






See further upon the giants: Quantifying intellectual lineage in science

Woo Seong Jo^{1,2,3} , Lu Liu^{1,2,3,4} , and Dashun Wang^{1,2,3,5} 

¹Center for Science of Science & Innovation, Northwestern University, Evanston, IL, USA

²Northwestern Institute on Complex Systems, Northwestern University, Evanston, IL, USA

³Kellogg School of Management, Northwestern University, Evanston, IL, USA

⁴College of Information Sciences and Technology, Pennsylvania State University, University Park, PA, USA

⁵McCormick School of Engineering, Northwestern University, Evanston, IL, USA

an open access  journal



Citation: Jo, W. S., Liu, L., & Wang, D. (2022). See further upon the giants: Quantifying intellectual lineage in science. *Quantitative Science Studies*, 3(2), 319–330. https://doi.org/10.1162/qss_a_00186

DOI: https://doi.org/10.1162/qss_a_00186

Peer Review: https://publons.com/publon/10.1162/qss_a_00186

Supporting Information: https://doi.org/10.1162/qss_a_00186

Received: 21 April 2021
Accepted: 7 February 2022

Corresponding Authors:
Lu Liu
luliu2131@gmail.com
Dashun Wang
dashun.wang@northwestern.edu

Handling Editor:
Staša Milojević

Copyright: © 2022 Woo Seong Jo, Lu Liu, and Dashun Wang. Published under a Creative Commons Attribution 4.0 International (CC BY 4.0) license.



Keywords: citation network, cocitation, science of science, scientific impact

ABSTRACT

Newton’s centuries-old wisdom of standing on the shoulders of giants raises a crucial yet underexplored question: Out of all the prior works cited by a discovery, which one is its giant? Here, we develop a discipline-independent method to identify the giant for any individual paper, allowing us to better understand the role and characteristics of giants in science. We find that across disciplines, about 95% of papers appear to stand on the shoulders of giants, yet the weight of scientific progress rests on relatively few shoulders. Defining a new measure of giant index, we find that, while papers with high citations are more likely to be giants, for papers with the same citations, their giant index sharply predicts a paper’s future impact and prize-winning probabilities. Giants tend to originate from both small and large teams, being either highly disruptive or highly developmental. Papers that did not have a giant tend to do poorly on average, yet interestingly, if such papers later became a giant for other papers, they tend to be home-run papers that are highly disruptive to science. Given the crucial importance of citation-based measures in science, the developed concept of giants may offer a useful dimension in assessing scientific impact that goes beyond sheer citation counts.

1. INTRODUCTION

“If I have seen further, it is by standing on the shoulders of giants.” Isaac Newton’s famous 1675 letter to Robert Hooke highlights a fundamental feature of science: its cumulative nature. Indeed, new insights and discoveries rarely emerge in isolation; instead, they build on prior scientific work. Because scientists throughout the ages and across disciplines all acknowledge ideas that inspired their research, we have an opportunity to explore citation relationships to better understand how new research makes use of influential work (Clauset, Larremore, & Sinatra, 2017; Fortunato, Bergstrom et al., 2018; Garfield, 2006; Price, 1965; Radicchi, Fortunato, & Castellano, 2008; Redner, 2005; Waltman, 2016; Wang & Barabási, 2021; Wang, Song, & Barabási, 2013). Here, we hone in on a specific, underexplored question: Is there a way to estimate, given any paper, which reference is the “giant” whose shoulders the new research stands upon?

Here we take advantage of the citation relationships between papers and develop a network-based method that aims to estimate the relative intellectual significance of each reference to a paper, allowing us to estimate the potential giant for a paper. Our leading

hypothesis is that, while each paper has many references, their importance to the paper can be uneven, and the relationships among the listed references, revealed through the overall citation network, may help us understand the intellectual significance of each reference in the context of the overall scientific discourse.

Citations are essential to scientific communication, allowing scientists to condense knowledge, bolster the strength of their evidence, attribute prior ideas to appropriate sources, and more (Wang & Barabási, 2021). Partly because canonical papers tend to inspire follow-up research which builds on them, citation relationships have been widely used to quantify scientific impact (Bergstrom, West, & Wiseman, 2008; Cole & Cole, 1974; Garfield, 2006; Hirsch, 2005; King, 2004; Liu, Wang et al., 2018; Radicchi et al., 2008; Sinatra, Wang et al., 2016; Uzzi, Mukherjee et al., 2013; Waltman, 2016; Wang et al., 2013; Way, Morgan et al., 2019; Wu, Wang, & Evans, 2019; Wuchty, Jones, & Uzzi, 2007). At the same time, citations can be affected by myriad factors: Publication venue, year, and field of study are among many other reasons why authors cite a given paper (Aksnes, 2006; Moravcsik & Murugesan, 1975; Radicchi, 2012; Simkin & Roychowdhury, 2002), contributing to noise in evaluating and comparing their relative importance.

The need to quantify scientific impact by considering citations of different importance has inspired various methods to identify the key references. Some examined the citation context by analyzing the number of mentions and the sections and relative positions that a reference appears (Bornmann & Daniel, 2008; Boyack, van Eck et al., 2018; Ding, Liu et al., 2013; Ding, Zhang et al., 2014; Hu, Chen, & Liu, 2013; Jones & Hanney, 2016). Some combined features from the full text, citation counts, and abstract similarity, and trained a classifier to predict important references (Hassan, Akram, & Haddawy, 2017; Zhu, Turney et al., 2015). Some developed a local diffusion method on the citation network and ranked references by their diffusion score (Cui, Zeng et al., 2020). Other researchers conducted surveys for authors and asked them to identify references that shaped the research idea and influenced the research (Tahamtan & Bornmann, 2018; Zhu et al., 2015). These studies have also contributed to the development of new metrics, such as weighting citations by the number of mentions and the diffusion score on the citation network.

Here we study 33 million papers indexed by the Web of Science (WoS) between 1955 and 2014, and 962 million citations among them (Supplementary Material, Figure S1). To quantify the intellectual significance of each reference to a given paper, we first use cocitation relationships to establish a measure of proximity between references, which measures how many times two papers are cited together by other papers. Cocitation has been used for a variety of purposes from quantifying the topical relevance among papers (Chen, 2006; Liu & Chen, 2012; Small, 1973) to author credit allocation within a paper (Shen & Barabási, 2014) to evaluating impacts of authors (Ding, Yan et al., 2009) and suggesting relevant references (Sarol, Liu, & Schneider, 2018). To identify the “giant” for a given paper, we first take all the papers ever published up to the publication year of the focal paper, and identify all cocitation relationships (Figure 1A), with link weight indicating the number of times two papers are cocited. This cocitation network can act as a proxy for the overall knowledge space at the time (Figure 1B). Next, for the focal paper, we locate all its references to obtain the reference subnetwork embedded in the overall cocitation network, approximating the intellectual context in which the paper is placed (Figure 1C). Here we focus on papers that contain at least five references to ensure we have enough nodes for each reference subnetwork, resulting in 25 million papers in total.

To identify local significance of a reference within the subnetwork while taking into account its overall influence in the global network (Lü, Chen et al., 2016; Lü & Zhou,

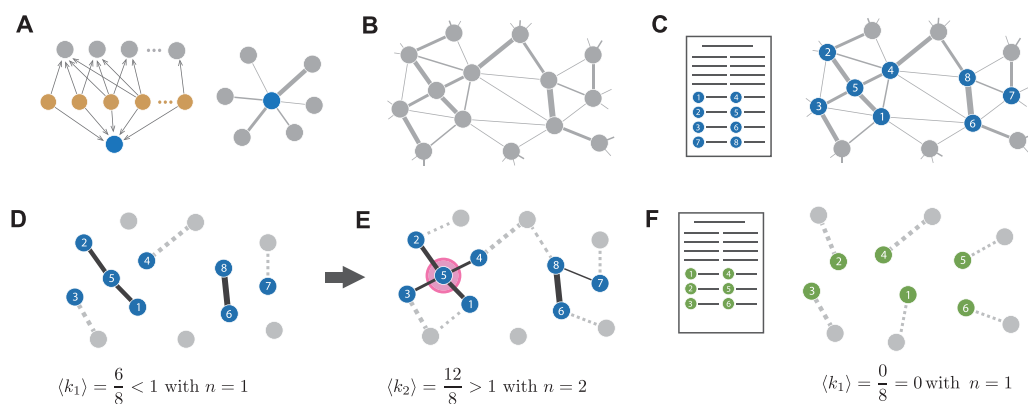


Figure 1. Identifying the giant paper. (A) For a given reference (blue), we find all other papers that are cocited with this reference (cocited papers in grey) with link weight indicating the number of cocitations. (B) We can follow the procedure in (A) to construct the entire cocitation network for all papers published up to a certain year, representing the overall knowledge space at the time. (C) To identify the giant for a given paper, we use the cocitation network in its publication year, and locate all its references (blue nodes) to pinpoint the reference subnetwork from the overall cocitation network. (D) We give each reference the same number of “votes” to link to their most relevant paper in the global cocitation network (both solid and dashed lines), but keep only the links formed within the reference subnetwork (solid line). We begin by adding one connection (a vote) for each reference, and monitor the connectivity within the subnetwork. (E) We increase the number of connections iteratively, and stop when the average degree of the reference subnetwork is greater than 1: the minimum requirement to have a cluster. We then pick the node with the largest degree as the giant paper (highlighted in purple). (F) If, however, the reference subnetwork remains completely isolated at $n = 1$, we stop our algorithm without identifying giants for such cases.

2011), we develop a new method by borrowing concepts from democratic voting and percolation theory. Our key insight is to give each reference the same number of “votes” to link to their most relevant paper in the global cocitation network, but keep only the links formed within the reference subnetwork (only counting the votes within the specific context). We hypothesize that the giant of a paper should appear in its reference list and have high topical relevance to the body of work, prompting us to focus on the votes in the reference subnetwork. At the same time, the giant paper may also be captured by collective recognition within a scientific community. To this end, our method recognizes the collective nature of the intellectual lineage by exploring rich information embedded in the global cocitation network. Specifically, we begin by adding one connection (a vote) for each reference, and monitor the connectivity within the subnetwork (Figure 1D). We increase the number of connections iteratively, and stop when the reference subnetwork starts to coalesce into a cluster, suggesting plausible knowledge structures beginning to emerge within this subnetwork. Resorting to percolation theory, we stop at the minimal n votes needed for the subnetwork to cross the percolation threshold for the corresponding random network (Newman, 2010, i.e., $\langle k_n \rangle > 1$), where references within the subnetwork have at least one connection on average (Figure 1E). We then select the node with the largest degree (k) as the giant paper (Supplementary Material, Figure S2). The idea is to approximate the local importance of a reference by counting the “votes” it received from other references. And the one with the most “votes” suggests its importance among other references. At the same time, the idea of a percolation threshold also affords us the possibility of not always identifying a giant for a paper. Indeed, if, however, the reference subnetwork remains completely isolated at $n = 1$ (Figure 1F), suggesting that the paper’s references all belong to distinctive parts of the overall knowledge space, we stop our algorithm at $n = 1$ without identifying giants for such cases.

One advantage of our method is to overcome mixed signals introduced by the skewed citation distribution (Barabási & Albert, 1999; Cui et al., 2020; Price, 1965; Radicchi et al., 2008). Indeed, because highly cited papers tend to dominate the subnetwork of references, directly

applying existing network-based methods to the cocitation network tends to favor papers with overall high citation counts, which may or may not be the specific giant to a paper. One can somewhat mitigate this issue by keeping, for example, the most essential links using network sparsification methods such as backbone extraction (Radicchi, Ramasco, & Fortunato, 2011; Serrano, Boguná, & Vespignani, 2009), yet it appears insufficient to overcome the dominance of highly cited papers. This highlights one of the many challenges of this task: To identify the giant for a specific paper, we need to balance between a reference's local importance to the specific context and its overall scientific impact. At the same time, prior work reveals that creative works may not necessarily build from existing literature (Tahamtan & Bornmann, 2018), and our method allows us to identify papers without giants, which is difficult to achieve using citation count or the backbone algorithm. Overall, we recognize that high-citation papers are important and often instrumental in the development of the field, but it does not guarantee that they are the most relevant for a specific paper in a field. This prompts us to take into account both the local relevance of a paper within the reference list and its global impact in the overall citation network, which may offer a complementary dimension of a paper's impact to its citation count. We also note that each of the steps in the method can be further refined with different stopping criteria and a different number of giants to pick for a paper. Here we choose this specification of the method for its simplicity and leave for future work to systematically study the different variants of the method, which we shall discuss in more detail in the discussion and limitation section of the paper.

2. RESULTS

We apply our method to millions of papers published over the past 60 years, allowing us to study the giants we identify and examine patterns of intellectual lineage in science. First, we find that about 95% of papers have a giant associated with them (Figure 2A). This overwhelming proportion of papers with identifiable giants offers quantitative evidence for the cumulative nature of science. We further find that, despite the exponential growth of science and rise of interdisciplinary research and collaborations, this prevalence of giants has held true over the

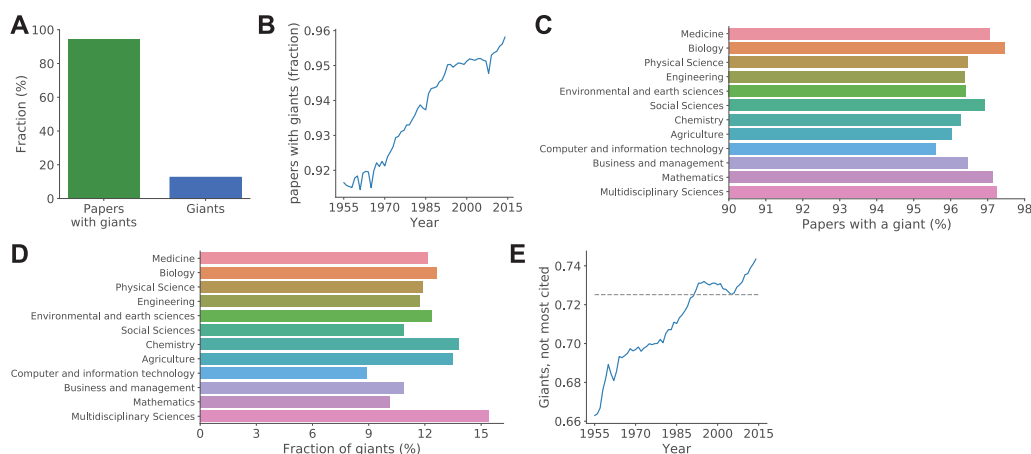


Figure 2. Prevalence of giants in science. (A) Out of the 25 million papers we studied, 95% of them have a giant (green). By contrast, only 12% of papers later become giants for other papers (blue). (B) The fraction of papers that have an associated giant increases over the years. (C) Breaking down the fraction of papers with giants by 12 different fields. (D) Breaking down the fraction of papers that later become a giant by 12 different fields. Papers in multidisciplinary sciences have the highest likelihood to become a giant. (E) We calculate the fraction of papers whose giant is not the most cited paper within its reference list, and show this fraction over time.

ages but has risen gradually over time (Figure 2B), growing from 91.6% in 1955 to 95.8% in 2014, which implies that research today is increasingly conducted on the shoulders of giants. Yet at the same time, the papers that serve as giants are remarkably concentrated, as only 12% of the papers we analyzed fell into this category (Figure 2A), consistent with what the literature suggests (Bornmann, de Moya Anegón, & Leydesdorff, 2010; Cole & Cole, 1972). We further find that these results are remarkably consistent across the disciplines we studied (Figures 2C–D). Hence, despite the inherent differences in norm and culture across disciplines, the shoulders of giants appear universally appreciated. Interestingly, of the 12 different scientific fields, multidisciplinary science has the highest fraction of giant papers (15.4%, Figure 2D), suggesting a premium in bringing together diverse approaches. Overall, these results suggest that while the vast majority of papers stand on the shoulders of giants, the weight of scientific progress rests on only a few shoulders.

Is the identified giant for a given paper the most cited within its references? For each paper that has a giant, we calculate the citation counts of all its references when the paper was published. We find that the vast majority of giant papers are *not* the most cited within the reference list (72.5%, dashed line in Figure 2E), and this fraction has mostly been trending upward over the past 60 years (Figure 2E). For instance, in 1955, 44% of giants were also the most cited references, but that number decreased to only 26% in 2014. These results suggest being the most cited paper does not guarantee being the giant, implying that our measures of giant may capture complementary dimensions of impact to citation counts, which prompt us to further examine the overall characteristics of a giant paper.

To understand what kinds of papers tend to become a giant, we introduce a new index: the giant index (G), which calculates the number of times a paper is a giant for other papers. We then compare a paper's G with its citation count (C). We first find that the likelihood of being a giant posts a highly nonlinear relationship with its citation counts. Indeed, among papers with exceptional impact ($C > 1,000$), virtually all of them have a positive giant index. The probability for a paper to be a giant ($P(G > 0)$) increases with citation counts but undergoes a sharp crossover for papers with respectable but more moderate citations (Figure 3A). To ensure that the observed trend is not affected by self-citations, we remove self-citations and repeat our analysis, arriving at the same conclusion (Supplementary Material, Figure S3). These results suggest that not all papers can be a giant, and a paper's potential to become a giant may be related to its citation impact, prompting us to further examine the correlation between a paper's giant index and its citations (Figure 3B). We find that for papers with dozens to hundreds of citations, their giant index and citations follow a superlinear relationship, suggesting that an increase in a paper's citations is associated with an increasing return in the rate at which a paper becomes a giant. Yet for papers on the right-hand tail of citation impacts, their giant index roughly follows a linear relationship with their citations, consistent with the trend shown in Figure 3A, suggesting that papers with exceptional impact are disproportionately more likely to be the shoulders that carry scientific progress in their field.

At the same time, the overall correlation between the giant index and citations (Figure 3B) also masks heterogeneous relationships between the two. Indeed, as both $P(G)$ and $P(C)$ follow a fat-tailed distribution (Supplementary Material, Figure S4) and by design $G \leq C$, we calculate the conditional probability $P(G|C)$ (Figure 3C). We find that, as C increases, the distribution systematically shifts to the right, consistent with the correlations observed in Figures 3A–B. Yet, for papers with the same level of citations, their giant indexes are still characterized by a high degree of heterogeneity, suggesting that a high citation count does not necessarily guarantee a high giant index. Next, we show that this discrepancy between the giant index and citation counts offers signals for a paper's future impact.

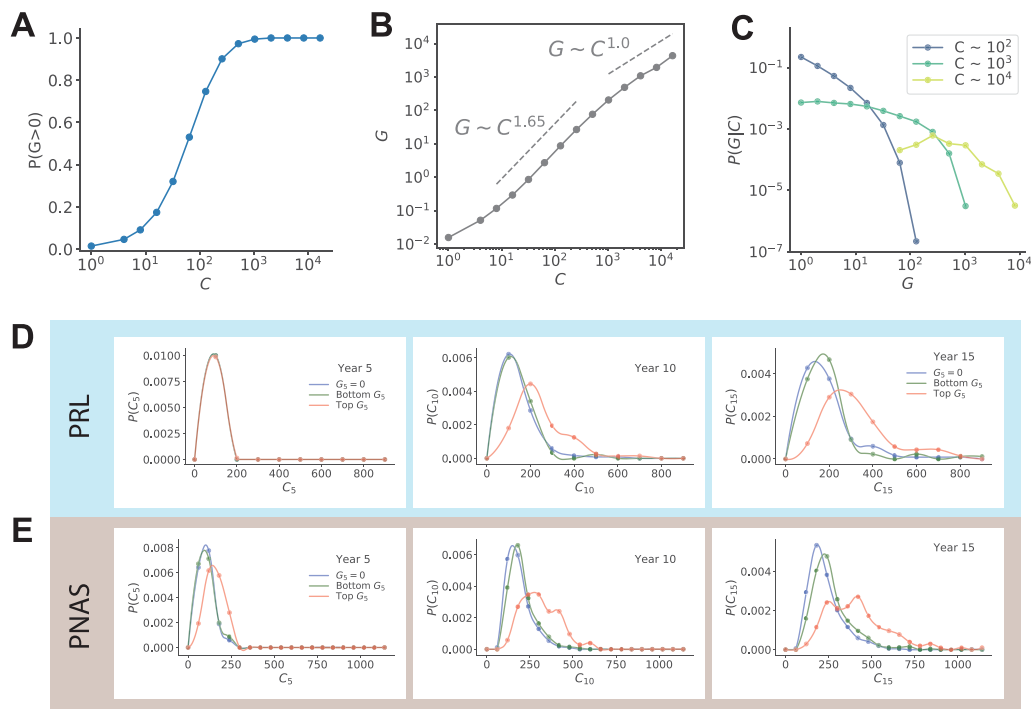


Figure 3. The giant index. (A) The fraction of papers with a nonzero giant index as a function of citations. (B) The giant index G as a function of C . The two quantities follow a superlinear relationship when $C < 1,000$ but exhibit a linear relationship for high C regime ($C > 1,000$). (C) Conditional distribution $P(G|C)$ shows that papers are characterized by a heterogeneous giant index even though they have the same level of citation impacts. (D) We pick papers published in *PRL* between 1990 and 2000 that have between 100 and 200 citations in 5 years ($100 \leq C_5 \leq 200$). Depending on their giant index at year 5, G_5 , these papers exhibit different levels of future impacts. (E) The same analysis as (D) but for a group of papers published in *PNAS* ($100 \leq C_5 \leq 200$).

We select papers published in *Physical Review Letters (PRL)* between 1990 and 2000 but which have a similar level of citation impact after 5 years of publication (within the range of $100 \leq C_5 \leq 200$ by Year 5) (Figure 3D). We divide these papers into three groups based on their giant index at year 5 (G_5): high G group (top 10% in giant score G_5 , $\langle G_5 \rangle \approx 31.3$), low G group (bottom 10% in G_5 , $\langle G_5 \rangle \approx 1.0$), and nongiant group ($G_5 = 0$). We then trace the citations of these papers over the next 10 years. We find that at year 5, the three groups follow a similar citation distribution, by construction. Yet with time, the high G group clearly stood out from the pack, collecting citations at a much faster rate than the other two groups. Interestingly, there is a statistically significant difference between the $G \approx 1$ and $G = 0$ groups by year 15, suggesting that papers with a small giant index are likely to have higher future impact. We repeat the same analysis for a multidisciplinary journal, selecting papers published in *Proceedings of the National Academy of Sciences (PNAS)* ($100 \leq C_5 \leq 200$, Figure 3E), finding again the same patterns. Taken together, Figure 3 shows that while a paper's giant index and citation counts are overall correlated, papers with the same citations can have vastly different giant indexes and that a difference in giant index appears to substantially reveal a paper's potential for future impact. These results suggest that the giant index offers additional information on a paper's role in science that goes beyond its citation counts. Together, they reflect the idea that not all citations are the same, and those that frequently lend their shoulders to others tend to distinguish themselves from those that do not.

To understand potential forces that might facilitate the production of giants, we further examine the organization of scientific activity, probing the role of teams in shouldering

scientific progress. Indeed, research shows that small and large teams are differentially positioned for innovation (Wu et al., 2019): Large teams tend to excel at furthering existing ideas and design, whereas small teams tend to disrupt current ways of thinking with new ideas and opportunities. This distinction prompts us to measure whether the likelihood of producing giants varies by team size. To control the effect of field and time, we normalize citations and giant indexes with the average value for papers published in the same field and year (Radicchi et al., 2008), computing the normalized citation $C/\langle C \rangle_{f,y}$ and normalized giant index $G/\langle G \rangle_{f,y}$. We first find that, across different fields, the normalized citation increases with team size M (Figure 4A), confirming previous studies showing the citation premium conferred upon large teams (Wu et al., 2019; Wuchty et al., 2007). We then repeat this analysis for giant papers ($G > 0$), finding that the normalized giant index posts a U-shaped curve (Figure 4B), and this nonlinear relationship holds the same across different fields (shaded curves in Figure 4B). These results suggest that while works by large teams tend to garner higher citations, giants that frequently lend their shoulders tend to originate from both small and large teams.

The relationship between team size and research outcomes prompted us to ask if the observed nonlinear relationship between giant index and team size is related to the character of work that teams of different size produce. Here, we measure the relationship between a paper's giant index and its disruption percentile DP . We calculate the disruption score following prior work (Funk & Owen-Smith, 2017; Wu et al., 2019). For each paper, we calculate the number of subsequent papers citing the paper but not its references (n_i), the number of subsequent papers citing both the paper and its references (n_j), and the number of subsequent papers citing the references but not the paper (n_k). The raw disruption score (D) is defined as $D = \frac{n_i - n_j}{n_i + n_j + n_k}$. We further normalized the raw score to its percentile, the




Figure 4. Teams, disruption, and giant index. (A) We normalize citation by field and year $C/\langle C \rangle_{f,y}$ and plot it as a function of team size M . (B) We normalize the giant index by field and year, finding that it follows a U-shaped curve with team size. (C) The normalized giant index as a function of disruption percentile DP , showing that either highly developmental ($DP \leq 20$) or highly disruptive ($DP \geq 80$) papers tend to a high giant index. (D) Distributions of citations for papers with or without a giant. (E) We categorize papers without giants into two groups based on their giant index: $G = 0$ vs. $G > 0$. The two groups show clearly different relationships with the citation impacts of these papers. (F) The distribution of disruption percentile DP for papers with (red circle) or without (yellow squares) a giant, showing that papers without a giant are disproportionately more likely to be a highly disruptive paper. (Inset) The fraction of papers without a giant as a function of team size, showing that such papers are more likely to be produced by small teams.

disruption percentile (DP). DP measures the relative ranking among all papers published in the same year, with 100 indicating the most disruptive and 0 as the most developmental. We find that giant index sharply increases for both highly developmental ($DP < 20$) and highly disruptive ($DP > 80$) work (Figure 4C). These results illustrate the divergent characters of giants in science, being either highly disruptive or highly developmental. Both types of giants shoulder scientific progress, but they move science forward in different ways.

One advantage of our method is to allow the flexibility of not identifying a giant for a paper. Indeed, as discussed in Figure 1, if all the references of a paper are initially isolated in the reference subnetwork, it suggests that the paper draws upon disparate rather than established knowledge clusters, and for these papers our method proceeds without assigning a giant (i.e., papers without giants). Of the 25 million papers we studied, 5.2% of them fall into this category. Figure 4D plots the citation distributions for papers with and without a giant, respectively. We find that papers without a giant tend to have overall fewer citations than those that stand upon a giant's shoulder, further suggesting the importance of giants in the production of knowledge. Yet at the same time, Figure 4D also reveals an intriguing observation: Even among papers without a giant, their citations are characterized by a high degree of heterogeneity, indicating that some papers do eventually garner high impacts, albeit uncommonly. Indeed, we find that, among the papers without giants, 10% of them go on to become giants for others. We separate the papers without giants into two groups, calculating the relative probability of observing a $G > 0$ vs. $G = 0$ paper as a function of citations, and find that the two groups follow clearly divergent patterns (Figure 4E): If a paper neither became a giant for others ($G = 0$) nor stood on one's shoulder when published, its impact is mostly concentrated within the low-citation region, and the probability for such papers to garner higher citations diminishes rather rapidly. By contrast, those that did not have a specific shoulder to rely upon but later became a giant for others are systematically overrepresented in the high-citation region, as their relative abundance rises precipitously with citations. These results paint a highly polarized view for papers without a giant. On the one hand, such papers have a rather limited impact on average, suggesting that skipping the shoulders substantially limits a paper's ability to "see further." Yet on the other hand, perhaps counterintuitively, papers without a giant may also become home-run papers with right-tail citation impact. One conjecture is that papers that emerge from a seeming vacuum that lies between knowledge clusters may reorient the existing knowledge in a way that offers new ideas and opportunities. To test this conjecture, we measure the disruption percentile DP , and find that papers without giants are sharply overrepresented in the highly disruptive region (Figure 4F), and much more likely to be produced by small teams (Figure 4F inset).

Last, as a further validation, we show that the giant index offers a simple yet additional early signal for the Noble prize-winning papers. As the most prestigious prize in science, the Nobel Prize recognizes some of the most crucial scientific breakthroughs. There have been constant attempts to identify Nobel-Prize-winning discoveries based on citation counts (Garfield & Malin, 1968; Revesz, 2015; Zakhlebin & Horvát, 2017). Despite their occasional success, citations appear to be a noisy signal for the Nobel, for a simple reason: While Nobel-Prize-winning papers all tend to be highly cited, having high citations does not guarantee a Nobel. This raises an intriguing question: Could the giant index offer additional information in differentiating the prize-winning papers beyond citation counts? To answer this question, we identified 370 Nobel-Prize-winning papers for physics, chemistry, and medicine awarded between 1955 and 2014, and compared their giant index with papers published in the same year and field with similar citations (the comparison group). We find that, by construction, the citation distribution $P(C)$ is largely indistinguishable between the two groups of papers

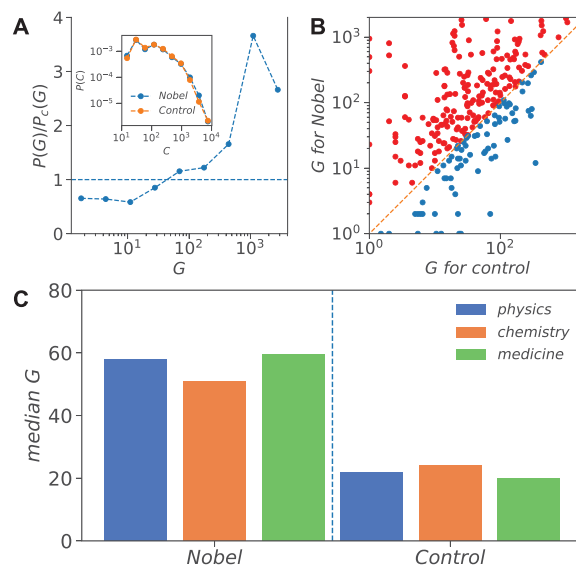


Figure 5. Giant index and Nobel-Prize-winning papers. (A) For each Nobel-Prize-winning paper, we construct a comparison set of papers that are published in the same year and field and have a similar number of citations (inset). We then compare the distributions of giant index for these two groups of papers, calculating the relative probability $P(G)/P_C(G)$ as a function of the giant index, where $P(G)$ is the distribution of giant index for Nobel Prize papers and $P_C(G)$ is the distribution of giant index for their nonprize-winning counterparts. (B) We compare the giant index between the prize-winning papers and their nonprize-winning counterparts in the comparison set. Red dots indicate prize-winning papers that have a higher G than those in the comparison group. (C) The median giant index for the Nobel-Prize-winning papers is 58 in physics, 51 in chemistry, and 59.5 in medicine, respectively. By contrast, the median giant index for papers in the comparison group with a similar level of citations is 20 in physics, 22 in chemistry, and 24 in medicine, respectively.

(Figure 5A, inset), yet the prevalence of Nobel-Prize-winning papers systematically increases with their giant index (Figure 5A). Indeed, if we just compare the giant index of prize-winning papers with that of their nonprize-winning counterparts, we find a majority of prize-winning papers have a higher G (67%) (Figure 5B). Hence, even though the two groups have the same citations, simply comparing their giant index offers a stronger signal to distinguish them. Indeed, we further compared the median giant index between the two samples, finding that across physics, chemistry, and medicine, the median giant index for prize-winning papers is more than twice of that of the control group (Figure 5C). Note that the exercise shown in Figure 5 illustrates a relatively simple approach, suggesting that the utility of the giant index can be further improved with additional features and more sophisticated models. Together, Figure 5 not only offers further evidence that incorporating the giant index may offer an additional signal to identify influential work than citation count alone; it also suggests distinctions between getting cited and being the reliable shoulder for ensuing science.

3. DISCUSSION

In summary, here we present a quantitative framework to test Newton's canonical insight on standing on the shoulders of giants. Several past efforts have attempted to identify a paper's influential references. While these attempts are mostly limited to small samples and require domain knowledge or manual processes for classification, our method offers an alternative way to quantitatively identify crucial references within a paper, which has several advantages that are worth noting. First, the method is discipline or journal independent, easily applicable

to any paper corpus where citation information is available. Second, it does not require ex-post information, meaning that the giant paper can be identified at the time of publication, which increases the practical utility of the method. Third, as a framework, the method can be further extended to allow for more flexibilities. For example, although the proposed method places an implicit assumption that each paper has just one giant, it can be easily extended to incorporate multiple giants by picking the top k references of a paper or references whose weight in the cociting network is above a threshold. Indeed, it is reasonable to assume that a paper may build upon multiple giants. Understanding papers with multiple giants is an important direction for future work. It would be interesting to analyze these cases and compare with the results obtained by picking one giant, which could be especially insightful when it comes to multidisciplinary work and could further illuminate the roles of giants in the development of new ideas.

The overall intuition behind the proposed measure is to recognize the uneven intellectual influence of the references to a specific paper by quantifying the relationships among the references cited in the paper through cocitation networks. As such, the method offers one way to balance between a reference's local importance to the specific context and its overall scientific impact, helping us better understand the intellectual significance of each reference in the context of the overall scientific discourse. Nevertheless, as an initial attempt to understand the role and characteristics of giants in science, there are several important limitations of our work, each suggesting important directions for future work. First and foremost is the validation of the proposed framework. In this paper, we calculated the giant score for Nobel-Prize-winning papers. Future work may compare the giant identified by our method to existing metrics, such as the number of mentions and the appearance in difference sections of papers through full text analysis. One could also assess the validity of the giant index through surveys, by asking, for example, the lead authors of a paper to identify the most important reference for their work, as well as asking whether there was any giant at all. This direction is especially important given the algorithmic nature of our method, which may represent a crude approximation for a paper's intellectual lineage that inherently depends on a range of social and institutional factors. Second, while we propose one way to quantify the intellectual lineage among papers, there could potentially be several other ways to quantify the shoulders of giants in science, representing fruitful directions for future work, which may lead to more robust methods and further insights. It is also important for future work to compare the performance and validity of different methods in identifying giants. Third, it is important to keep in mind that, as with many citation-based indicators, the giant index does not account for the many individual and institutional factors influencing a paper's future impact. Further, citation-based measures may also have inherent biases against recent work, as it takes time for citations to accumulate, suggesting that it may be more difficult to identify the giant(s) if a paper builds on a more recent body of work. Last, the prediction task in the paper offers correlational evidence supporting the relevance of the giant index. Future work with causal design may help improve the causative interpretation of the idea of standing on the shoulders of giants.

Overall, given the crucial importance of citations in science decision-making, including hiring, promotion, granting, and rewards, the developed concept of giant and its associated giant index may offer a useful dimension in our quantitative understanding of science by allowing us to appreciate those who shoulder scientific progress. As such, this measure is not limited to individual discoveries but offers a complementary dimension to the growing literature of the science of science (Fortunato et al., 2018; Wang & Barabási, 2021), and can be fruitfully applied to assess the role of giants in careers, teams, institutions, and more, pointing to promising future directions.

ACKNOWLEDGMENTS

The authors thank all members of the Center for Science of Science and Innovation (CSSI) for invaluable comments.

AUTHOR CONTRIBUTIONS

Woo Seong Jo: Data curation, Formal analysis, Investigation, Methodology, Visualization, Writing—original draft, Writing—review & editing. Lu Liu: Formal analysis, Investigation, Methodology, Visualization, Writing—original draft, Writing—review & editing. Dashun Wang: Conceptualization, Funding acquisition, Investigation, Methodology, Project administration, Supervision, Writing—original draft, Writing—review & editing.

COMPETING INTERESTS

The authors have no competing interests.

FUNDING INFORMATION

This work is supported by Air Force Office of Scientific Research under award nos. FA9550-15-1-0162, FA9550-17-1-0089, and FA9550-19-1-0354.

DATA AVAILABILITY

WoS data are available via Clarivate Analytics. The Nobel Prize data are from <https://www.nature.com/articles/s41597-019-0033-6>.

REFERENCES

- Aksnes, D. W. (2006). Citation rates and perceptions of scientific contribution. *Journal of the American Society for Information Science and Technology*, 57(2), 169–185. <https://doi.org/10.1002/asi.20262>
- Barabási, A.-L., & Albert, R. (1999). Emergence of scaling in random networks. *Science*, 286(5439), 509–512. <https://doi.org/10.1126/science.286.5439.509>, PubMed: 10521342
- Bergstrom, C. T., West, J. D., & Wiseman, M. A. (2008). The Eigenfactor (TM) Metrics. *Journal of Neuroscience*, 28(45), 11433–11434. <https://doi.org/10.1523/jneurosci.0003-08.2008>, PubMed: 18987179
- Bornmann, L., & Daniel, H.-D. (2008). Functional use of frequently and infrequently cited articles in citing publications. A content analysis of citations to articles with low and high citation counts. *European Science Editing*, 34(2), 35–38.
- Bornmann, L., de Moya Anegón, F., & Leydesdorff, L. (2010). Do scientific advancements lean on the shoulders of giants? A bibliometric investigation of the Ortega hypothesis. *PLOS ONE*, 5(10), e13327. <https://doi.org/10.1371/journal.pone.0013327>, PubMed: 20967252
- Boyack, K. W., van Eck, N. J., Colavizza, G., & Waltman, L. (2018). Characterizing in-text citations in scientific articles: A large-scale analysis. *Journal of Informetrics*, 12(1), 59–73. <https://doi.org/10.1016/j.joi.2017.11.005>
- Chen, C. (2006). CiteSpace II: Detecting and visualizing emerging trends and transient patterns in scientific literature. *Journal of the American Society for Information Science and Technology*, 57(3), 359–377. <https://doi.org/10.1002/asi.20317>
- Clauset, A., Larremore, D. B., & Sinatra, R. (2017). Data-driven predictions in the science of science. *Science*, 355(6324), 477–480. <https://doi.org/10.1126/science.aal4217>, PubMed: 28154048
- Cole, J. R., & Cole, S. (1972). The Ortega hypothesis: Citation analysis suggests that only a few scientists contribute to scientific progress. *Science*, 178(4059), 368–375. <https://doi.org/10.1126/science.178.4059.368>, PubMed: 17815351
- Cole, J. R., & Cole, S. (1974). *Social stratification in science*. Chicago and London: University of Chicago Press. <https://doi.org/10.1119/1.1987897>
- Cui, H., Zeng, A., Fan, Y., & Di, Z. (2020). Identifying the key reference of a scientific publication. *Journal of Systems Science and Systems Engineering*, 29(4), 429–439. <https://doi.org/10.1007/s11518-020-5455-3>
- Ding, Y., Liu, X., Guo, C., & Cronin, B. (2013). The distribution of references across texts: Some implications for citation analysis. *Journal of Informetrics*, 7(3), 583–592. <https://doi.org/10.1016/j.joi.2013.03.003>
- Ding, Y., Yan, E., Frazho, A., & Caverlee, J. (2009). PageRank for ranking authors in co-citation networks. *Journal of the American Society for Information Science and Technology*, 60(11), 2229–2243. <https://doi.org/10.1002/asi.21171>
- Ding, Y., Zhang, G., Chambers, T., Song, M., Wang, X., & Zhai, C. (2014). Content-based citation analysis: The next generation of citation analysis. *Journal of the Association for Information Science and Technology*, 65(9), 1820–1833. <https://doi.org/10.1002/asi.23256>
- Fortunato, S., Bergstrom, C. T., Börner, K., Evans, J. A., Helbing, D., ... Uzzi, B. (2018). Science of science. *Science*, 359(6379). <https://doi.org/10.1126/science.aao0185>, PubMed: 29496846
- Funk, R. J., & Owen-Smith, J. (2017). A dynamic network measure of technological change. *Management Science*, 63(3), 791–817. <https://doi.org/10.1287/mnsc.2015.2366>
- Garfield, E. (2006). Citation indexes for science. A new dimension in documentation through association of ideas. *International*

- Journal of Epidemiology*, 35(5), 1123–1127. <https://doi.org/10.1093/ije/dyl189>, PubMed: 16987841
- Garfield, E., & Malin, M. V. (1968). Can Nobel Prize winners be predicted? *135th Annual Meeting of the American Association for the Advancement of Science*, Dallas, TX.
- Hassan, S.-U., Akram, A., & Haddawy, P. (2017). Identifying important citations using contextual information from full text. *2017 ACM/IEEE Joint Conference on Digital Libraries (JCDL)*. <https://doi.org/10.1109/JCDL.2017.7991558>
- Hirsch, J. E. (2005). An index to quantify an individual's scientific research output. *Proceedings of the National Academy of Sciences*, 102(46), 16569–16572. <https://doi.org/10.1073/pnas.0507655102>, PubMed: 16275915
- Hu, Z., Chen, C., & Liu, Z. (2013). Where are citations located in the body of scientific articles? A study of the distributions of citation locations. *Journal of Informetrics*, 7(4), 887–896. <https://doi.org/10.1016/j.joi.2013.08.005>
- Jones, T. H., & Hanney, S. (2016). Tracing the indirect societal impacts of biomedical research: Development and piloting of a technique based on citations. *Scientometrics*, 107(3), 975–1003. <https://doi.org/10.1007/s11192-016-1895-4>, PubMed: 27340306
- King, D. A. (2004). The scientific impact of nations. *Nature*, 430(6997), 311–316. <https://doi.org/10.1038/430311a>, PubMed: 15254529
- Liu, L., Wang, Y., Sinatra, R., Giles, C. L., Song, C., & Wang, D. (2018). Hot streaks in artistic, cultural, and scientific careers. *Nature*, 559(7714), 396. <https://doi.org/10.1038/s41586-018-0315-8>, PubMed: 29995850
- Liu, S., & Chen, C. (2012). The proximity of co-citation. *Scientometrics*, 91(2), 495–511. <https://doi.org/10.1007/s11192-011-0575-7>
- Lü, L., Chen, D., Ren, X.-L., Zhang, Q.-M., Zhang, Y.-C., & Zhou, T. (2016). Vital nodes identification in complex networks. *Physics Reports*, 650, 1–63. <https://doi.org/10.1016/j.physrep.2016.06.007>
- Lü, L., & Zhou, T. (2011). Link prediction in complex networks: A survey. *Physica A: Statistical Mechanics and its Applications*, 390(6), 1150–1170. <https://doi.org/10.1016/j.physa.2010.11.027>
- Moravcsik, M. J., & Murugesan, P. (1975). Some results on the function and quality of citations. *Social Studies of Science*, 5(1), 86–92. <https://doi.org/10.1177/030631277500500106>
- Newman, M. (2010). *Networks: An introduction*. New York: Oxford University Press. <https://doi.org/10.1093/acprof:oso/9780199206650.001.0001>
- Price, D. J. D. S. (1965). Networks of scientific papers. *Science*, 149(3683), 510–515. <https://doi.org/10.1126/science.149.3683.510>, PubMed: 14325149
- Radicchi, F. (2012). In science “there is no bad publicity”: Papers criticized in comments have high scientific impact. *Scientific Reports*, 2(1), 1–5. <https://doi.org/10.1038/srep00815>, PubMed: 23139864
- Radicchi, F., Fortunato, S., & Castellano, C. (2008). Universality of citation distributions: Toward an objective measure of scientific impact. *Proceedings of the National Academy of Sciences*, 105(45), 17268–17272. <https://doi.org/10.1073/pnas.0806977105>, PubMed: 18978030
- Radicchi, F., Ramasco, J. J., & Fortunato, S. (2011). Information filtering in complex weighted networks. *Physical Review E*, 83(4), 046101. <https://doi.org/10.1103/PhysRevE.83.046101>, PubMed: 21599234
- Redner, S. (2005). Citation statistics from 110 years of *Physical Review*. *Physics Today*, 58(physics/0506056), 49. <https://doi.org/10.1063/1.1996475>
- Revesz, P. Z. (2015). Data mining citation databases: A new index measure that predicts Nobel Prizewinners. *Proceedings of the 19th International Database Engineering & Applications Symposium*. <https://doi.org/10.1145/2790755.2790763>
- Sarol, M. J., Liu, L., & Schneider, J. (2018). Testing a citation and text-based framework for retrieving publications for literature reviews. *CEUR Workshop Proceedings, 2080*, 22–33.
- Serrano, M. Á., Boguná, M., & Vespignani, A. (2009). Extracting the multiscale backbone of complex weighted networks. *Proceedings of the National Academy of Sciences*, 106(16), 6483–6488. <https://doi.org/10.1073/pnas.0808904106>, PubMed: 19357301
- Shen, H.-W., & Barabási, A.-L. (2014). Collective credit allocation in science. *Proceedings of the National Academy of Sciences*, 111(34), 12325–12330. <https://doi.org/10.1073/pnas.1401992111>, PubMed: 25114238
- Simkin, M. V., & Roychowdhury, V. P. (2002). Read before you cite! *arXiv preprint cond-mat/0212043*.
- Sinatra, R., Wang, D., Deville, P., Song, C., & Barabási, A.-L. (2016). Quantifying the evolution of individual scientific impact. *Science*, 354(6312), aaf5239. <https://doi.org/10.1126/science.aaf5239>, PubMed: 27811240
- Small, H. (1973). Co-citation in the scientific literature: A new measure of the relationship between two documents. *Journal of the American Society for Information Science*, 24(4), 265–269. <https://doi.org/10.1002/asi.4630240406>
- Tahamtan, I., & Bornmann, L. (2018). Creativity in science and the link to cited references: Is the creative potential of papers reflected in their cited references? *Journal of Informetrics*, 12(3), 906–930. <https://doi.org/10.1016/j.joi.2018.07.005>
- Uzzi, B., Mukherjee, S., Stringer, M., & Jones, B. (2013). Atypical combinations and scientific impact. *Science*, 342(6157), 468–472. <https://doi.org/10.1126/science.1240474>, PubMed: 24159044
- Waltman, L. (2016). A review of the literature on citation impact indicators. *Journal of Informetrics*, 10(2), 365–391. <https://doi.org/10.1016/j.joi.2016.02.007>
- Wang, D., & Barabási, A.-L. (2021). *The science of science*. Cambridge, MA: Cambridge University Press. <https://doi.org/10.1017/9781108610834>
- 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>, PubMed: 24092745
- Way, S. F., Morgan, A. C., Larremore, D. B., & Clauset, A. (2019). Productivity, prominence, and the effects of academic environment. *Proceedings of the National Academy of Sciences*, 116(22), 10729–10733. <https://doi.org/10.1073/pnas.1817431116>, PubMed: 31036658
- Wu, L., Wang, D., & Evans, J. A. (2019). Large teams develop and small teams disrupt science and technology. *Nature*, 566(7744), 378–382. <https://doi.org/10.1038/s41586-019-0941-9>, PubMed: 30760923
- Wuchty, S., Jones, B. F., & Uzzi, B. (2007). The increasing dominance of teams in production of knowledge. *Science*, 316(5827), 1036–1039. <https://doi.org/10.1126/science.1136099>, PubMed: 17431139
- Zakhlebin, I., & Horvát, E.-Á. (2017). Network signatures of success: Emulating expert and crowd assessment in science, art, and technology. In: C. Cherifi, H. Cherifi, M. Karsai, & M. Musolesi (Eds.), *Complex Networks & Their Applications VI. COMPLEX NETWORKS 2017*. Springer, Cham. https://doi.org/10.1007/978-3-319-72150-7_36
- Zhu, X., Turney, P., Lemire, D., & Vellino, A. (2015). Measuring academic influence: Not all citations are equal. *Journal of the Association for Information Science and Technology*, 66(2), 408–427. <https://doi.org/10.1002/asi.23179>