

# BIP! DB: A Dataset of Impact Measures for Scientific Publications

Thanasis Vergoulis  
vergoulis@athenarc.gr  
IMSI, ATHENA RC  
Athens, Greece

Ilias Kanellos  
ilias.kanellos@athenarc.gr  
IMSI, ATHENA RC  
Athens, Greece

Claudio Atzori  
claudio.atzori@isti.cnr.it  
ISTI, CNR  
Pisa, Italy

Andrea Mannocci  
andrea.mannocci@isti.cnr.it  
ISTI, CNR  
Pisa, Italy

Serafeim Chatzopoulos  
schatz@athenarc.gr  
IMSI, ATHENA RC  
Athens, Greece

Sandro La Bruzzo  
sandro.labruzzo@isti.cnr.it  
ISTI, CNR  
Pisa, Italy

Natalia Manola  
natalia@di.uoa.gr  
OpenAIRE  
Athens, Greece

Paolo Manghi  
paolo.manghi@isti.cnr.it  
ISTI, CNR  
Pisa, Italy

## ABSTRACT

The growth rate of the number of scientific publications is constantly increasing, creating important challenges in the identification of valuable research and in various scholarly data management applications, in general. In this context, measures which can effectively quantify the scientific impact could be invaluable. In this work, we present BIP! DB, an open dataset that contains a variety of impact measures calculated for a large collection of more than 100 million scientific publications from various disciplines.

## KEYWORDS

scientometrics, impact, research assessment

### ACM Reference Format:

Thanasis Vergoulis, Ilias Kanellos, Claudio Atzori, Andrea Mannocci, Serafeim Chatzopoulos, Sandro La Bruzzo, Natalia Manola, and Paolo Manghi. 2021. BIP! DB: A Dataset of Impact Measures for Scientific Publications. In *Proceedings of ACM Conference (Conference'17)*. ACM, New York, NY, USA, 5 pages. <https://doi.org/10.1145/nnnnnnn.nnnnnnn>

## 1 INTRODUCTION

The growth rate of the number of published scientific articles is constantly increasing [2]. At the same time, studies suggest that, among the vast number of published works, many are of low impact or may even contain research of questionable quality [7]. Consequently, identifying the most valuable publications for any given research topic has become extremely tedious and time consuming.

Quantifying the impact of scientific publications could facilitate this and other related tasks, which make up the daily routine of researchers and other professionals of the broader scientific and academic community. For instance, most contemporary search engines for research publications (e.g., Google Scholar, Semantic Scholar) combine keyword-based relevance with a scientific impact measure (usually citation counts) to rank their search results, in an attempt to help their users prioritise reading for literature review.

However, many impact measures, which are widely used in various applications, have inherent drawbacks. For instance, citation counts cannot differentiate citations based on the importance of the citing articles. This is an important drawback since citation counts may, for example, present a publication in a predatory journal, heavily cited by other trivial works, as invaluable, while disregarding the seminal importance of an otherwise sparsely cited publication that has influenced (and is cited by) a breakthrough article. For the same reason, citation counts are also vulnerable to various malpractices (e.g., excessive self-citation, citation cartels).

Another important issue, often overlooked by most existing academic search engines, is the importance of capturing publication impact from a broader set of perspectives. It is an oversimplification to rely on one impact measure only, since there are many different aspects of scientific impact [1, 8]. Indicatively, a publication's *influence* is its overall, long-term importance, which can be calculated based on its whole citation history; its *popularity*, on the other hand, is the publication's current attention in the research community, which is indicative of its expected short-term impact. These impact aspects are not entirely correlated and each one may be preferable for different applications. Consider, for example, an experienced researcher who needs to reexamine a topic of interest to learn about its latest developments. Ranking articles based on their popularity would be preferable for her. On the other hand, a young researcher wanting to delve into the same topic to prepare a survey would prefer to rank the relevant articles based on their influence. Although using citation counts would satisfy the needs of the young researcher to an extent, they would fail to help the needs of the experienced one, since citation counts are biased

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from [permissions@acm.org](mailto:permissions@acm.org).

Conference'17, July 2017, Washington, DC, USA  
© 2021 Association for Computing Machinery.  
ACM ISBN 978-x-xxxx-xxxx-x/YY/MM... \$15.00  
<https://doi.org/10.1145/nnnnnnn.nnnnnnn>

against recent articles. This is because any recent article (irrespective of its current attention in the research community) usually requires months or even years to receive its first citations [15], and eventually gain momentum.

Overlooking the above aspects is problematic. On the one hand, recently published research may be at the center of the current attention of the scientific community, although this is not reflected by the traditional measures. On the other hand, scientific impact should not be examined through a limited set of measures. This is important not only because a larger set of measures captures a wider range of impact aspects, providing a more complete picture about a publication's impact; since any individual measure used for research assessment is bound to be abused, based on Goodhart's law<sup>1</sup>, consulting multiple measures can work as a countermeasure to reduce the effects of relevant attacks/malpractices. Taking all these into consideration, it is part of the good practices in research assessment to consider a range of quantitative measures as inclusive as possible (e.g., this is also emphasized in the Declaration on Research Assessment<sup>2</sup>).

In this work, we present BIP! DB, an open dataset that contains various impact measures calculated for more than 104M scientific articles, taking into consideration more than 1.25B citation links between them. The production of this dataset is based on the integration of three major datasets, which provide citation data: OpenCitation's [13] COCI dataset, Microsoft Academic Graph, and Crossref. Based on these data, we perform citation network analysis to produce five useful impact measures, which capture three distinct aspects of scientific impact. This set of measures makes up a valuable resource that could be leveraged by various applications in the field of scholarly data management.

The remainder of this manuscript is structured as follows. In Section 2 we discuss various aspects of scientific impact and present the impact measures in the BIP! DB dataset. In Section 3 we provide technical details on the production of our dataset and in Section 4 we empirically show how distinct the different measures are and elaborate on some interesting issues about the dataset's potential uses and extensions. Finally, in Section 5 we conclude the work.

## 2 IMPACT ASPECTS & MEASURES

As mentioned in Section 1, there are many perspectives from which one can study scientific impact [1, 8]. Consequently, a multitude of impact measures and indicators have been introduced in the literature, each of them better capturing a different aspect. In this work, we present an open dataset of different impact measures, which we freely provide. We focus on measures that quantify three aspects of scientific impact, which can be useful for various real-life applications: *popularity*, which reflects a publication's current attention (and its expected impact in the near future), *influence* which reflects its overall, long-term importance, and *impulse*, which better captures its initial impact during its "incubation phase", i.e., during the first years following its publication.

We focus on the first two aspects because they are well-studied: a recent experimental study [8] has revealed the strengths and weaknesses of various measures in terms of quantifying popularity

and influence. Based on this analysis, and on a set of subsequent experiments [9], we select the following measures: Citation Count, PageRank, RAM, and AttRank. We also include the impulse, since it is a distinct impact aspect that is reflected by a Citation Count variant, which is utilised for the production of some useful metrics, like the FWCI<sup>3</sup>. This variant considers only citations received during the "incubation" phase of a publication, i.e., during the first  $y$  years after its publication (usually,  $y = 3$ ); we refer to this useful Citation Count variant as *Incubation Citation Count*, and we select it for inclusion in the BIP! DB collection, as well. In the next paragraphs, we elaborate on all the impact measures of our dataset.

**Citation Count (CC).** This is the most widely used scientific impact indicator, which sums all citations received by each article. The citation count of a publication  $i$  corresponds to the in-degree of the corresponding node in the underlying citation network:  $s_i = \sum_j A_{i,j}$ , where  $A$  is the adjacency matrix of the network (i.e.,  $A_{i,j} = 1$  when paper  $j$  cites paper  $i$ , while  $A_{i,j} = 0$  otherwise). Citation count can be viewed as a measure of a publication's overall impact, since it conveys the number of other works that directly drew on it.

**"Incubation" Citation Count (iCC).** This measure is essentially a time-restricted version of the citation count, where the time window is distinct for each paper, i.e., only citations  $y$  years after its publication are counted (usually,  $y = 3$ ). The "incubation" citation count of a paper  $i$  is calculated as:  $s_i = \sum_{j, t_j \leq t_i+3} A_{i,j}$ , where  $A$  is the adjacency matrix and  $t_j, t_i$  are the citing and cited paper's publication years, respectively. iCC can be seen as an indicator of a paper's initial momentum (impulse) directly after its publication.

**PageRank (PR).** Originally developed to rank Web pages [11], PageRank has been also widely used to rank publications in citation networks (e.g., [3, 10, 16]). In this latter context, a publication's PageRank score also serves as a measure of its influence. In particular, the PageRank score of a publication is calculated as its probability of being read by a researcher that either randomly selects publications to read or selects publications based on the references of her latest read. Formally, the score of a publication  $i$  is given by:

$$s_i = \alpha \cdot \sum_j P_{i,j} \cdot s_j + (1 - \alpha) \cdot \frac{1}{N} \quad (1)$$

where  $P$  is the stochastic transition matrix, which corresponds to the column normalised version of adjacency matrix  $A$ ,  $\alpha \in [0, 1]$ , and  $N$  is the number of publications in the citation network. The first addend of Equation 1 corresponds to the selection (with probability  $\alpha$ ) of following a reference, while the second one to the selection of randomly choosing any publication in the network. It should be noted that the score of each publication relies on the score of publications citing it (the algorithm is executed iteratively until all scores converge). As a result, PageRank differentiates citations based on the importance of citing articles, thus alleviating the corresponding issue of the Citation Count.

**RAM.** RAM [5] is essentially a modified Citation Count, where recent citations are considered of higher importance compared to older ones. Hence, it better captures the popularity of publications. This "time-awareness" of citations alleviates the bias of methods

<sup>1</sup>Also known as Campbell's law (an example of "Cobra effect").

<sup>2</sup>The Declaration on Research Assessment (DORA), <https://sfedora.org/read>

<sup>3</sup>More info on this measure can be found in the Snowball metrics cookbook: <https://snowballmetrics.com>.

like Citation Count and PageRank against recently published articles, which have not had “enough” time to gather as many citations. The RAM score of each paper  $i$  is calculated as follows:

$$s_i = \sum_j R_{i,j} \quad (2)$$

where  $R$  is the so-called Retained Adjacency Matrix (RAM) and  $R_{i,j} = \gamma^{t_c - t_j}$  when publication  $j$  cites publication  $i$ , and  $R_{i,j} = 0$  otherwise. Parameter  $\gamma \in (0, 1)$ ,  $t_c$  corresponds to the current year and  $t_j$  corresponds to the publication year of citing article  $j$ .

**AttRank.** AttRank [9] is a PageRank variant that alleviates its bias against recent publications (i.e., it is tailored to capture popularity). AttRank achieves this by modifying PageRank’s probability of randomly selecting a publication. Instead of using a uniform probability, AttRank defines it based on a combination of the publication’s age and the citations it received in recent years. The AttRank score of each publication  $i$  is calculated based on:

$$s_i = \alpha \cdot \sum_j P_{i,j} \cdot s_j + \beta \cdot Att(i) + \gamma \cdot c \cdot e^{-\rho \cdot (t_c - t_i)} \quad (3)$$

where  $\alpha + \beta + \gamma = 1$  and  $\alpha, \beta, \gamma \in [0, 1]$ .  $Att(i)$  denotes a recent attention-based score for publication  $i$ , which reflects its share of citations in the  $y$  most recent years,  $t_j$  is the publication year of article  $i$ ,  $t_c$  denotes the current year, and  $c$  is a normalisation constant. Finally,  $P$  is the stochastic transition matrix.

### 3 THE BIP! DB DATASET

#### 3.1 Data Collection & Integration

BIP! DB’s impact measures all rely on citation network analysis. Hence, a major challenge in our work was to construct an interdisciplinary and as inclusive as possible citation network on which all impact measures would be calculated. To achieve this, we gathered citation data and metadata from three data sources: OpenCitations’ [13] COCI dataset, Microsoft’s Academic Graph (MAG) [14, 17], and Crossref [6]. The current dataset version (regular updates are scheduled in the future) exploits the latest version of the COCI dataset (Sep 2020), and recent snapshots of MAG (Aug 2020) and Crossref (May 2020). Our dataset production workflow collects, cleans, and integrates data from these sources to produce a citation graph based on the distinct DOI-to-DOI relationships found. Since the publication year is required for some of the measures to be calculated, publications lacking this information were excluded from the final network. Table 1 summarises some statistics for the original data sources and the complete, integrated dataset.

#### 3.2 Calculation of Impact Measures

As discussed in Section 3.1, the volume of processed data exceeds 100 million publications and 1 billion references. Hence, particular care must be taken by any algorithms developed for the calculation of the required impact scores, to allow for the handling of time-efficient and scalable updates.

By examining the formulas of the impact measures presented in Section 2, we can observe that all of them rely on the analysis of the underlying citation network, by calculating a (possibly weighted) sum of scores, which are received by each publication by its citing articles. Hence, we can take advantage of data parallelism

Dataset	DOIs	Citations
COCI	59,455,882	733,366,727
CrossRef	96,703,144	596,803,579
MAG	90,224,789	1,177,733,277
Unified Graph	104,769,307	1,254,817,030

**Table 1: Distinct DOIs and citations per data source.**

when implementing the algorithms calculating these measures. In particular, each measure can be implemented as a set of MapReduce operations where each publication *maps* its score (e.g., a citation, its PageRank score, etc) to its cited papers. The final score of each publication, in turn, results from an aggregation (*reduce*) of all scores mapped to it. Additionally, PageRank, and AttRank in particular, are iterative processes, which require such sets of operations to repeat until the calculated publication scores converge. Therefore, we chose Spark, which is particularly suitable for data parallel iterative processes, as our development platform of reference. In particular, we implemented all algorithms as PySpark scripts, running on Spark version 2.3. All impact measure calculations are performed on a cluster of 10 VMs, each with 4 cores and 8GB RAM, and each script runs with up to 35 Spark workers.

#### 3.3 Published Data Records

The current version of the BIP! DB dataset consists of five compressed TSV files, one for each impact measure provided. All files follow the same format: each line contains two data columns, where the first corresponds to the DOI of a publication, followed by the column which corresponds to the score of the measure.

For the sake of clarity, each of the files published contains in its name the configuration of the corresponding parameter setting of each algorithm, which produced the respective measure scores. For example, the file named “PR\_graph\_universe2\_1.txt\_a0.5\_error1e-12.gz” contains the PageRank scores calculated with parameter  $\alpha = 0.5$  and with a convergence error set to  $\epsilon \leq 10^{-12}$ . All files published are freely available in Zenodo<sup>4</sup> under the Creative Commons Attribution 4.0 International license.

#### 3.4 Updated BIP! API

Extra effort was given to update the existing BIP! API [16] with the most recent version of the BIP! DB dataset so to provide programmatic access to the impact measures of the same set of publications. As a result, all calculated impact measures are also accessible via a public REST API<sup>5</sup>. It supports retrieving impact scores for a given article or for a number of articles given a list of DOI identifiers. The API response includes all five impact scores in a simple JSON object, one for each requested DOI.

### 4 DISCUSSION

To highlight the fact that the different measures capture semantically diverse impact aspects we present, in Table 2, the pairwise

<sup>4</sup>BIP! DB dump, <https://doi.org/10.5281/zenodo.4386934>

<sup>5</sup>BIP! API documentation, <https://bip-api.imsi.athenarc.gr/documentation>

	iCC	CC	PR	AttRank	RAM
iCC	1	0.0985	-0.3468	0.3141	0.3042
CC		1	0.4144	0.4583	0.2774
PR			1	-0.0675	-0.2598
AttRank				1	0.9056
RAM					1

**Table 2: Top-1% pairwise correlations of impact measures.**

top- $k$  correlations<sup>6</sup> of the top-ranking 1,000,000 papers (corresponding roughly to the top-1% of papers) for each pair of impact measures we calculated. Intuitively, we expect to see high correlations between measures that capture the same impact aspect. In general, our findings confirm this intuition since the popularity measures (AttRank & RAM) appear to be highly correlated ( $\rho > 0.9$ ), while the influence ones (CC & PR) appear to have a moderate correlation ( $\rho > 0.4$ ). Of course, as discussed, there are differences between measures of the same aspect, thus we do not expect pairs of measures for the same impact aspect to correlate perfectly. In addition, in this experiment we only examine the correlation of the top-1% publications; the full set would reveal larger correlations for the measures of the same aspect (e.g., using the top 10% publications we measured  $\rho$  values greater than 0.6 for CC and PR). Finally, based on the same measurements, iCC at best correlates weakly ( $\rho < 0.4$ ) to other measures, supporting the intuition that it captures a distinct impact aspect (impulse).

Our data are openly available both on Zenodo and through an open API, so to facilitate their utilization by third-party research teams and to enable building useful services on top of them. Furthermore, the files in the repository are available in TSV format, to allow for easy editing and/or processing for import in various database management systems. Finally, in line with the need for any relevant applications to use fresh data, we plan to update the provided files regularly, taking into consideration the update rate of the integrated data sources.

There are many possible applications that can leverage the data provided by BIP! DB. For instance, academic search engines may utilize impact measures to rank publications based on different impact aspects<sup>7</sup>, while science monitors (like the Open Science Observatory [12]) may use them to produce intuitive reports and visualisations. Furthermore, publication-level impact measures can be propagated and aggregated to related entities (e.g., datasets, software packages, individual researchers) to quantify their expected impact, which may be useful for various applications (e.g., planning for officers in funding organisations, decision support for HR departments of research institutions). Finally, the calculated impact measures may be used as features in various machine learning applications that apply data mining on citation networks.

At this point, we would like to highlight that researchers should always have a large toolbox of impact measures available, in order to get the full picture about a publication's impact and to successfully secure themselves from various types of attacks and malpractices in the filed of research assessment. This is why we envisage to

<sup>6</sup>We use Spearman's top- $k$  rank correlation  $\rho_{min}$  as defined in [4] ( $\rho \in [-1, 1]$ , with  $-1, 0, 1$  indicating perfect inverse, no, and perfect correlation, respectively).

<sup>7</sup>Our academic search engine BIP! Finder [16], is an attempt in this direction.

continuously update and extend the BIP! DB dataset to always contain up-to-date data and to be inline with the latest developments in the field of research assessment and scientometrics. To this end, we plan to update our set of included measures with new ones performing better in terms of effectiveness, and to extend BIP! DB to include measures that reflect additional impact aspects. Finally, since scientific impact is not always entirely correlated with publication quality, we intend to also include measures that capture other aspects of a publication's merit (e.g., novelty, readability).

## 5 CONCLUSIONS

We presented BIP! DB, an open dataset containing various impact measures, calculated for hundreds of millions of scientific articles. Our dataset provides a multidimensional view of article impact, and thus may be potentially beneficial for many different applications and diverse stakeholders. Furthermore, we aim to deliver regular updates of our dataset in line with the updates of the data sources we use. Finally, in the future, we additionally plan to extend the published dataset not only with further impact measures, but also with other indicators capturing aspects other than the ones strictly related to scientific impact, such as readability and novelty.

## REFERENCES

- [1] Johan Bollen, Herbert Van de Sompel, Aric Hagberg, and Ryan Chute. 2009. A principal component analysis of 39 scientific impact measures. *PLoS one* 4, 6 (2009), e6022.
- [2] Lutz Bornmann and Rüdiger Mutz. 2015. Growth rates of modern science: a bibliometric analysis based on the number of publications and cited references. *JASIST* 66, 11 (2015), 2215–2222.
- [3] Peng Chen, Huafeng Xie, Sergej Maslov, and Sidney Redner. 2007. Finding scientific gems with Google's PageRank algorithm. *Journal of Informetrics* 1, 1 (2007), 8–15.
- [4] Ronald Fagin, Ravi Kumar, and Dakshinamurthi Sivakumar. 2003. Comparing top  $k$  lists. *SIAM Journal on discrete mathematics* 17, 1 (2003), 134–160.
- [5] Rumi Ghosh, Tsung-Ting Kuo, Chun-Nan Hsu, Shou-De Lin, and Kristina Lerman. 2011. Time-aware ranking in dynamic citation networks. In *Data Mining Workshops (ICDMW), 2011 IEEE 11th International Conference on*. IEEE, 373–380.
- [6] Ginny Hendricks, Dominika Tkaczyk, Jennifer Lin, and Patricia Feeney. 2020. Crossref: The sustainable source of community-owned scholarly metadata. 1, 1 (2020), 414–427. [https://doi.org/10.1162/qss\\_a\\_00022](https://doi.org/10.1162/qss_a_00022)
- [7] John PA Ioannidis. 2005. Why most published research findings are false. *PLoS medicine* 2, 8 (2005), e124.
- [8] Ilias Kanellos, Thanasis Vergoulis, Dimitris Sacharidis, Theodore Dalamagas, and Yannis Vassiliou. 2019. Impact-Based Ranking of Scientific Publications: A Survey and Experimental Evaluation. *IEEE TKDE* (2019).
- [9] Ilias Kanellos, Thanasis Vergoulis, Dimitris Sacharidis, Theodore Dalamagas, and Yannis Vassiliou. 2020. Ranking Papers by their Short-Term Scientific Impact. arXiv:2006.00951 [cs.DL]
- [10] Nan Ma, Jiancheng Guan, and Yi Zhao. 2008. Bringing PageRank to the citation analysis. *Information Processing & Management* 44, 2 (2008), 800–810.
- [11] Lawrence Page, Sergey Brin, Rajeev Motwani, and Terry Winograd. 1999. The PageRank citation ranking: Bringing order to the web. (1999).
- [12] George Papastefanatos, Elli Papadopoulou, Marios Meimaris, Antonis Lempesis, Stefania Martziou, Paolo Manghi, and Natalia Manola. 2020. Open Science Observatory: Monitoring Open Science in Europe. In *AIMinScience 2020 (Communications in Computer and Information Science, Vol. 1260)*. Springer, 341–346.
- [13] Silvio Peroni and David M. Shotton. 2020. OpenCitations, an infrastructure organization for open scholarship. *Quant. Sci. Stud.* 1, 1 (2020), 428–444.
- [14] Arnab Sinha, Zhihong Shen, Yang Song, Hao Ma, Darrin Eide, Bo-June (Paul) Hsu, and Kuansan Wang. 2015. An Overview of Microsoft Academic Service (MAS) and Applications. In *Proceedings of WWW '15 Companion*. 243–246.
- [15] Derek R Smith. 2009. A 30-year citation analysis of bibliometric trends at the Archives of Environmental Health, 1975–2004. *Archives of environmental & occupational health* 64, sup1 (2009), 43–54.
- [16] Thanasis Vergoulis, Serafeim Chatzopoulos, Ilias Kanellos, Panagiotis Deligiannis, Christos Tryfonopoulos, and Theodore Dalamagas. 2019. BIP! Finder: Facilitating Scientific Literature Search by Exploiting

Impact-Based Ranking. In *Proceedings of CIKM 2019*. ACM, 2937–2940. <https://doi.org/10.1145/3357384.3357850>

[17] Kuansan Wang, Zhihong Shen, Chiyuan Huang, Chieh-Han Wu, Yuxiao Dong, and Anshul Kanakia. 2020. Microsoft Academic Graph: When experts are not enough. *QSS* 1, 1 (2020), 396–413.