Network Signatures of Success: Emulating Expert and Crowd Assessment in Science, Art, and Technology

Igor Zakhlebin[1,2](✉) and Emőke-Ágnes Horvát[1]

[1] Northwestern University, Evanston, IL, USA
izakhlebin@u.northwestern.edu, a-horvat@northwestern.edu
[2] NRU Higher School of Economics, Moscow, Russia

Abstract. The success of scientific, artistic, and technological works is typically judged by human experts and the public. Recent empirical literature suggests that exceptionally creative works might have distinct patterns of citation. Given the recent availability of large citation and reference networks, we investigate how highly successful works differ from less successful ones in terms of a broad selection of centrality indices. Our experiments show that expert opinion is better emulated than popular judgment even with a single well-chosen index. Our findings further provide insights into otherwise implicit assumptions about indicators of success by evaluating the success of works based on the patterns of references that they receive.

1 Introduction

Key areas of human activity such as science, art, and technology advance by producing works and artifacts that build on top of each other. Drawing on ideas from previous works enables a recombination process that has been found to aid the production of novelty [21,24,27,32]. Dependencies on prior work display an intricate web of knowledge transfer and suggest that a network-based modeling framework could indicate which works accumulate the highest impact over time [15,25,28]. In all three areas of interest, research shows that based on various criteria of success some works are much more appreciated than others: the impact of scientific papers, adoption of technologies, and popularity of various forms of art differ widely between the most and least successful works [17,20,23,31]. To understand the extreme inequality in these "winner-take-all" settings and to capitalize on the emerging superstars [6,22], it is important to be able to recognize outstanding works and study why they are perceived as such.

In the absence of an agreed upon definition for the intrinsic value of works, the two main points of reference in assessing success are *(i)* expert judgment and *(ii)* public perception. Trained experts have been traditionally entrusted with a crucial role in determining which works are exceptional. However, specialists in various fields are often inaccurate and biased [1,10,16]. These misjudgments

C. Cherifi et al. (eds.), *Complex Networks & Their Applications VI*,
Studies in Computational Intelligence 689, https://doi.org/10.1007/978-3-319-72150-7_36

have contributed to an increased support for approaches that explore the wisdom of crowds coming from popular belief [5].

In this paper, we formalize notions of success in the areas of science, art, and technology using referencing patterns between works. For that, we build networks in which connections represent references between the works [25,29]. We consider three different settings: *(1)* citations between scientific publications, *(2)* references between feature films, and *(3)* dependencies between software projects. We argue that these networks contain the information about dissemination of knowledge in all three domains. First, we select a set of different centrality indices inspired by literature on citation networks to quantify importance. Then, we ask whether those indices are able to emulate the costly expert assessment as well as difficult-to-obtain public judgment.

Our results indicate that a single well-chosen centrality index is often enough to emulate the critical acclaim of works. At the same time, understanding popularity is a more difficult problem, for which inclusion of multiple centrality indices is beneficial. This research contributes to the recognition of intrinsically influential works instead of encouraging the introduction of narrow measures of impact, which are likely to create feedback loops that prompt people to optimize for those specific metrics in their pursuit of success [12] and undermine the creation of genuine value in the long term.

The rest of the paper is structured as follows. First, we review different centrality indices previously applied to citation networks and select a representative sample of them in Sect. 2. Second, we collect networks of references from three different areas of creative human activity (Sect. 3) on which we run a set of analyses to emulate expert and crowd evaluation (Sect. 4). Results described in Sect. 5 show which indices are important for emulating such evaluations. We conclude with discussion of findings and possible extensions of this work.

2 Background

Network definitions. We model networks as graphs $G = (\mathscr{V}, \mathscr{E})$ which consist of a set of nodes $\mathscr{V}$ and a set of edges $\mathscr{E}$ that connect those nodes. For directed edges, we write $(v \rightarrow w)$ if the edge goes from node v to node w. $\mathscr{N}(v)$ denotes the set of in-neighbors of node v such that $\mathscr{N}(v) = \{u \in \mathscr{V} : (u \rightarrow v) \in \mathscr{E}\}$. A path between nodes v_1 and v_k is defined as $P(v_1, v_k) = \{v_1, v_2, \ldots, v_k\}$, i.e., a sequence of nodes such that $v_j \in \mathscr{V}$ and $(v_j \rightarrow v_{j+1}) \in \mathscr{E}, \forall 1 \leq j < k$. A directed path that starts and ends at the same node is called a cycle. A directed graph without any cycles is called a directed acyclic graph (DAG). $\mathscr{I}(v) = \{u \in \mathscr{V} : \exists P(u, v)\}$ denotes the set of all nodes from which v can be reached by directed paths, and is said to be the in-component of v [18, p. 143–145] except v itself. Each node v is associated with a time stamp t_v, which is equal to release date of the artifact represented by node v.

Nodes in networks of references represent individual works and directed edges represent references between them. Since networks of references grow over time by addition of new nodes with references to previously existing nodes, edges can only go from newer works to older ones, as shown on Fig. 1. Real-world networks may not satisfy this assumption fully, so additional processing to remove non-permitted edges is required (see Sect. 3 for details).

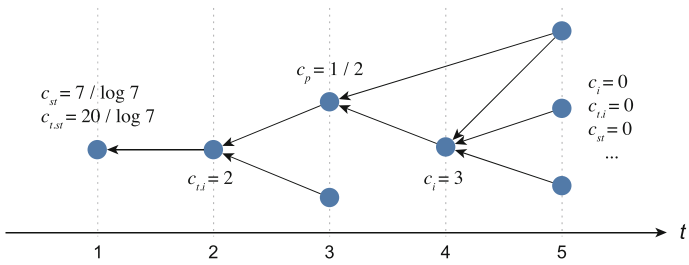

Fig. 1. Schematic representation of a reference network annotated with the centrality values described below

Centrality in citation networks. There is a significant body of literature studying references between scientific papers. The typical focus of those works is on predicting the influence of individual papers [4], their authors [11], and publication venues [9]. More recently, there has been a stream of studies doing similar research on cinematic references [25,30] and citations between patents [17].

The most commonly studied feature of nodes in citation networks is the number of references: the more incoming references a work receives, the more impactful it is [8]. That number is captured by the node's *indegree*. A more recent idea is to incorporate the presumed quality of references into the centrality index. If a work receives two incoming references—one from a very important work and one from an average work—the former reference is a stronger indicator of this work being itself important. This intuition stands behind *PageRank* [19] and has been applied to assess the importance of scientific journals [9] and papers within [4]. Finally, one could count the total number of works that can be traced back to a particular work by either direct references or paths of references going through other works. Two centrality indices based on this number are *subtree centrality* [25] and *global reaching centrality* [15]. Since earlier works have first-mover advantage in that they have more opportunities to accumulate incoming edges, these centralities are usually normalized, for example, by the logarithm of total number of works that have appeared after the considered artifact.

The second most studied feature of citation networks is the distribution of references with respect to the amount of time passed since the work's release. For example, it is known that people tend to cite mostly recent works [3]. This tendency has been shown to hold for both scientific citations [3,28] and movie references [30]. Hence, longevity of the work can serve as a proxy for its importance, which can be captured by the *long-gap citation count*, i.e., the number of citations with time difference larger than a certain threshold [30], or *temporal degree centrality*, the sum of differences between the referenced work's and referencing works' publication times [25]. Applying the same temporal intuition to other centrality indices yields the so-called *propagation centrality* [25], which is similar to PageRank and Katz centrality and assigns a score to each node according to the scores of its neighbors after weighing them by the time difference between the nodes. In this paper we additionally introduce the *temporal subtree centrality* based on previously discussed subtree centrality, having each component weighted by the time difference between referencing and referenced works.

Accordingly, we work with six different centrality indices (Table 1) that are meaningful to apply to networks of references. We group them in a 3×2 matrix with respect to the two aforementioned properties: locality (local to global) and temporality (temporal or not). Since global centrality indices should incorporate more information about the network than the local ones, we expect them to be better predictors of the works' importance. Similarly, temporal indices contain information that is complimentary to the bare structure of references, so they might be better predictors than their non-temporal counterparts.

Table 1. Centrality indices used throughout this study; these formulas assume a DAG as an input

Locality	Regular	Temporal
Local	Indegree centrality [20] $c_i(v) = \lvert\mathscr{N}(v)\rvert$	Temporal indegree centrality [25] $c_{t.i}(v) = \sum_{u \in \mathscr{N}(v)} (t_u - t_v)$
Mid-range	PageRank centrality [19] $c_{PR}(v) = \sum_{u \in \mathscr{N}(v)} \frac{c_{PR}(u)}{deg^{out}(u)} + \sum_{u \in \{l \in \mathscr{V} : deg^{out}(l)=0\}} \frac{c_{PR}(u)}{N}$	Propagation centrality [25] $c_p(v) = \sum_{u \in \mathscr{N}(v)} \alpha(v,u) \cdot \beta(u)$ $\beta(u) = \begin{cases} 1 & \text{if } deg^{in}(u) = 0 \\ c_p(u) & \text{otherwise} \end{cases}$ $\alpha(v,u) = \max(1 - \frac{1}{t_u - t_v}, 0)$
Global	Subtree centrality [15,25] $c_{st}(v) = \frac{\lvert\mathscr{I}(v)\rvert}{\log(\lvert\{u \in \mathscr{V} : t_u > t_v\}\rvert)}$	Temporal subtree centrality $c_{t.st}(v) = \sum_{u \in \mathscr{I}(v)} \frac{t_u - t_v}{\log(\lvert\{u \in \mathscr{V} : t_u > t_v\}\rvert)}$

Indicators of success. Indicators of works' importance can be divided into two general groups based on the qualification of people who are consulted [14]: *(i) Expert opinion*. In case of scientific papers, the recognition of experts comes in form of highly selective prizes that distinguish "foundational" works [11]. For movies, expert appraisal can entail inclusion in critically curated lists or receipt of awards [30]. In case of some technologies, experts can provide access to them on a variety of platforms. *(ii) Public opinion*. In the case of papers, their number of readers is a good proxy for popularity [26]. For movies, widely-used indicators of popularity are user ratings and box office revenue [30]. Finally, the public appreciation of a technology is captured by the number of people who use it and their feedback about it. Due to considerable conceptual and functional distinctions between these two forms of evaluation, we examine them separately.

3 Data Descriptives

Data. We have collected data from three different domains as summarized in Table 2. The three resulting datasets are released along with this paper.[1]

Physics paper citations (APS papers). We use a dataset of over 400,000 papers published by the American Physical Society (APS). The data includes basic information such as paper titles, authors, publication dates, and DOI numbers. It also provides citations between all papers published in journals run by APS between 1893 and 2009. These citations are used to construct the network of references. To quantify expert acclaim, we use the manually collected information about Nobel prize winners in Physics 1995–2013 and the papers that led to awarding those prizes. As an indicator of popularity, we use the numbers of readers on the popular bibliographic management service Mendeley, which have been retrieved through their public API[2] for all APS papers cited at least once.

Cinematic references (IMDb movies). We use a set of references between movies taken from the Internet Movie Database (IMDb), the oldest and most complete online movie catalogue. We have extracted a snapshot of this data on June 18, 2017.[3] To build the network of references we preprocessed the data by limiting nodes to feature films released until 2017 and connections to the types labeled "references", "features", "spoofs", and "edited from." The proxies for success in this case are *(i)* inclusion in the National Film Preservation Board's film registry[4] as of 2016 and *(ii)* aggregated IMDb user ratings that signal public perception of the movies.

Python package dependencies (PyPI packages). Many programming languages have adopted the notion of packages, which are self-contained modules that provide additional bits of functionality to the core language. When one package builds on another, it is said to *depend* on it. In case of the Python language, such information is stored in the Python Package Index (PyPI). We construct a network where nodes are individual Python packages and edges are dependencies between them. The dependency data is publicly available through the Libraries.io project[5], which provides snapshots of data from PyPI. The snapshot used here was retrieved on 21 July, 2017. To approximate the expert evaluation, we assume that a package is important if it is included in a popular software distribution like Anaconda[6] for Python 2.7. To quantify the public appeal, we count "stars" that package repositories have received on the widely used software development website GitHub through their public API[7] In the rare cases when several packages share the same repository, we split the number of stars between them equally.

---

[1] https://github.com/inguar/network-signatures-of-success-data.
[2] https://api.mendeley.com/apidocs/docs.
[3] http://www.imdb.com/interfaces.
[4] https://www.loc.gov/programs/national-film-preservation-board/film-registry/.
[5] https://libraries.io/data.
[6] https://docs.continuum.io/anaconda/packages/pkg-docs.
[7] https://developer.github.com/.

Table 2. Descriptives of networks and corresponding success indicators with non-permitted edges and isolate nodes removed from the networks

	APS papers	IMDb movies	PyPI packages
Nodes	449,463	36,335	5,211
Components	342	1,350	158
Year range	1893–2009	1904–2017	2002–2017
Edges	4,690,075	103,993	10,355
Non-permitted edges	17,883	2,068	636
Expert evaluations (no. of successes)	Won Nobel prize 43	Included in National Film Registry 680	Distributed by Anaconda 261
Popularity indicators (no. of observations)	Readers on Mendeley 270,082	IMDb user ratings 21,367	Stars on GitHub 1,302

Preprocessing. When constructing networks of references based on these data, we perform two cleaning steps: *(i)* We remove non-permitted edges, i.e., edges that go from older works to newer ones or between the works that appeared simultaneously. One could reasonably expect that such edges should not exist, but they legitimately occur in practice[8]. *(ii)* We remove isolate nodes, i.e., works that don't reference and are not referenced by any other works.

Next, we compute the six centrality indices listed in Table 1 for every node in the cleaned networks. We further exclude the nodes for which all centrality indices are equal to 0 *and* they are neither critically acclaimed nor highly popular. Their inclusion in further analyses would inflate the estimates of model fitness without meaningfully contributing to it.

Relationships between variables. Fig. 2 shows correlations among different centrality indices. The Kendall rank correlation coefficient compares centralities of nodes relative to each other and is thus a good choice given that the distributions of centrality values are highly skewed. According to it, all indices exhibit moderate correlation with each other. On the one hand, in the APS data, the first four indices (indegree, temporal indegree, PageRank, and propagation centrality) and the last two (regular and temporal subtree centrality) correlate with each other more highly than with others. In the IMDb and PyPI data, on the other hand, temporal centrality indices generally exhibit higher correlation with each other than with other indices resulting in a checkerboard pattern. Additionally, the indegree centrality shows a consistently higher correlation with every other index.

[8] Authors of scientific papers can reference their forthcoming work or papers in the same volume. Movies can have anomalous references due to delayed release dates and avid marketing strategies in pre-release stage. Developers can freely change dependencies after publishing their packages, sometimes choosing more recent packages.

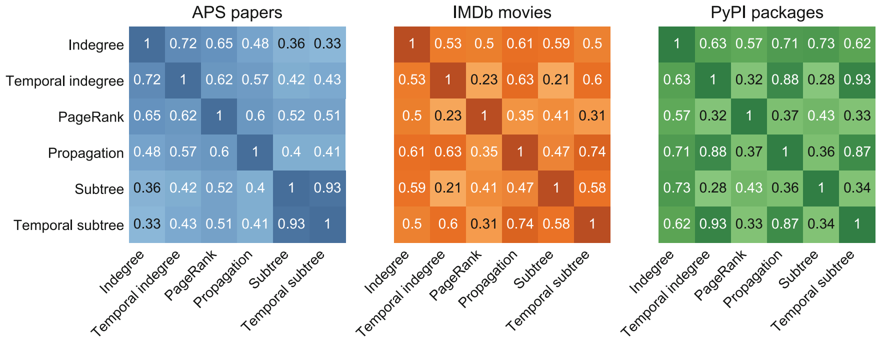

Fig. 2. Kendall rank correlation coefficients between centrality indices show considerable differences between the three scenarios

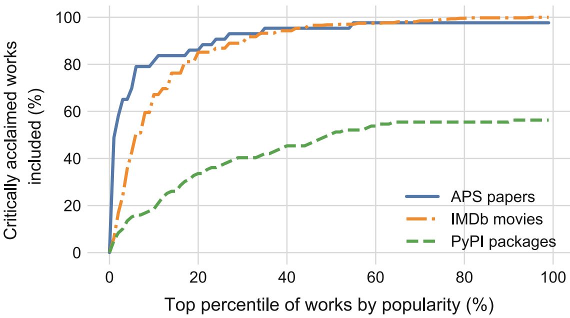

Fig. 3. Correspondence between the expert and public evaluation of works

Figure 3 shows the percentage of critically acclaimed works that is included in the expanding list of top $k\%$ most popular works. All critically acclaimed movies in the IMDb data set have a crowd rating associated with them; nearly all Nobel prize-winning APS papers have a Mendeley readership; and around 56% of the Anaconda-distributed packages could be matched to GitHub repositories and star counts. Most of the critically acclaimed works are usually included in a small portion of the most popular works: half of the critically acclaimed works is contained in the 2% most popular papers, 6% most popular movies, and 59% most popular packages. To keep the expert and public evaluations substantially different, in further analyses we consider a work popular if it lies within the top 1% of APS papers, top 5% of IMDb movies, or top 10% of PyPI packages.

4 Experimental Setup

To emulate expert and public judgment based on networks of references, we consider the six centrality indices associated with individual works to be independent variables. From the collected indicators of success we deduce two binary dependent variables: *(i)* critical acclaim is defined to be 1 for works that have received a Nobel prize, are included in the National Film Registry's list, or are distributed with Anaconda; and *(ii)* popularity is considered to be 1 if the work is in the top 1% of APS papers, 5% of IMDb movies, or 10% of PyPI packages by popularity.

Using the network-based independent variables we formulate binary classification problems aimed at discerning highly successful works from the rest, as defined by our dependent variables. For each dataset, our analysis consists of two parts. First, we fit univariate logistic regressions using each of the centrality indices *separately.* We evaluate the performance of each centrality index using the area under the ROC curve (AUC) [7]. This measure aggregates classification accuracy over all possible settings of thresholds given the specificity-sensitivity trade-off and thus results in a more complete assessment of classification accuracy. An AUC=1 means perfect classification, while AUC=0.5 indicates random guessing.

In the second part, we deploy random forest classifiers [2] with the same dependent variables and *all* centrality indices at once. Random forest is a state-of-the-art ensemble method typically deployed in a supervised learning setting that has been used previously in network science, e.g., for link prediction. It relies on building n decision trees (in our case, $n = 100$) and uses majority voting to assign class probabilities based on individual trees' decisions. Beyond prediction, the random forest enables us to assess variable importance according to the Gini index [13]. The basic idea is that the more some particular variable is used across the trees to perform sample splits and the larger those samples are, the more important the variable is.

To ensure that our results are representative, we run a 5-fold cross-validation, i.e., we randomly partition 80% of nodes into the training and 20% into the test set. The classifiers learn on the training set and are then evaluated on the test set. This procedure is then repeated 20 times to improve reliability of estimates. Additionally, we under-sample observations from the prevalent class (0) to match the number of observations in the minority class (1) in all sets. This step is crucial in fitting and evaluating models since class imbalance in some cases is extreme (i.e., up to about 1:10,000 when predicting the expert acclaim of APS papers). We have tested the classifiers on imbalanced test sets as well and the results have been similar to the ones reported in the next section.

5 Results

Emulating expert opinion. Results of our univariate ordinal linear regressions evaluated with the AUC are shown in Fig. 4, which displays the means and standard deviations over all runs of each regression. For IMDb movies and PyPI packages, temporal versions of centrality indices perform better than their regular counterparts with a small margin. For APS papers, on the contrary, temporality does not seem to have a distinct effect. Additionally, more local indices like PageRank and temporal indegree perform relatively better.

A similar ordering of indices is generated through the relative importance of variables in random forest experiments (see Fig. 5). One interesting feature of these models is that their predictive power as quantified by AUC is on the same level as that of the best single-variable predictor. This fact might indicate that there is no emergent effect in combining centrality indices. Another possibility is that the predictability of success as assessed by experts is theoretically low and a single well-chosen index is enough to bring us sufficiently close to the maximum attainable accuracy. However, depending on the setting, the best performing index is different.

Emulating popularity. Most single-variable logistic regressions have significantly lower predictive capabilities at popularity emulation than at expert evaluation (see Fig. 6). There is no consensus over all data sets on whether temporal or local versions are better at predicting the popularity of works. For instance, in the PyPI data temporal indices outperform non-temporal ones. In the APS and IMDb data, however, non-temporal indices are systematically better.

The results of random forest analysis shown in Fig. 7 reveal that AUC values for APS and IMDb data are higher than the ones obtained by single-variable logistic regressions. Furthermore, the order of variables based on their relative importance has changed in comparison to the order determined by the regression analysis. Yet the random forest's accuracy at predicting popularity is lower than the accuracy of the same model when emulating expert evaluations. This pattern indicates that emulating popularity is a harder task that requires more information than what is contained in a single centrality index. Thus, the usage of multiple indices is beneficial.

Overall, we notice that indegree, despite its simplicity, is a strong predictor of both critical acclaim and popularity on its own, but when combined with other centrality indices its effect seems to diminish. This tendency can be explained by the fact that it has consistently higher correlations with other variables in both IMDb and PyPI data (cf. Fig. 2). Therefore, the information in it is partly contained in other variables and thus random forests tend to find it less important.

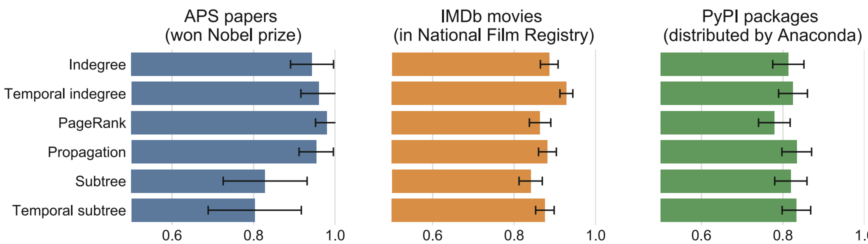

Fig. 4. Average AUC values computed from logistic regressions quantify the ability of centrality indices to emulate *expert opinion*. Error bars show standard deviations

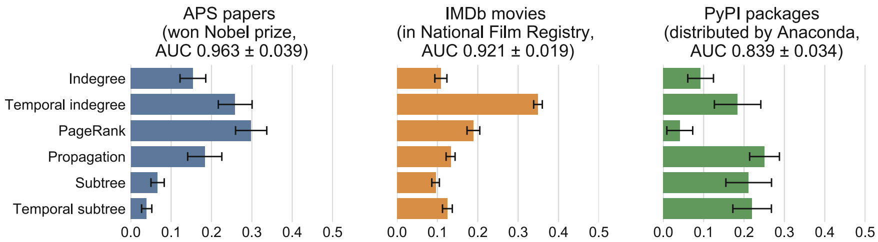

Fig. 5. Importance of different centrality indices according to the random forest classifier when emulating *expert opinion*. Error bars show standard deviations

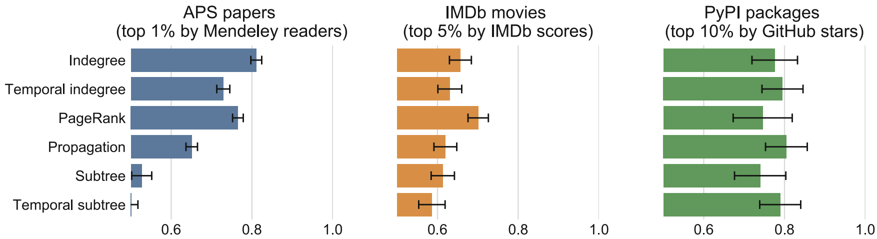

Fig. 6. Average AUC values computed from logistic regressions quantify the ability of centrality indices to emulate *the most popular* works. Error bars show standard deviations

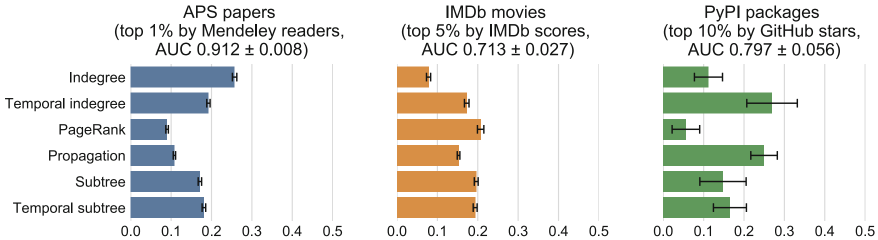

Fig. 7. Importance of different centrality indices according to the random forest classifier when emulating *the most popular* works. Error bars show standard deviations

6 Discussion

Despite the plethora of research aimed at determining the most important works in science, few efforts have been made to generalize those findings to other domains. Using novel data, this paper contributes to broadening the scope of citation analysis by further extending it to the areas of art and technology. We have fitted multiple models to uncover the ability of centrality indices computed from citation networks to emulate expert and public recognition in these three domains.

Our findings point to substantial differences between the areas in terms of predictability of success as well as which index is the best predictor. When emulating expert opinion, the best-performing index is as good as all of them taken together, but it is different in different settings: PageRank for physics papers, temporal indegree for movies, and propagation centrality for Python packages. When emulating popularity, the accuracy improves from inclusion of additional centrality indices, but it remains below the accuracy of emulating expert opinion.

These initial findings invite discussion within the research community about the reasons behind these patterns. They also bring up a number of technical and conceptual issues, such as different temporal scales, lack of consistent criteria for evaluation of success, as well as differences in the subject matter and definition of excellence/creativity in these areas. All these challenges are inevitable when conducting a systematic inquiry that overarches these fields, and our work attempts to perform it for the first time.

This work is not without limitations. A key shortcoming is the lack of a general framework for assessing the importance of works. Since such data is scarce and disputed, we use the two most common indicators—expert acclaim and general popularity. However, there is a much broader set of measures that might better estimate the value or creativity of works. This prompts a more systematic investigation into which success indicators should be used and to what extent they can be captured based on the structure of references. Further work should also address the problem of not knowing the theoretical limits to predictability of success.

Acknowledgements. Authors would like to thank Roberta Sinatra for valuable discussion and data on Nobel prize winners. Special thanks go to Noshir Contractor and SONIC lab members for supporting the preparation of this manuscript.I.Z. has been partially funded through the Russian academic excellence project “5–100”.

References

1. Bonabeau, E.: Decisions 2.0: the power of collective intelligence. MIT Sloan Manag. Rev. 45–52 (2009)
2. Breiman, L.: Random forests. Mach. Learn. **45**(1), 5–32 (2001)
3. Burton, R.E., Kebler, R.W.: The “half-life” of some scientific and technical literatures. Am. Doc. **11**(1), 18–22 (1960). https://doi.org/10.1002/asi.5090110105
4. Chen, P., Xie, H., Maslov, S., Redner, S.: Finding scientific gems with google’s pagerank algorithm. J. Inf. **1**(1), 8–15 (2007). https://doi.org/10.1016/j.joi.2006.06.001
5. Dreber, A., et al.: Using prediction markets to estimate the reproducibility of scientific research. Proc. Natl. Acad. Sci. **112**(50), 15343–15347 (2015)
6. Elberse, A.: Blockbusters: Hit-making, Risk-taking, and the Big Business of Entertainment. Henry Holt and Company, LLC (2013)
7. Fawcett, T.: An introduction to ROC analysis. Pattern Recognit. Lett. **27**(8), 861–874 (2006). https://doi.org/10.1016/j.patrec.2005.10.010

8. Garfield, E.: Citation indexes for science: a new dimension in documentation through association of ideas. Science **122**(3159), 108–111 (1955). https://doi.org/10.1126/science.122.3159.108
9. Garfield, E.: The history and meaning of the journal impact factor. JAMA **295**(1), 90–93 (2006). https://doi.org/10.1001/jama.295.1.90
10. Hahn, R., Tetlock, P.C.: Using information markets to improve public decision making. Harv. J. Law Public Policy **29**(1), 213–289 (2005)
11. Hirsch, J.E.: An index to quantify an individual's scientific research output. Proc. Natl. Acad. Sci. U. S. A. **102**(46), 16569–16572 (2005). https://doi.org/10.1073/pnas.0507655102
12. Lehmann, S., Jackson, A.D., Lautrup, B.E.: Measures for measures. Nature **444**(7122), 1003–1004 (2006). https://doi.org/10.1038/4441003a
13. Louppe, G., Wehenkel, L., Sutera, A., Geurts, P.: Understanding variable importances in forests of randomized trees. In: Advances in Neural Information Processing Systems, pp. 431–439 (2013)
14. Monechi, B., Gravino, P., Servedio, V.D.P., Tria, F., Loreto, V.: Significance and popularity in music production. R. Soc. Open Sci. **4**(7) (2017). https://doi.org/10.1098/rsos.170433
15. Mones, E., Pollner, P., Vicsek, T.: Universal hierarchical behavior of citation networks. J. Stat. Mech. Theory Exp. **2014**(5), P05023 (2014). https://doi.org/10.1088/1742-5468/2014/05/P05023
16. Moore, D.A., et al.: Confidence calibration in a multiyear geopolitical forecasting competition. Manag. Sci. (2016). https://doi.org/10.1287/mnsc.2016.2525
17. Mukherjee, S., Romero, D.M., Jones, B., Uzzi, B.: The nearly universal link between the age of past knowledge and tomorrow's breakthroughs in science and technology: the hotspot. Sci. Adv. **3**(4), e1601315 (2017)
18. Newman, M.E.: Networks. An introduction. Oxford (2010)
19. Page, L., Brin, S., Motwani, R., Winograd, T.: The PageRank citation ranking: Bringing order to the web. Tech. rep, Stanford InfoLab (1999)
20. Price, D.J.D.S.: Networks of scientific papers. Science pp. 510–515 (1965)
21. Reader, S.M., Morand-Ferron, J., Flynn, E.: Animal and human innovation: novel problems and novel solutions. Philos. Trans. R. Soc. Lond. B. Biol. Sci. **371**(1690) (2016). https://doi.org/10.1098/rstb.2015.0182
22. Rosen, S.: The economics of superstars. Am. Econ. Rev. **71**(5), 845–858 (1981)
23. Salganik, M.J., Dodds, P.S., Watts, D.: Experimental study of inequality and unpredictability in an artificial cultural market. Science **311**(10), 854–856 (2006)
24. Schilling, M.A.: A "small-world" network model of cognitive insight. Creativity Res. J. **17**(2–3), 131–154 (2005). https://doi.org/10.1080/10400419.2005.9651475
25. Spitz, A., Horvát, E.Á.: Measuring long-term impact based on network centrality: unraveling cinematic citations. PloS one **9**(10), e108857 (2014)
26. Thelwall, M., Maflahi, N.: When are readership counts as useful as citation counts? Scopus versus Mendeley for LIS journals. J. Assoc. Inf. Sci. Technol. (2014)
27. Uzzi, B., Mukherjee, S., Stringer, M., Jones, B.: Atypical combinations and scientific impact. Science **342**(6157), 468–472 (2013). https://doi.org/10.1126/science.1240474
28. Wang, D., Song, C., Barabási, A.L.: Quantifying long-term scientific impact. Science **342**(6154), 127–132 (2013). https://doi.org/10.1126/science.1237825
29. Wasserman, M., et al.: Correlations between user voting data, budget, and box office for films in the internet movie database. J. Assoc. Inf. Sci. Technol. **66**(4), 858–868 (2015). https://doi.org/10.1002/asi.23213

30. Wasserman, M., Zeng, X.H.T., Amaral, L.A.N.: Cross-evaluation of metrics to estimate the significance of creative works. Proc. Natl. Acad. Sci. **112**(5), 1281–1286 (2015). https://doi.org/10.1073/pnas.1412198112
31. Watts, D.: Everything is Obvious: Once You Know the Answer. Crown Business (2011)
32. Weitzman, M.: Recombinant growth. Q. J. Econ. **113**(2) (1998)