynamics in fields other than chemical engineering, or consider even more control variables, to account, for example, for the presence of sleeping beauties or for possible geographical biases (Wuestman et al. 2019). Working with different citation timeframes, could reveal new factors impacting paper success. Subcategories of articles could also be considered, distinguishing between research papers and reviews of literature (we have already excluded the other categories of documents). Moreover, it might be that open access papers are cited more, as they are more easily accessible than paywalled ones (Eysenbach 2006)—even if, nowadays, this effect is mitigated by many factors, such as the increased availability of pre-print versions of published papers<sup>4</sup> and the existence of (pirate) websites like Sci-Hub (Himmelstein et al. 2018). To the best of our knowledge, this is one of the first studies where sentiment analysis of scientific abstracts is carried out. Indeed, sentiment analysis of scientific papers is a new and interesting problem (Athar 2011; Athar and Teufel 2012). Scientific communication is usually fact-based, and more technical, than texts that can be mined from other sources (Athar 2011)—for example social media. In this sense, our research is partially exploratory and tries to see whether the sentiment metric conveys any useful information for the prediction of future citations. We calculated sentiment using the

---

<sup>4</sup> [https://arxiv.org/stats/monthly\\_submissions](https://arxiv.org/stats/monthly_submissions).



VADER lexicon (Hutto and Gilbert 2014), whereas future research could consider different, or context-specific, approaches.

## References

- Abramo, G., Cicero, T., & D'Angelo, C. A. (2011). Assessing the varying level of impact measurement accuracy as a function of the citation window length. *Journal of Informetrics*, 5(4), 659–667. <https://doi.org/10.1016/j.joi.2011.06.004>.
- Abramo, G., & D'Angelo, C. A. (2015). The relationship between the number of authors of a publication, its citations and the impact factor of the publishing journal: Evidence from Italy. *Journal of Informetrics*, 9(4), 746–761. <https://doi.org/10.1016/j.joi.2015.07.003>.
- Abramo, G., & D'Angelo, C. A. (2017). Does your surname affect the citability of your publications? *Journal of Informetrics*, 11(1), 121–127. <https://doi.org/10.1016/j.joi.2016.12.003>.
- Abramo, G., D'Angelo, C. A., & Cicero, T. (2012). What is the appropriate length of the publication period over which to assess research performance? *Scientometrics*, 93(3), 1005–1017. <https://doi.org/10.1007/s11192-012-0714-9>.
- Abramo, G., D'Angelo, C. A., & Di Costa, F. (2016). The effect of a country's name in the title of a publication on its visibility and citability. *Scientometrics*, 109(3), 1895–1909. <https://doi.org/10.1007/s11192-016-2120-1>.
- Abramo, G., D'Angelo, C. A., & Felici, G. (2019a). Predicting publication long-term impact through a combination of early citations and journal impact factor. *Journal of Informetrics*, 13(1), 32–49. <https://doi.org/10.1016/j.joi.2018.11.003>.
- Abramo, G., D'Angelo, C. A., & Reale, E. (2019b). Peer review versus bibliometrics: Which method better predicts the scholarly impact of publications? *Scientometrics*, 121(1), 537–554. <https://doi.org/10.1007/s11192-019-03184-y>.
- Allen, T. J., Gloor, P., Fronzetti Colladon, A., Woerner, S. L., & Raz, O. (2016). The power of reciprocal knowledge sharing relationships for startup success. *Journal of Small Business and Enterprise Development*, 23(3), 636–651. <https://doi.org/10.1108/JSBED-08-2015-0110>.
- Antonacci, G., Fronzetti Colladon, A., Stefanini, A., & Gloor, P. (2017). It is rotating leaders who build the swarm: Social network determinants of growth for healthcare virtual communities of practice. *Journal of Knowledge Management*, 21(5), 1218–1239. <https://doi.org/10.1108/JKM-11-2016-0504>.
- Athar, A. (2011). Sentiment analysis of citations using sentence structure-based features. In *ACL HLT 2011: 49th annual meeting of the association for computational linguistics: Human language technologies, proceedings of student session* (pp. 81–87). Association for Computational Linguistics.
- Athar, A. (2014). *Sentiment analysis of scientific citations* (No. UCAM-CL-TR-856). University of Cambridge, Computer Laboratory <https://www.cl.cam.ac.uk/techreports/UCAM-CL-TR-856.pdf>. Accessed March 19, 2020.<Hyperlink></Hyperlink>
- Athar, A., & Teufel, S. (2012). Context-enhanced citation sentiment detection. In *NAACL HLT 2012–2012 conference of the North American chapter of the association for computational linguistics: Human language technologies, proceedings of the conference* (pp. 597–601). Association for Computational Linguistics.
- Barabási, A. L., Jeong, H., Néda, Z., Ravasz, E., Schubert, A., & Vicsek, T. (2002). Evolution of the social network of scientific collaborations. *Physica A: Statistical Mechanics and its Applications*, 311(3–4), 590–614. [https://doi.org/10.1016/S0378-4371\(02\)00736-7](https://doi.org/10.1016/S0378-4371(02)00736-7).
- Beaver, D. de B. (2004). Does collaborative research have greater epistemic authority? *Scientometrics*, 60(3), 399–408. <https://doi.org/10.1023/B:SCIE.0000034382.85360.c0>.
- Bennett, L. M., & Gadlin, H. (2012). Collaboration and team science: from theory to practice. *Journal of investigative medicine : the official publication of the American Federation for Clinical Research*, 60(5), 768–775. <https://doi.org/10.2310/JIM.0b013e318250871d>.
- Borgatti, S. P., Everett, M. G., & Johnson, J. C. (2013). *Analyzing social networks*. New York, NY: SAGE Publications.
- Bornmann, L., & Daniel, H. (2008). What do citation counts measure? A review of studies on citing behavior. *Journal of Documentation*, 64(1), 45–80. <https://doi.org/10.1108/00220410810844150>.
- Bornmann, L., Leydesdorff, L., & Wang, J. (2014). How to improve the prediction based on citation impact percentiles for years shortly after the publication date? *Journal of Informetrics*, 8(1), 175–180. <https://doi.org/10.1016/j.joi.2013.11.005>.
- Bozeman, B., & Corley, E. (2004). Scientists' collaboration strategies: Implications for scientific and technical human capital. *Research Policy*, 33(4), 599–616. <https://doi.org/10.1016/j.respol.2004.01.008>.

- Bozeman, B., Dietz, J., & Gaughan, M. (2001). Scientific and technical human capital: An alternative model for research evaluation. *International Journal of Technology Management*, 22(8), 716–740. <https://doi.org/10.1504/IJTM.2001.002988>.
- Breiman, L. (2001). Random forests. *Machine Learning*, 45(1), 5–32. <https://doi.org/10.1023/A:1010933404324>.
- Brönnimann, L. (2014). *Analyse der Verbreitung von Innovationen in sozialen Netzwerken*. [www.twittpolitiker.ch/documents/Master\\_Thesis\\_Lucas\\_Broennimann.pdf](http://www.twittpolitiker.ch/documents/Master_Thesis_Lucas_Broennimann.pdf). Accessed March 19, 2020.
- Bruns, S. B., & Stern, D. I. (2016). Research assessment using early citation information. *Scientometrics*, 108(2), 917–935. <https://doi.org/10.1007/s11192-016-1979-1>.
- Brunswick, S., Matei, S. A., Zentner, M., Zentner, L., & Klimeck, G. (2017). Creating impact in the digital space: digital practice dependency in communities of digital scientific innovations. *Scientometrics*, 110(1), 417–442. <https://doi.org/10.1007/s11192-016-2106-z>.
- Burt, R. S. (1995). *Structural holes: The social structure of competition*. Cambridge, MA: Harvard University Press.
- Burt, R. S. (2004). Structural holes and good ideas. *American Journal of Sociology*, 110(2), 349–399.
- Chen, T., & Guestrin, C. (2016). XGBoost: Reliable Large-scale Tree Boosting System. In *Proceedings of the 22nd ACM SIGKDD international conference on knowledge discovery and data mining* (pp. 785–794). San Francisco, CA: ACM New York. <https://doi.org/10.1145/2939672.2939785>
- Davis, J. P., & Eisenhardt, K. M. (2011). Rotating leadership and collaborative innovation: Recombination processes in symbiotic relationships. *Administrative Science Quarterly*, 56(2), 159–201. <https://doi.org/10.1177/0001839211428131>.
- Defazio, D., Lockett, A., & Wright, M. (2009). Funding incentives, collaborative dynamics and scientific productivity: Evidence from the EU framework program. *Research Policy*, 38(2), 293–305. <https://doi.org/10.1016/j.respol.2008.11.008>.
- Dietz, J. S. (2000). Building a social capital model of research development: the case of the Experimental Program to Stimulate Competitive Research. *Science and Public Policy*, 27(2), 137–145. <https://doi.org/10.3152/147154300781782093>.
- Dubitzky, W., Granzow, M., & Berrar, D. (2007). *Fundamentals of data mining in genomics and proteomics*. *Fundamentals of Data Mining in Genomics and Proteomics*. New York, NY: Springer Science + Business Media. <https://doi.org/10.1007/978-0-387-47509-7>
- Eysenbach, G. (2006). Citation advantage of open access articles. *PLoS Biology*, 4(5), e157. <https://doi.org/10.1371/journal.pbio.0040157>.
- Falahati, M. R., Goltaji, M., & Parto, P. (2015). The impact of title length and punctuation marks on article citations. *Annals of Library & Information Studies*, 62, 126–132.
- Franceschet, M., & Costantini, A. (2010). The effect of scholar collaboration on impact and quality of academic papers. *Journal of Informetrics*, 4(4), 540–553. <https://doi.org/10.1016/j.joi.2010.06.003>.
- Freeling, B., Doubleday, Z. A., & Connell, S. D. (2019). Opinion: How can we boost the impact of publications? Try better writing. *Proceedings of the National Academy of Sciences*, 116(2), 341–343. <https://doi.org/10.1073/pnas.1819937116>.
- Freeman, L. C. (1979). Centrality in social networks conceptual clarification. *Social networks*, 1, 215–239.
- Fronzetti Colladon, A., & Vagaggini, F. (2017). Robustness and stability of enterprise intranet social networks: The impact of moderators. *Information Processing & Management*, 53(6), 1287–1298. <https://doi.org/10.1016/j.ipm.2017.07.001>.
- Gloor, P., Fronzetti Colladon, A., Giacomelli, G., Saran, T., & Grippa, F. (2017a). The impact of virtual mirroring on customer satisfaction. *Journal of Business Research*, 75, 67–76. <https://doi.org/10.1016/j.jbusres.2017.02.010>.
- Gloor, P., Fronzetti Colladon, A., Grippa, F., & Giacomelli, G. (2017b). Forecasting managerial turnover through e-mail based social network analysis. *Computers in Human Behavior*, 71, 343–352. <https://doi.org/10.1016/j.chb.2017.02.017>.
- Guan, J., Yan, Y., & Zhang, J. J. (2017). The impact of collaboration and knowledge networks on citations. *Journal of Informetrics*, 11(2), 407–422. <https://doi.org/10.1016/j.joi.2017.02.007>.
- Guerrero-Bote, V. P., & Moya-Aneón, F. (2012). A further step forward in measuring journals' scientific prestige: The SJR2 indicator. *Journal of Informetrics*, 6(4), 674–688. <https://doi.org/10.1016/j.joi.2012.07.001>.
- Habibzadeh, F., & Yadollahie, M. (2010). Are shorter article titles more attractive for citations? Cross-sectional study of 22 scientific journals. *Croatian Medical Journal*, 51(2), 165–170. <https://doi.org/10.3325/cmj.2010.51.165>.
- Haggan, M. (2004). Research paper titles in literature, linguistics and science: Dimensions of attraction. *Journal of Pragmatics*, 36(2), 293–317. [https://doi.org/10.1016/S0378-2166\(03\)00090-0](https://doi.org/10.1016/S0378-2166(03)00090-0).

- Himmelstein, D. S., Romero, A. R., Levernier, J. G., Munro, T. A., McLaughlin, S. R., Greshake Tzovaras, B., & Greene, C. S. (2018). Sci-Hub provides access to nearly all scholarly literature. *eLife*, 7, 1–48. <https://doi.org/10.7554/eLife.32822>
- Hoekman, J., Frenken, K., & Tijssen, R. J. W. (2010). Research collaboration at a distance: Changing spatial patterns of scientific collaboration within Europe. *Research Policy*, 39(5), 662–673. <https://doi.org/10.1016/j.respol.2010.01.012>
- Huang, W. (2015). Do ABCs get more citations than XYZs? *Economic Inquiry*, 53(1), 773–789. <https://doi.org/10.1111/ecin.12125>
- Hutto, C. J., & Gilbert, E. (2014). VADER: A Parsimonious Rule-based Model for Sentiment Analysis of Social Media Text. In *Proceedings of the eighth international AAAI conference on weblogs and social media* (pp. 216–225). Ann Arbor, Michigan, USA: AAAI Press.
- Jacques, T. S., & Sebire, N. J. (2010). The impact of article titles on citation hits: An analysis of general and specialist medical journals. *JRSM Short Reports*, 1(1), 1–5. <https://doi.org/10.1258/shorts.2009.100020>
- Jamali, H. R., & Nikzad, M. (2011). Article title type and its relation with the number of downloads and citations. *Scientometrics*, 88(2), 653–661. <https://doi.org/10.1007/s11192-011-0412-z>
- Jha, Y., & Welch, E. W. (2010). Relational mechanisms governing multifaceted collaborative behavior of academic scientists in six fields of science and engineering. *Research Policy*, 39(9), 1174–1184. <https://doi.org/10.1016/j.respol.2010.06.003>
- Justeson, J. S., & Katz, S. M. (1995). Technical terminology: some linguistic properties and an algorithm for identification in text. *Natural Language Engineering*, 1(01), 9–27.
- Karlsson, A., Hammarfelt, B., Steinhauer, H. J., Falkman, G., Olson, N., Nelhans, G., et al. (2015). Modeling uncertainty in bibliometrics and information retrieval: An information fusion approach. *Scientometrics*, 102(3), 2255–2274. <https://doi.org/10.1007/s11192-014-1481-6>
- Katz, J. S., & Martin, B. R. (1997). What is research collaboration? *Research Policy*, 26(1), 1–18. [https://doi.org/10.1016/S0048-7333\(96\)00917-1](https://doi.org/10.1016/S0048-7333(96)00917-1)
- Kawashima, H., & Tomizawa, H. (2015). Accuracy evaluation of Scopus Author ID based on the largest funding database in Japan. *Scientometrics*, 103(3), 1061–1071. <https://doi.org/10.1007/s11192-015-1580-z>
- Kidane, Y. H., & Gloor, P. (2007). Correlating temporal communication patterns of the Eclipse open source community with performance and creativity. *Computational and Mathematical Organization Theory*, 13(1), 17–27.
- Kim, J., & Diesner, J. (2015). Coauthorship networks: A directed network approach considering the order and number of coauthors. *Journal of the Association for Information Science and Technology*, 66(12), 2685–2696. <https://doi.org/10.1002/asi.23361>
- Knorr-Cetina, K. D. (1981). *The manufacture of knowledge: An essay on the constructivist and contextual nature of science*. Oxford, UK: Pergamon Press.
- Knorr-Cetina, K. D. (1991). Merton's sociology of science: The first and the last sociology of science? *Contemporary Sociology*, 20(4), 522–526. <https://doi.org/10.2307/2071782>
- Larivière, V., Gingras, Y., Sugimoto, C. R., & Tsou, A. (2015). Team size matters: Collaboration and scientific impact since 1900. *Journal of the Association for Information Science and Technology*, 66(7), 1323–1332. <https://doi.org/10.1002/asi.23266>
- Latour, B., & Woolgar, S. (1986). *Laboratory life: The construction of scientific facts*. Princeton, NJ: Princeton University Press. <https://doi.org/10.1017/CBO9781107415324.004>
- Laudel, G. (2002). What do we measure by co-authorships? *Research Evaluation*, 11(1), 3–15. <https://doi.org/10.3152/147154402781776961>
- Leimu, R., & Koricheva, J. (2005). What determines the citation frequency of ecological papers? *Trends in Ecology & Evolution*, 20(1), 28–32. <https://doi.org/10.1016/j.tree.2004.10.010>
- Letchford, A., Preis, T., & Moat, H. S. (2016). The advantage of simple paper abstracts. *Journal of Informetrics*, 10(1), 1–8. <https://doi.org/10.1016/j.joi.2015.11.001>
- Levitt, J. M., & Thelwall, M. (2011). A combined bibliometric indicator to predict article impact. *Information Processing & Management*, 47(2), 300–308. <https://doi.org/10.1016/j.ipm.2010.09.005>
- Li, E. Y., Liao, C. H., & Yen, H. R. (2013). Co-authorship networks and research impact: A social capital perspective. *Research Policy*, 42(9), 1515–1530. <https://doi.org/10.1016/j.respol.2013.06.012>
- Li, M., Wu, J., Wang, D., Zhou, T., Di, Z., & Fan, Y. (2007). Evolving model of weighted networks inspired by scientific collaboration networks. *Physica A: Statistical Mechanics and its Applications*, 375(1), 355–364. <https://doi.org/10.1016/j.physa.2006.08.023>
- Lundberg, J., Tomson, G., Lundkvist, I., Skar, J., & Brommels, M. (2006). Collaboration uncovered: Exploring the adequacy of measuring university-industry collaboration through co-authorship and funding. *Scientometrics*, 69(3), 575–589. <https://doi.org/10.1007/s11192-006-0170-5>

- Lundberg, S. M., Erion, G., Chen, H., DeGrave, A., Prutkin, J. M., Nair, B., et al. (2020). From local explanations to global understanding with explainable AI for trees. *Nature Machine Intelligence*, 2(1), 56–67. <https://doi.org/10.1038/s42256-019-0138-9>.
- Lundberg, S. M., & Lee, S. I. (2017). A Unified Approach to Interpreting Model Predictions. In *Proceedings of the 31st conference on neural information processing system* (pp. 1–10). Long Beach, CA.
- Luwel, M., & Moed, H. F. (1998). Publication delays in the science field and their relationship to the ageing of scientific literature. *Scientometrics*, 41(1–2), 29–40. <https://doi.org/10.1007/BF02457964>.
- Maglaughlin, K. L., & Sonnenwald, D. H. (2005). Factors that impact interdisciplinary natural science research collaboration in academia. In *Proceedings of the international society for scientometrics and informatics* (pp. 499–508).
- Malvern, D., Richards, B., Chipere, N., & Durán, P. (2004). *Lexical diversity and language development: Quantification and assessment*. London, UK: Palgrave Macmillan. <https://doi.org/10.1057/9780230511804>.
- Marx, W., & Bornmann, L. (2013). Journal impact factor: “the poor man’s citation analysis” and alternative approaches. *European Science Editing*, 39(3), 62–63.
- Matveeva, N., & Poldin, O. (2016). Citation of scholars in co-authorship network: Analysis of google scholar data. *Applied Econometrics*, 44, 100–118.
- Melin, G., & Persson, O. (1996). Studying research collaboration using co-authorships. *Scientometrics*, 36(3), 363–377. <https://doi.org/10.1007/BF02129600>.
- Melkers, J., & Kiopa, A. (2010). The social capital of global ties in science: The added value of international collaboration. *Review of Policy Research*, 27(4), 389–414. <https://doi.org/10.1111/1j.1541-1338.2010.00448.x>.
- Merton, R. K. (1957). Priorities in scientific discovery: A chapter in the sociology of science. *American Sociological Review*, 22(6), 635. <https://doi.org/10.2307/2089193>.
- Merton, R. K. (1968). The Matthew effect in science: The reward and communication systems of science are considered. *Science*, 159(3810), 56–63. <https://doi.org/10.1126/science.159.3810.56>.
- Milojević, S. (2014). Principles of scientific research team formation and evolution. *Proceedings of the National Academy of Sciences*, 111(11), 3984–3989. <https://doi.org/10.1073/pnas.1309723111>.
- Mingers, J., & Xu, F. (2010). The drivers of citations in management science journals. *European Journal of Operational Research*, 205(2), 422–430. <https://doi.org/10.1016/j.ejor.2009.12.008>.
- Moed, H. F., & van Leeuwen, T. N. (1996). Impact factors can mislead. *Nature*, 381(6579), 186–186. <https://doi.org/10.1038/381186a0>.
- Murray, C. (2005). Social capital and cooperation in Central and Eastern Europe: A theoretical perspective. *EconPapers*, 9(9), 25.
- Nahapiet, J., & Ghoshal, S. (1998). Social capital, intellectual capital, and the organizational advantage. *Academy of management review*, 23(2), 242–266.
- Nair, L. B., & Gibbert, M. (2016). What makes a ‘good’ title and (how) does it matter for citations? A review and general model of article title attributes in management science. *Scientometrics*, 107(3), 1331–1359. <https://doi.org/10.1007/s11192-016-1937-y>.
- Newman, H., & Joyner, D. (2018). Sentiment analysis of student evaluations of teaching. In C. P. Rosé, R. Martínez-Maldonado, H. U. Hoppe, R. Luckin, M. Mavrikis, K. Porayska-Pomsta, et al. (Eds.), *Artificial Intelligence in Education. AIED 2018* (pp. 246–250). Cham, Switzerland: Springer. [https://doi.org/10.1007/978-3-319-93846-2\\_45](https://doi.org/10.1007/978-3-319-93846-2_45)
- Olson, G., & Olson, J. (2000). Distance matters. *Human-Computer Interaction*, 15(2), 139–178. [https://doi.org/10.1207/S15327051HCI1523\\_4](https://doi.org/10.1207/S15327051HCI1523_4).
- Ong, D., Chan, H. F., Torgler, B., & Yang, Y. (2018). Collaboration incentives: Endogenous selection into single and coauthorships by surname initial in economics and management. *Journal of Economic Behavior & Organization*, 147, 41–57. <https://doi.org/10.1016/j.jebo.2018.01.001>.
- Paiva, C., Lima, J., & Paiva, B. (2012). Articles with short titles describing the results are cited more often. *Clinics*, 67(5), 509–513. [https://doi.org/10.6061/clinics/2012\(05\)17](https://doi.org/10.6061/clinics/2012(05)17).
- Perc, M. (2010). Growth and structure of Slovenia’s scientific collaboration network. *Journal of Informetrics*, 4(4), 475–482. <https://doi.org/10.1016/j.joi.2010.04.003>.
- Perkins, J. (2014). *Python 3 Text Processing With NLTK 3 Cookbook. Python 3 Text Processing With NLTK 3 Cookbook*. Birmingham, UK: Packt Publishing.
- Persson, O., Melin, G., Danell, R., & Kaloudis, A. (1997). Research collaboration at Nordic universities. *Scientometrics*, 39(2), 209–223. <https://doi.org/10.1007/bf02457449>.
- Petersen, A. M., Pan, R. K., Pammolli, F., & Fortunato, S. (2019). Methods to account for citation inflation in research evaluation. *Research Policy*, 48(7), 1855–1865. <https://doi.org/10.1016/j.respo.1.2019.04.009>.

- Plavén-Sigra, P., Matheson, G. J., Schiffler, B. C., & Thompson, W. H. (2017). The readability of scientific texts is decreasing over time. *Elife*, 6, 1–14. <https://doi.org/10.7554/eLife.27725>.
- Rostami, F., Mohammadpoorasl, A., & Hajizadeh, M. (2014). The effect of characteristics of title on citation rates of articles. *Scientometrics*, 98(3), 2007–2010. <https://doi.org/10.1007/s11192-013-1118-1>.
- Sekara, V., Deville, P., Ahnert, S. E., Barabási, A.-L., Sinatra, R., & Lehmann, S. (2018). The chaperone effect in scientific publishing. *Proceedings of the National Academy of Sciences*, 115(50), 12603–12607. <https://doi.org/10.1073/pnas.1800471115>.
- Shevlin, M., & Davies, M. N. O. (1997). Alphabetical listing and citation rates. *Nature*, 388(6637), 14–14. <https://doi.org/10.1038/40253>.
- Spearman, C. (1904). The proof and measurement of association between two things. *The American Journal of Psychology*, 15(1), 72–101. <https://doi.org/10.2307/1412159>.
- Stegehuis, C., Litvak, N., & Waltman, L. (2015). Predicting the long-term citation impact of recent publications. *Journal of Informetrics*, 9(3), 642–657. <https://doi.org/10.1016/j.joi.2015.06.005>.
- Subotic, S., & Mukherjee, B. (2014). Short and amusing: The relationship between title characteristics, downloads, and citations in psychology articles. *Journal of Information Science*, 40(1), 115–124. <https://doi.org/10.1177/0165551513511393>.
- Trivedi, P. K. (1993). An analysis of publication lags in econometrics. *Journal of Applied Econometrics*, 8(1), 93–100. <https://doi.org/10.1002/jae.3950080108>.
- Uddin, S., Hossain, L., Abbasi, A., & Rasmussen, K. (2012). Trend and efficiency analysis of co-authorship network. *Scientometrics*, 90(2), 687–699. <https://doi.org/10.1007/s11192-011-0511-x>.
- Uddin, S., & Khan, A. (2016). The impact of author-selected keywords on citation counts. *Journal of Informetrics*, 10(4), 1166–1177. <https://doi.org/10.1016/j.joi.2016.10.004>.
- van Raan, A. F. J. (2004). Sleeping beauties in science. *Scientometrics*, 59(3), 467–472. <https://doi.org/10.1023/B:SCIE.0000018543.82441.f1>.
- van Wesel, M., Wyatt, S., & ten Haaf, J. (2014). What a difference a colon makes: How superficial factors influence subsequent citation. *Scientometrics*, 98(3), 1601–1615. <https://doi.org/10.1007/s11192-013-1154-x>.
- VTR. (2006). *Italian Triennial Research Evaluation. VTR 2001–2003. Risultati delle valutazioni dei Panel di Area*. <https://vtr2006.cineca.it>
- Wagner, C. S., Park, H. W., & Leydesdorff, L. (2015). The continuing growth of global cooperation networks in research: A conundrum for national governments. *PLoS ONE*, 10(7), 1–15. <https://doi.org/10.1371/journal.pone.0131816>.
- Wallace, M. L., Larivière, V., & Gingras, Y. (2012). A Small world of citations? The influence of collaboration networks on citation practices. *PLoS ONE*, 7(3), 1–10. <https://doi.org/10.1371/journal.pone.0033339>.
- Waltman, L., & van Eck, N. J. (2015). Field-normalized citation impact indicators and the choice of an appropriate counting method. *Journal of Informetrics*, 9(4), 872–894. <https://doi.org/10.1016/j.joi.2015.08.001>.
- Wang, D., Song, C., & Barabási, A.-L. (2013). Quantifying long-term scientific impact. *Science*, 342(6154), 127–132. <https://doi.org/10.1126/science.1237825>.
- Wang, J. (2014). Unpacking the Matthew effect in citations. *Journal of Informetrics*, 8(2), 329–339. <https://doi.org/10.1016/j.joi.2014.01.006>.
- Wasserman, S., & Faust, K. (1994). *Social network analysis: Methods and applications*. New York, NY: Cambridge University Press. <https://doi.org/10.1525/ae.1997.24.1.219>.
- Weinberger, C. J., Evans, J. A., & Allesina, S. (2015). Ten simple (empirical) rules for writing science. *PLOS Computational Biology*, 11(4), 1–6. <https://doi.org/10.1371/journal.pcbi.1004205>.
- Weingart, P. (2005). Impact of bibliometrics upon the science system: Inadvertent consequences? *Scientometrics*, 62(1), 117–131. <https://doi.org/10.1007/s11192-005-0007-7>.
- Wuestman, M. L., Hoekman, J., & Frenken, K. (2019). The geography of scientific citations. *Research Policy*, 48(7), 1771–1780. <https://doi.org/10.1016/j.respol.2019.04.004>.