

Understanding the Impact of Early Citers on Long-Term Scientific Impact

Mayank Singh
Dept. of Computer Science and Engg.
IIT Kharagpur, India
mayank.singh@cse.iitkgp.ernet.in

Ajay Jaiswal
Dept. of Computer Science and Engg.
IIT Kharagpur, India
ajayjaiswal@iitkgp.ac.in

Priya Shree
Dept. of Computer Science and Engg.
IIT Kharagpur, India
priya.shree@iitkgp.ac.in

Arindam Pal
TCS Innovation Labs, India
arindam.pal1@tcs.com

Animesh Mukherjee
Dept. of Computer Science and Engg.
IIT Kharagpur, India
animeshm@cse.iitkgp.ernet.in

Pawan Goyal
Dept. of Computer Science and Engg.
IIT Kharagpur, India
pawang@cse.iitkgp.ernet.in

ABSTRACT

This paper explores an interesting new dimension to the challenging problem of predicting long-term scientific impact (*LTSI*) usually measured by the number of citations accumulated by a paper in the long-term. It is well known that early citations (within 1–2 years after publication) acquired by a paper positively affects its *LTSI*. However, there is no work that investigates if the set of authors who bring in these early citations to a paper also affect its *LTSI*. In this paper, we demonstrate for the first time, the impact of these authors whom we call *early citers* (EC) on the *LTSI* of a paper. Note that this study of the complex dynamics of EC introduces a brand new paradigm in citation behavior analysis. Using a massive computer science bibliographic dataset we identify two distinct categories of EC – we call those authors who have high overall publication/citation count in the dataset as *influential* and the rest of the authors as *non-influential*. We investigate three characteristic properties of EC and present an extensive analysis of how each category correlates with *LTSI* in terms of these properties. In contrast to popular perception, we find that influential EC negatively affects *LTSI* possibly owing to *attention stealing*. To motivate this, we present several representative examples from the dataset. A closer inspection of the collaboration network reveals that this stealing effect is more profound if an EC is nearer to the authors of the paper being investigated. As an intuitive use case, we show that incorporating EC properties in the state-of-the-art supervised citation prediction models leads to high performance margins. At the closing, we present an online portal to visualize EC statistics along with the prediction results for a given query paper. The portal is accessible online at: <http://www.cnergres.iitkgp.ac.in/earlyciters/>. To facilitate reproducible research, we shall soon make all the codes and the processed dataset available in the public domain.

CCS CONCEPTS

•Information systems →Data mining; Digital libraries and archives;

KEYWORDS

Long-term scientific impact, citation count, early citers, supervised regression models

ACM Reference format:

Mayank Singh, Ajay Jaiswal, Priya Shree, Arindam Pal, Animesh Mukherjee, and Pawan Goyal. 2017. Understanding the Impact of Early Citers on Long-Term Scientific Impact. In *Proceedings of The ACM/IEEE-CS Joint Conference on Digital Libraries, Toronto, Ontario, Canada, June 2017 (JCDL '17)*, 10 pages.
DOI: 10.475/1234

1 INTRODUCTION

Success of a research work is estimated by its scientific impact. Quantifying scientific impact through citation counts or metrics [2, 10, 12, 14] has received much attention in the last two decades. This is primarily owing to the exponential growth in the literature volume requiring the design of efficient impact metrics for policy making concerning with recruitment, promotion and funding of faculty positions, fellowships etc. Although these approaches are quite popular, they appear to be highly debatable [15, 17]. Additionally, they fail to take into account the future accomplishments of a researcher/article. A natural and intriguing question is – *why should one be concerned about the future accomplishments of a researcher/article?* When an early-career researcher is selected for a tenure-track position, it is an investment. More likely, an organization will largely invest on a researcher who has higher chances of accomplishing more in future. Similarly, to ensure high quality search/recommendation results, search engines can rank recently published articles (low cited) higher than older articles (highly cited), if there is some guarantee that the recent article is going to be popular in the near future.

Prediction of future citation counts is an extremely challenging task because of the nature and dynamics of citations [8, 23, 32]. Recent advancement in prediction of future citation counts has led to the development of complex mathematical and machine learning based models. The existing supervised models have employed several paper, venue and author centric features that can be obtained at the publication time. There are equally many works [3, 26, 28]

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.

JCDL '17, Toronto, Ontario, Canada

© 2017 ACM. 123-4567-24-567/08/06...\$15.00

DOI: 10.475/1234

that leverage citation information generated within 1–2 years after publication to enhance the prediction. Despite this enormous interest, the characteristics of early citations generated immediately after publications have not been dealt with in-depth. In particular, to the best of our knowledge, there is no work that has studied the effect of the early citing authors on the long-term scientific impact (*LTSI*). We would like to stress that here we identify this social process for the first time that introduces a new paradigm in citation behavior analysis.

The aim of this work is to better understand the complex nature of the *early citers* (EC) and study their influence on *LTSI*. EC represents the set of authors who cite an article early after its publication (within 1–2 years). We investigate three characteristic properties of EC and present an extensive analysis to answer three interesting research questions:

- Do early citers influence the future citation count of the paper?
- How do early citations from influential authors impact the future citation count compared to the non-influential ones?
- How do citations from co-authors impact the future citation count compared to the others (influential as well as non-influential)?

In Section 4, we present a large-scale empirical study to answer these questions. Motivated by the empirical observations, in Section 5, we incorporate the EC features in a popular citation prediction framework proposed by Yan et al. [32]. In Section 6, we discuss the prediction outcomes and show that our extended framework outperforms the original framework by a high margin. In particular, we make the following contributions:

- (1) We identify two important categories of EC – we call those authors that have high publication/citation count in the data as *influential* and the rest of the authors as *non-influential*.
- (2) We analyze three different characteristic properties of EC.
- (3) We empirically show that early citations might not be always beneficial; in particular early citations from influential EC negatively correlates with the *LTSI* of a paper.
- (4) We build a citation prediction model incorporating the EC features; the prediction outcomes by far outperforms the baseline predictions.
- (5) We construct an online portal to present visualization of EC statistics and prediction results for a given query paper.

2 EARLY (NON-)INFLUENTIAL CITERS

The term *early citations* refers to citations accumulated immediately after the publication. In the literature, although, there seems to be no general definition of ‘early’, majority of the works kept it within ~ 2 years after publication [1, 23]. Multiple previous works assert that early citation count helps in better prediction of the *LTSI* [1, 3, 8]. Although these approaches are interesting, they fail to capture the existence of different types of *early citations* leading to more complex influence patterns on *LTSI*.

Given a candidate paper P published in the year T , we are interested in the citation information generated within δ year(s) after publication, i.e., within the time interval $[T, T + \delta]$. For example, for $\delta = 2$, if an article is published in the year 2000, we look into the citation information generated till 2002. *Early citation count*

$ECC_\delta(P)$ refers to the total number of citations received by the paper P from other articles within δ years after publication. Note, $ECC_\delta(P)$ quantitatively measures the early popularity of the paper P . However, $ECC_\delta(P)$ fails to capture the inherent nature of the individual early citations; for example, there exists no distinction between:

- originators (authors, journals etc.) of early citations.
- good (substantiating) and bad (criticizing) citations.
- self and non-self citations.

To incorporate some of the above distinctive characteristics in $ECC_\delta(P)$ and to better understand the inherent nature of the individual citations, we present the following three definitions:

Early citers ($EC_\delta(P)$): $EC_\delta(P)$ represents the set of authors that cite paper P within δ years after its publication. Figure 1 shows schematic representation of $EC_\delta(P)$ on a temporal scale. Here, $EC_\delta(P)$ consists of all authors that cite paper P within δ year after its publication. Further, we divide this set into two subsets – i) *influential*, and ii) *non-influential* early citers.

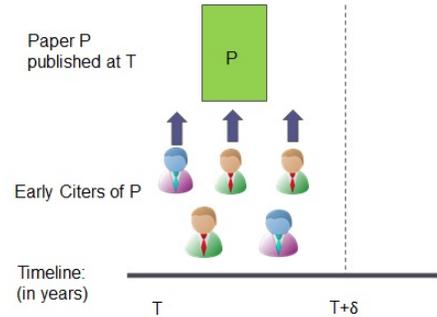


Figure 1: Schematic representation of early citers on a temporal scale. Early citers consist of all authors that cite paper P within δ year(s) after its publication. The set of early citers is divided into two subsets, namely, a) influential, and b) non-influential. Influential early citers are represented in purple color (online) whereas non-influential early citers are represented in green color (online).

Influential early citers ($IEC_\delta(P)$): This is a subset of $EC_\delta(P)$ in which each author either has a high publication count or a high citation count or both at the time of citation. Note that, in the current work, we consider top ~ 5% authors as influential early citers, both in terms of publication and citation counts. Empirically (from dataset described in Section 3), we find that top ~ 5% consists of authors who have authored at least 21 publications or acquired at least 250 citations or both. In Figure 1, for paper P , $IEC_\delta(P)$ are represented in the purple color.

Non-influential early citers ($NEC_\delta(P)$): Early citers that are not influential constitutes the set of non-influential citers, i.e.

$$NEC_\delta(P) = EC_\delta(P) \setminus IEC_\delta(P) \tag{1}$$

As described before, NEC_δ consists of the remaining ~ 95% of the authors in $EC_\delta(P)$. In figure 1, $NEC_\delta(P)$ authors are represented in green color. To study the impact of influential and non-influential EC on citations gained at a later point in time, we define long-term scientific impact as:

Long-term scientific impact ($LTSI_{\Delta}(P)$): Given a paper P , it represents cumulative citation count of P after Δ years of its publication. Section 4 demonstrates the effect of influential and non-influential EC on $LTSI$. Next, we describe the dataset we employ for the large scale experimental study and for the extended prediction framework.

3 DATASET DESCRIPTION

In this paper, we utilize two open source computer science datasets, both crawled from the Microsoft Academic Search (MAS)¹. First dataset (bibliographic dataset) was crawled by Chakraborty et al. [8] for a similar prediction work. The dataset consists of bibliographic information of more than 2.4 million papers, such as, the title, the abstract, the keywords, its author(s), the affiliation of the author(s), the year of publication, the publication venue, and the references. Second dataset (citation context dataset) was prepared by Singh et al. [23]. This dataset consists of more than 26 million citation contexts, pre-processed and annotated with the cited and the citing paper information. We combine the above two separately crawled datasets into a single *compiled dataset*.

We filter the compiled dataset by removing papers with incomplete information about the title, the abstract, the venue, the author(s), etc. Since the current study entirely focuses on early citers, we only include papers that consist of at least one citation within $\delta (= 2)$ years after publication. We term this dataset as *filtered dataset*. Table 1 outlines the various statistics for both the datasets. For the rest of this paper, we conduct all our experiments on the filtered dataset unless otherwise stated.

Table 1: General information about the datasets. We combine the two separately crawled datasets – a) the bibliographic dataset, and b) the citation context dataset into a single compiled dataset. We create the filtered dataset after removing incomplete information from the compiled dataset. Note, the filtered dataset consists of articles that have at least one citation within $\delta (= 2)$ years after publication.

	Compiled dataset	Filtered dataset
No. of publications	2,473,147	949,336
No. of authors	1,186,412	535,543
Year range	1859–2012	1970–2010
No. of citation contexts	26,037,804	11,532,780

4 EMPIRICAL STUDY

In this section, we plan to empirically investigate how the early citers impact the $LTSI$ of a paper. The section begins by introducing three properties of early citers, namely, the publication count, the citation count and the co-authorship distance. We describe each property in detail and present correlation statistics along with representative examples.

General Setting: Given a candidate paper P , we construct a set of early citing papers C_P that cite P within δ year(s) after publication. For the current study, we keep $\delta = 2$. From the definition presented in section 2, $EC_{\delta}(P)$ consists of all authors that have written papers present in C_P . Next, for each paper $c \in C_P$, we select

one representative author among all co-authors based on different selection criterion (described in Sections 4.1–4.3). More specifically, each selection criterion refers to one distinguishing property of EC. Further, we construct a representative author subset $REC_{\delta}(P)$ from the selected authors and present correlation statistics of this newly constructed subset with $LTSI$. Note that $REC_{\delta}(P) \subseteq EC_{\delta}(P)$. Next, we define the three key properties of EC that assist in distinguishing early citations.

4.1 Publication count

Publication count of an early citer refers to the number of articles written by her before citing the paper P . High publication count denotes high productivity of an early citer. For each paper $c \in C_P$, we select the author with the maximum publication count. The authors so selected constitute the set $REC_{\delta}(P)$. Note that in our experiments, authors with minimum, average and median publication counts have not shown significant correlations. Further, we aggregate early citers’ publication counts (PC_P) by averaging over the set of selected authors $REC_{\delta}(P)$. For each paper P present in our dataset, we compute PC_P and P ’s cumulative citation count at five later time periods after publication, $\Delta t = 5, 8, 10, 12, 15$. We utilize the definitions of influential and non-influential early citers described in section 2, i.e., a paper P is cited by a set of influential early citers, if $PC_P \geq 21$. Therefore, we split the entire paper set into two subsets: i) papers cited by non-influential EC ($PC_P < 21$), and ii) papers cited by influential EC ($PC_P \geq 21$). Figure 2 compares these two subsets correlating PC values with cumulative citation counts at five later time periods.

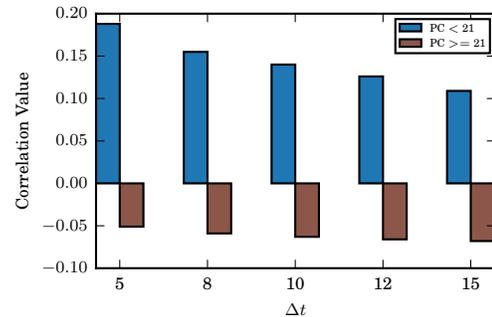


Figure 2: (Color online) Correlation between EC publication count and cumulative citation count at five later time periods after publication, $\Delta t = 5, 8, 10, 12, 15$. Papers with lower value of $PC (< 21)$ exhibit positive correlation diminishing over the time. Papers with high value of $PC (\geq 21)$ show an opposite trend. The overall separation decreases over time.

Observations: Figure 2 presents few interesting observations. Papers with lower value of $PC (< 21)$ exhibit positive correlation. However, as Δt progresses, this positive correlation starts diminishing. Surprisingly, papers with higher values of $PC (\geq 21)$, show negative correlation and this effect becomes more profound as Δt progresses. Thus, the overall separation between the two subsets decreases over time.

This study illustrates the fact that influential EC negatively affect the long-term citations. A plausible explanation could be that in general, researchers tend to cite works written by influential

¹<http://academic.research.microsoft.com>

authors. Therefore, once an influential author cites an article, researchers tend to cite the influential author’s paper, instead of the original paper. The attention from the original paper moves to the paper written by the influential citer toward the very beginning of the life-span of the original paper. Therefore, instead of flourishing, the long term citation count of the original paper gets negatively affected. This phenomenon of attention relaying from the less popular article to the more popular article is described as *attention stealing* [30]. In case of non-influential EC, the citation count of the candidate paper exhibits a positive correlation with PC. However, with the passage of time, this positive correlation diminishes due to ageing effect associated with paper’s life span [27]. In case of influential EC, same ageing effect leads to increase in the negative correlation over the passage of time.

Table 2 shows some specific examples of papers having the same early citation count in the first two years after publication but different PC values. In both cases, the paper having a low PC value receives a much higher citation count in the future.

Table 2: Example paper-pairs having a similar early citation count in the initial two years of publication but different PC values.

Paper ID	Early Citation Count	Early citer PC	Later Citation count
726084	13	18.9	79
140790	13	36.5	34
1663998	8	19.17	109
150167	8	65	38

4.2 Citation count

Citation count of an early citer refers to the number of citations received by her before citing paper P . High citation count denotes higher popularity of the early citer. Again, for each paper $c \in C_P$, we select the author with maximum citation count. Here again, the authors so selected constitute the set $REC_\delta(P)$. Further, we aggregate early citers’ citation counts (CC_P) by averaging over the set of selected authors $REC_\delta(P)$. For each paper P present in our dataset, we compute CC_P and P ’s cumulative citation count at five later time periods after publication, $\Delta t = 5, 8, 10, 12, 15$. Similar to previous section, we again split the entire paper set into two subsets: i) papers cited by non-influential EC ($CC_P < 250$), and ii) papers cited by influential EC ($CC_P \geq 250$). Figure 3 compares these two subsets by correlating CC values with the cumulative citation counts at five later time periods.

Observations: Figure 3 presents similar observations as reported in Figure 2. Papers with lower value of $CC (< 250)$ exhibit positive correlation diminishing over the time. Papers with high value of $CC (\geq 250)$ show an exactly opposite trend. Here also, the overall separation decreases with time. The results again confirm the existence of *attention stealing*, i.e. a popular citer steals the attention from a newly born paper by citing it. The temporal increase and decrease in correlation values of influential and non-influential early citers respectively relates to the ageing effect as discussed in the previous section.

Table 3 shows some specific examples of papers having the same early citation count in the first two years after publication but different CC values. Similar to publication count, here also, we

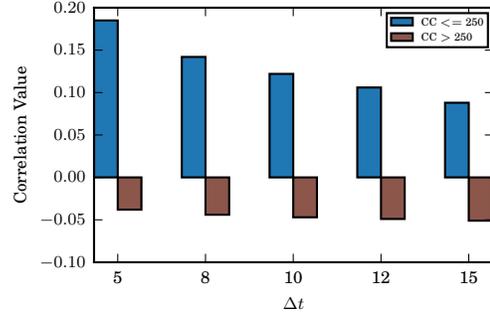


Figure 3: (Color online) Correlation between EC citation count and cumulative citation count at five later time periods after publication, $\Delta t = 5, 8, 10, 12, 15$. Papers with lower value of $CC (< 250)$ exhibit positive correlation diminishing over the time. Papers with high value of $CC (\geq 250)$ show an opposite trend. The overall separation decreases over time.

Table 3: Example paper-pairs having a similar early citation count in the initial two years of publication but different CC values.

Paper ID	Early Citation Count	Early citer CC	Later Citation count
2025205	4	124.75	51
287142	4	456	13
269672	18	74.45	61
1695635	18	623.17	29

observe that in both the cases, the paper having a low CC value receives a much higher citation count in the future.

4.3 Co-authorship distance

We construct a collaboration graph $G(V, E)$ to understand the effect of the co-authorship distance between EC and the authors of candidate paper P on $LTSI$. Here, V is the set of vertices representing authors and an edge $e \in E$ between two authors denotes that they have co-authored at least one article. We define the co-authorship distance (CA) between two authors as *the shortest distance between the two in the co-authorship network*. Again, for each paper $c \in C_P$, we select the author with the lowest CA from the authors of candidate paper P . The authors so selected constitutes the set $REC_\delta(P)$ here. Note that in our experiments, authors with highest, average and median co-authorship distance have not shown better correlations. We aggregate the co-authorship distance (CA_P) by averaging over the set of selected authors $REC_\delta(P)$. To understand the effect of co-authorship distance on $LTSI$, we divide CA into three buckets:

- **Bucket 1:** $0 \leq CA < 1$
- **Bucket 2:** $1 \leq CA < 2$
- **Bucket 3:** $CA \geq 2$

Note, $CA = 0$ represents self citations, i.e., one of the early citer is the author of the candidate paper P . The authors at $CA = 1$ are the co-authors of the authors in the candidate paper. Hence, **Bucket 1** mainly consists of authors of the candidate paper itself. **Bucket 2** mainly consists of the immediate co-authors of the author set of the candidate paper while **Bucket 3** mainly consists of co-authors of co-authors (distant neighbours) of the author set of the candidate paper.

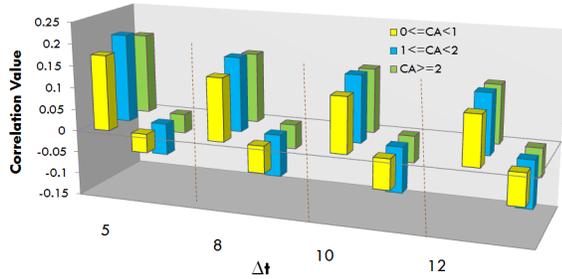


Figure 4: (Color online) Correlation between EC’s publication count and cumulative citation count for three co-authorship buckets at four later time periods after publication, $\Delta t = 5, 8, 10, 12$. For each time period, first three bars represent correlation for non-influential EC ($PC_P < 21$) whereas the next three bars represent correlation for influential EC ($PC_P \geq 21$). Influential immediate co-authors (Bucket 2) seem to badly affect the citation of the candidate paper P in the long term.

For each bucket, we present correlation statistics of EC’s publication count and citation count with $LTSI$. Figure 4 illustrates, for each bucket, correlation between EC’s publication count and cumulative citation count at four later time periods after publication, $\Delta t = 5, 8, 10, 12$. For each time period, the first three bars represent correlation for non-influential EC ($PC_P < 21$) whereas the next three bars represent correlation for influential EC ($PC_P \geq 21$). **Observations:** For each CA bucket, we observe similar trends as before, influential EC negatively affect the $LTSI$ while non-influential EC affect positively. The most striking observation from this experiment is the effect of immediate co-authors (**Bucket 2**) on $LTSI$. Even though, both influential or non-influential immediate co-authors maximally correlate with $LTSI$, influential immediate co-authors negatively affect the citation of the candidate paper P in the long term due to intensified *attention stealing* effect.

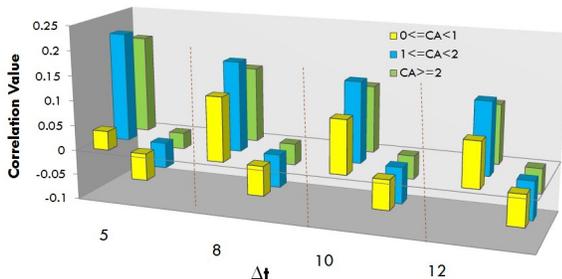


Figure 5: (Color online) Correlation between EC’s citation count and cumulative citation count for three co-authorship buckets at four later time periods after publication, $\Delta t = 5, 8, 10, 12$. For each time period, first three bars represent correlation for non-influential EC ($CC_P < 250$) whereas next three bars represent correlation for influential EC ($CC_P \geq 250$). Influential immediate co-authors (bucket 2) badly affect the attention of candidate paper P in long term.

Similarly, Figure 5 illustrates correlation between EC’s citation count and cumulative citation count at four later time periods after publication. For each time period, the first three bars represent

correlation for non-influential EC ($CC_P < 250$) whereas the next three bars represent correlation for influential EC ($CC_P \geq 250$).

Observations: In this case, the observations are very similar to the previous case. Motivated by these empirical observations, we incorporate the EC properties in a well recognized citation prediction framework as described in the next section.

5 CITATION PREDICTION FRAMEWORK

As an intuitive use case, we extend the long-term citation prediction framework proposed by [32] by including the three EC properties discussed in the previous sections. In addition, we also include two citation context based features proposed by Singh et al. [23]. Given a candidate paper, we predict its cumulative citation count at five different time-points ($\Delta t = 3, 5, 7, 9, 11$) after publication. Our citation prediction framework employs a set of features that can be computed at the time of publication plus a set of features that can be extracted from the citation information generated within two years after publication (section 5.1). We train four predictive models for comparative study, namely, linear regression, Gaussian process regression, classification and regression trees and support vector regression. We discuss each model briefly in Section 5.2. We compare our proposed prediction framework with three baselines in Section 5.3 using evaluation metrics outlined in section 5.4.

5.1 Feature definition

As described before, we utilize features available at the time of publication along with the features available within two years after publication. The feature set consists of 20 different features, out of which 14 features are available at the publication time, while the other six features utilize citation information generated within two years after publication. Features² available at the time of publication are the same as reported in [32]. Similarly early citation count and citation context features available after publication are same as reported in [23]. The entire feature set can be divided into seven categories: i) features based on early citer properties, ii) early citation count, iii) features based on paper information, iv) features based on author information, v) features based on venue information, vi) paper recency, and vii) features based on citation context. Given a candidate paper P published in the year T , we compute the following features:

5.1.1 Early citer centric features. Early citer centric features are computed within two years after the publication. Given a set of early citing papers C_P , we compute three features:

- (1) **Publication count (ECPC):** For each early citing article, we select the author with the maximum publication count. ECPC is computed by averaging this maximum publication count over all the early citing articles.
- (2) **Citation count (ECCC):** Here, for each early citing article, we select the author with the maximum citation count. ECCC is then computed by averaging this maximum citation count over all the early citing articles.
- (3) **Co-authorship distance (ECCA):** Here, we select the author with the minimum co-authorship distance from the authors of the candidate paper P . ECCA is computed by

²Some of these features might appear correlated; however, we use all of these in order to have a faithful reproduction of the model proposed in [32]

averaging this minimum co-authorship distance over all the early citing articles.

5.1.2 Early citation count (ECC). This feature simply includes the citation counts of paper P generated within the first two years after publication.

5.1.3 Paper centric features.

- (1) **Novelty (PCN):** Novelty measures the similarity between paper P and the other publications in the dataset. It is computed by measuring Kullback-Leibler divergence of an article against all its references. We assume that low similarity means high novelty and more novel article should attract more citations.
- (2) **Topic Rank (PCTR):** Topics are inferred from the paper title and abstract using unsupervised LDA. Each paper is assigned a topic and further each topic is ranked based on the average citations it has received.
- (3) **Diversity (PCD):** Diversity measures the breadth of an article inferred from its topic distribution. We measure diversity of an article by computing the entropy of the papers's topic distribution (see [32] for more details).

5.1.4 Author centric features.

- (1) **H-Index (ACHI):** H-index attempts to measure both the productivity and the impact of the published work of a researcher [14]. Yan et al. [32] observed high positive correlation between h-index and average citation counts of publications.
- (2) **Author rank (ACAR):** Author rank determines the "fame" of an author. Each author is assigned an author rank based on her current citation count. High rank authors have high citation counts.
- (3) **Past influence of authors (ACPI):** We measure the past influence of authors in two ways: previous (1) maximum citation counts, and (2) total citation counts. Previous maximum citation count of an author represents the citation count of author's most popular publication. Previous total citation count represents sum of the citation counts of all the author's publications.
- (4) **Productivity (ACP):** The more papers an author has published, the higher average citation counts she could expect. Productivity refers to the total number of articles published by an author.
- (5) **Sociality (ACS):** A widely connected author is more likely to be cited by her wide variety of co-authors. Sociality, thus, can be computed from the co-authorship network graph employing a formulation in a recursive form as in the PageRank algorithm.
- (6) **Authority (ACA):** A widely cited paper indicates peer acknowledgements, and hence indicates the 'authority' of its authors. We compute authority of paper in citation network graph using similar recursive algorithm as proposed for the sociality feature. The paper authority then is transmitted to all its authors.
- (7) **Versatility (ACV):** Versatility represents the topical breadth of an author. We measure the versatility of an author by computing the entropy of the author's topic distribution.

Higher versatility implies large volumes of audience from various research fields.

5.1.5 Venue centric features.

- (1) **Venue rank (VCVR):** The reputation of a venue relates to the volume of citations it receives. Similar to author rank, we rank venues based on its current citation count. High rank venues have high citation counts.
- (2) **Venue centrality (VCVC):** We create a venue connective graph $G(V, E)$ where V denotes the set of venues and the edges $e \in E$ denote the citing-cited relationships between venues. The in-degrees measure how many times a venue is cited by papers from other venues. Finally, venue centrality can be measured using a PageRank algorithm.
- (3) **Past influence of venues (VCPI):** Past influence of a venue is computed similar to the past influence of authors. As in the case of authors, we measure the past influence of venues in two ways: previous (1) maximum influence of venues, and (2) total influence of venues.

5.1.6 Recency (PR). Recency describes the temporal proximity of an article. It measures the age of a published article. The longer an article is published, the more citations it may receive.

5.1.7 Citation context centric features.

- (1) **Average countX (CCAC):** A high value of countX implies that the cited paper is referred multiple times by the citer paper in different sections of its text. Thus, cited paper might be quite relevant for citing paper. Singh et al. [23] argued that highly cited papers are cited more number of times in a single text.
- (2) **Average citeWords (CCAW):** Similar to countX, a high value of citeWords implies that the cited paper has been discussed in more details by the citer paper and therefore, cited paper might be quite relevant for the citing paper.

5.2 Predictive models

In this section, we describe four regression models. Each model is trained on features described in previous section. All models are trained using available implementations from the Weka toolkit [13].

5.2.1 Linear regression (LR). Linear regression is an approach to model the relationship between the dependent variable Y and one or more independent (explanatory) variables X . It attempts to model this relationship by fitting a linear equation to observed data. A linear regression line has an equation of the form:

$$Y = wX^T + b, \quad (2)$$

where Y is the dependent variable, X^T is a vector of explanatory variables, w is a vector of weights (parameters) of the linear regression and b represents the error. In the current work, we consider publication's predicted citation count to be the dependent variable and features (described in Section 5.1) are considered to be the explanatory variables.

5.2.2 Gaussian process regression (GPR). Due to the complex nature of the long-term citation impact estimation, it might well be the case that the dependent variable is a non-linear function of all the features used to represent the data. Gaussian processes [22]

provide formulations by which the prior information about the regression parameters can be easily encoded. This property makes them convenient for our problem formulation. Given a vector of input features X , the predicted citation counts $C(d)$ of the document d is:

$$C(d) = K(X, X^T)[K(X^T, X^T) + \sigma^2 I]^{-1} C(d^T), \quad (3)$$

where X^T is a matrix of feature vectors of the training set, K is a kernel function, I is the identity matrix, σ is the noise parameter and $C(d^T)$ is the vector of citation counts of the training set. Note, in our experiments, we keep $\sigma = 0.5$.

5.2.3 Classification and regression trees (CART). Classification and regression trees [4] are obtained by recursively partitioning the training data space and fitting a simple prediction model within each partition. As a result, the partitioning can be represented graphically as a decision tree. Regression trees are built for dependent variables (citation count in the present context) that take continuous or ordered discrete values, with prediction error typically measured by the squared difference between the observed and predicted values.

5.2.4 Support vector regression (SVR). Support vector regression [24] are derived from statistical learning theory and they work by solving a constrained quadratic problem where the convex objective function for minimization is given by the combination of a loss function with a regularization term. Support vector regression is the most common application form of SVMs. In the current study, we employ LIBSVM³ with default parameter settings. The best results were obtained for the linear kernel.

5.3 Baselines

5.3.1 Baseline I. The first baseline [32] is similar to our model except that it does not include any information generated after the publication. It includes paper, author and venue centric features along with recency.

5.3.2 Baseline II. The second baseline is similar to Baseline I plus one more feature – early citation counts. Chakraborty et al. [8] showed that inclusion of early citation counts enhances prediction accuracies mostly for the higher values of Δt .

5.3.3 Baseline III. In the third baseline, we include citation context centric features introduced by Singh et al. [23] to Baseline II. Thus, baseline III consists of paper, author, venue and citation context centric features along with recency and early citation count.

5.4 Evaluation metrics

5.4.1 Coefficient of determination (R^2). Coefficient of determination (R^2) [7] measures how well the data fits a statistical model of future outcome prediction. It determines the variability introduced by the statistical model. Let d be the document in the test document set D , we compute R^2 as:

$$R^2 = \frac{\sum_{d \in D} (C_p(d) - C_a(D))^2}{\sum_{d \in D} (C_a(d) - C_a(D))^2} \quad (4)$$

Here, $C_p(d)$ denotes the predicted citation count for document d . $C_a(D)$ denotes the mean of observed citation counts for the documents in D . $C_a(d)$ denotes actual citation count for document d . R^2 values range from 0 to 1. A larger value indicates better performance.

5.4.2 Pearson correlation coefficient (ρ). Pearson correlation coefficient (ρ) [18] measures the degree of linear dependence between two variables. Let d be the document in the test document set D , we compute ρ as:

$$\rho = \frac{\sum_{d \in D} (C_p(d) - C_p(D))(C_a(d) - C_a(D))}{\sqrt{\sum_{d \in D} (C_p(d) - C_p(D))^2} \sqrt{\sum_{d \in D} (C_a(d) - C_a(D))^2}} \quad (5)$$

Here, $C_p(d)$ and $C_a(d)$ represents predicted citation count and actual citation count of test document d respectively. $C_p(D)$ and $C_a(D)$ represent mean of the predicted and the observed citation counts for the documents in D . ρ ranges from -1 to 1, where $\rho = 1$ corresponds to a total positive correlation, 0 corresponds to no correlation, and -1 corresponds to total negative correlation. A larger value indicates better performance.

6 PREDICTION ANALYSIS

6.1 Experimental setup

Our experimental setup bears a close resemblance to [32]. We randomly select 10,000 training sample papers published in and before the year 1995. We opted for a small sample size because of associated computational complexities. Since, our prediction framework utilizes information generated within first two years after publication, we perform prediction task from 1998 – 2010. The reason behind choosing 1998 as the start year is to counter information leakage due to the training papers published at 1995 since prediction framework utilizes early citation data till 1997 for papers published in the year 1995. To evaluate, we select three random sets of 10,000 sample papers (published between 1998 – 2010). Note that for $\Delta t = 11$, we can only consider papers published between 1998 – 1999, for $\Delta t = 9$, we can consider papers published between 1998 – 2001 and so on. Given a candidate paper, we predict its cumulative citation count at five different time-points after publication, $\Delta t = 3, 5, 7, 9, 11$. For example, given a candidate paper P published in 1998, $\Delta t = 3$ represents prediction at 2001, $\Delta t = 5$ represents prediction at 2003 and so on. In the next section, we present a comprehensive analysis of our proposed framework.

6.2 Prediction results

6.2.1 Comparison between predictive models. Our model: To begin with, we incorporate all features described in section 5.1 for the prediction task (includes early citer centric, paper centric, author centric, venue centric, citation context centric features plus early citation count and recency features). However, we observe marginal performance gain in all models after removing the citation context based features. Therefore, it was decided that the best framework (hereafter ‘our model’) for this prediction task would consist of all features except the citation context based features. Table 4 compares the four predictive models (LR, GPR, CART and SVR) at five different time-points after publication, $\Delta t = 3, 5, 7, 9, 11$. Overall, SVR achieves the best performance, while GPR seems to

³<http://www.csie.ntu.edu.tw/~cjlin/libsvm/>

have the worst performance. As expected, in all the models, the performance diminishes as Δt increases.

Table 4: Performance comparison among the four predictive models – LR, GPR, CART and SVR. Two evaluation metrics R^2 and ρ are used. A high value of R^2 and ρ represent an efficient prediction. Prediction is performed over five time periods, $\Delta t = 3, 5, 7, 9, 11$.

Model	$\Delta T = 3$		$\Delta T = 5$		$\Delta T = 7$		$\Delta T = 9$		$\Delta T = 11$	
	ρ	R^2	ρ	R^2	ρ	R^2	ρ	R^2	ρ	R^2
LR	0.95	0.82	0.91	0.79	0.84	0.74	0.81	0.68	0.75	0.61
GPR	0.83	0.57	0.80	0.55	0.71	0.48	0.66	0.47	0.64	0.30
CART	0.95	0.73	0.87	0.68	0.79	0.62	0.75	0.55	0.63	0.43
SVR	0.97	0.84	0.91	0.82	0.88	0.76	0.82	0.69	0.76	0.65

6.2.2 Comparison with the baseline models. Next, we compare the performance of the three baselines (described in section 5.3) with our model. Due to high performance gain discussed in the previous section, we use SVR for modeling the three baselines as well as our model. Table 5 compares Baseline I, Baseline II and Baseline III with our model. Prediction is made over five time periods, $\Delta t = 3, 5, 7, 9, 11$. Each cell represents mean and standard deviation (in parenthesis) of the metric values for the three random samples. Even though, as highlighted, our model by far outperform all three baselines at each time period for both metrics, it slightly under estimates LTSI (see Figure 6).

6.2.3 Effect of different early time periods. So far, we have performed experiments for a fixed early time period ($\delta = 2$). In this section, we experiment with $\delta = 1, 2, 3$ for estimating the early citer features⁴. Table 6 compares the prediction results for the SVR model using three different values of δ . The table presents an interesting finding that increasing the value of δ does not always improve prediction accuracy. R^2 values at $\delta = 2$ always outperform $\delta = 1, 3$ in the later time points.

6.3 Feature analysis

We now study how the various features correlate with the actual citation counts. As described in Section 6.2.1, our model is trained on 18 features out of 20 features (described in Section 5.1); therefore, we perform feature analysis for 18 features. We train SVR with individual features and rank them based on Pearson’s correlation values of each feature with the actual citation count for $\Delta t = 3$ years after publication in descending order. Table 7 reports ranked list of features at $\Delta t = 3$. We can observe from the table that the first six in the rank list consists of all the three EC features, indicating importance of the EC features. As expected, early citation count is the most distinctive feature.

Figure 7 presents cross-correlation between features. Diagonal entries have maximum positive correlation (self) values = 1. Overall, features seem to be not much correlated with each other except a few cases. Interestingly, we observe that the EC features negatively correlate with the early citation count feature, the two being very distinct sources of information. Thus, including the EC features

⁴Note that the early citation count however is obtained using $\delta = 2$ as suggested in the literature.

enhances the prediction performance significantly over and above the early citation count feature.

7 ONLINE PORTAL

We have also built an online portal to showcase the different results from our current work. Given a query paper present in our dataset, the portal displays different statistics related to the paper; in particular, each query result is accompanied by the statistics of the EC properties and other paper details. In addition, the portal also presents with a visualization comparing the actual and the predicted citation count of the paper. The current system is hosted on our research group server and can be accessed at <http://www.cnergres.iitkgp.ac.in/earlyciters/>.

8 RELATED WORK

In recent years, several researchers have investigated the problem of *LTSI* [8, 23, 27, 32]. While some works propose complex mathematical models [21, 25, 27–29, 31] incorporating ageing assumptions, majority of the works focused on supervised machine learning models. Moreover, there are few recent works [3, 28] that present an empirical analysis of the correlation between short-term and long-term citation counts. Interestingly, Stern [26] reports that shortly after the appearance of a publication the combined use of early citations and impact factors yields a better prediction of the *LTSI* of the publication than the use of early citations only. Recently, Didegah et al. [9] presented an overview of the literature on predicting *LTSI*.

Mathematical models: The use of early citations to predict *LTSI* has been studied in various papers using mathematical models. Wang et al. [28] and Mingers et al. [21] proposed models that described how publications accumulate citations over the time. Stegehuis et al. [25] employed two predictor models (journal impact factor and early paper citations) to predict a probability distribution for the future citation count of a publication. They only considered accumulated citations within one year after publication. This is in contrast to the approach proposed by Wang et al. [27] where they allow predictions to be made fairly soon after the appearance of a publication. They propose three fundamental citation driving mechanisms – a) preferential attachment, b) ageing and novelty, and c) importance of a discovery. Their proposed model collapses the citation histories of papers from different journals and disciplines into a single curve indicating that all papers tend to follow the same universal temporal pattern. More recent work by Xiao et al. [31] explored paper-specific covariates and a point process model to account for the ageing effect and triggering role of recent citations.

Machine learning models: Among machine learning (ML) based prediction models, majority of the works have utilized support vector regression (SVR) [8, 23], classification and regression tree (CART) [6, 33] and linear and multiple regression models [16, 20]. Among ML models, we categorize works into three types based on the temporal availability of features – (a) features available at the time of publication [6, 11, 16, 19, 32], (b) features available after publication [5], and (c) combination of (a) and (b) [8, 23]. Callahan et al. [6] used features like journal impact factor, research design, number of subjects, rated subjectivity for scientific quality, news-worthiness etc. Further, they train decision trees to predict

Table 5: Performance comparison among Baseline I, Baseline II, Baseline III and our model. Two evaluation metrics ρ and R^2 are used. A high value of both metrics represent an efficient model. Prediction is made over five time periods, $\Delta t = 3, 5, 7, 9, 11$. Each cell represents mean and standard deviation (in parenthesis) of the metric values for three random samples. Bold numbers in the table indicate the best performing model for a given time period. Our model by far outperforms all three baselines at each time period for both metrics.

Δt	Baseline I		Baseline II		Baseline III		Our model	
	ρ	R^2	ρ	R^2	ρ	R^2	ρ	R^2
3	0.793 (0.003)	0.654 (0.019)	0.856 (0.021)	0.724 (0.001)	0.895 (0.012)	0.769 (0.017)	0.971 (0.002)	0.841 (0.001)
5	0.745 (0.021)	0.644 (0.006)	0.792 (0.007)	0.699 (0.012)	0.814 (0.019)	0.788 (0.001)	0.915 (0.015)	0.819 (0.019)
7	0.691 (0.016)	0.593 (0.003)	0.752 (0.004)	0.688 (0.019)	0.754 (0.023)	0.690 (0.026)	0.877 (0.007)	0.765 (0.013)
9	0.543 (0.008)	0.588 (0.015)	0.646 (0.009)	0.639 (0.002)	0.684 (0.002)	0.643 (0.001)	0.819 (0.003)	0.687 (0.021)
11	0.591 (0.015)	0.544 (0.002)	0.633 (0.010)	0.542 (0.006)	0.675 (0.008)	0.582 (0.021)	0.758 (0.005)	0.651 (0.016)

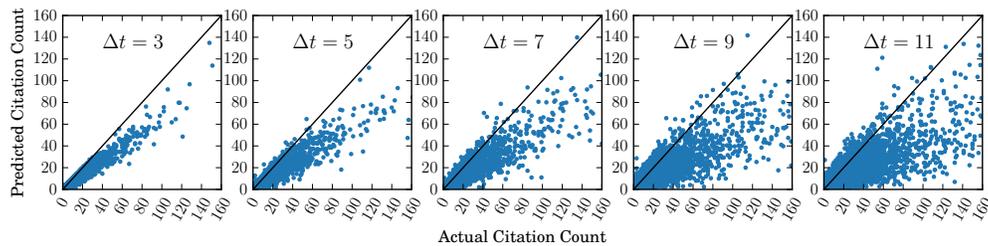


Figure 6: Change in prediction results over five time-periods. Scatter plots showing correlation between SVR predictions with real citation count values at $\Delta t = 3, 5, 7, 9, 11$. The black color line represents $y = x$ line passing through origin. Our model performs best for $\Delta T = 3$ with majority of the points on $y = x$ line. It performs worst for $\Delta T = 11$ with high divergence from the line. Our model under estimates *LTSI* as majority of the points lie below the line. However, this prediction is considerably better than all the other baselines.

Table 6: Performance of the model assuming different values of δ . Prediction is made over three early time periods, $\delta = 1, 2, 3$, and at three later time points, $\Delta t = 5, 7, 9$. Best results are obtained at $\delta = 2$. The added information does not always improve prediction accuracy.

ΔT	$\delta = 1$		$\delta = 2$		$\delta = 3$	
	ρ	R^2	ρ	R^2	ρ	R^2
5	0.882	0.68	0.915	0.82	0.911	0.76
7	0.841	0.61	0.877	0.77	0.884	0.72
9	0.765	0.58	0.819	0.69	0.822	0.64

Table 7: Ranked list of features based on Pearson’s correlation values between the predicted citation count and the actual citation count for $\Delta t = 3$ years after publication. Each SVR model is trained with individual feature.

1	ECC	6	ECCA	11	ACAR	16	PCN
2	ECCC	7	ACHI	12	ACP	17	ACV
3	ECPC	8	VCVR	13	PCTR	18	VCVC
4	VCPI	9	ACS	14	PR		
5	ACPI	10	PCD	15	ACA		

citation counts of 204 publications from emergency medicine specialty meeting. Livne et al. [19] used five group of features – authors, institutions, venue, references network and content similarity to train an SVR model. Similarly, Kulkarni et al. [16] also used information present at the publication time. They train linear regression to predict citation count for five year ahead window using 328 medical articles. Yan et al. [32] introduced features covering venue prestige,

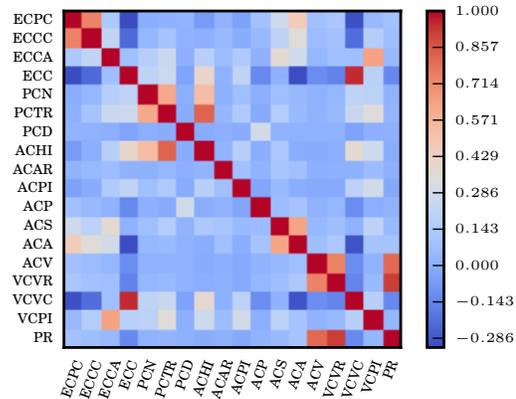


Figure 7: (Color online) Cross correlation between features: Red color represents highly correlated features (=1). Blue represents uncorrelated to weakly negatively correlated features. Diagonal entries have maximum correlation (self) values = 1.

content novelty and diversity, and authors’ influence and activity. Another work used data generated after the publication to predict citation count [5]. In this study, the downloaded data within the first six months after publication was used as a predictive feature. Chakraborty et al. [8] claimed that stratified learning approach leads to higher prediction accuracy. They proposed a two-stage prediction model that consumes information present at the publication time as well as citation information generated within the first two years after publication. Singh et al. [23] proposed extension to

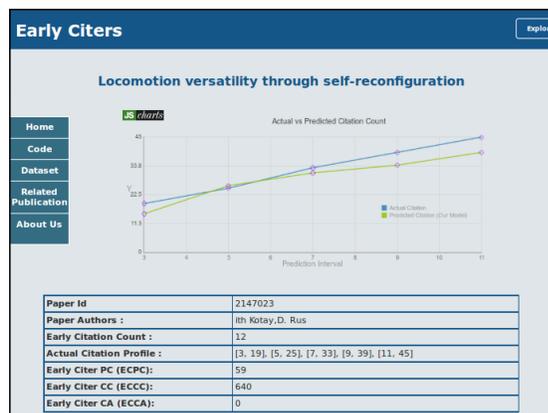


Figure 8: (Color online) Snapshot of online portal: For input candidate paper, the portal presents visualization of prediction results along with EC statistics. It compares SVR predictions with real values at $\Delta t = 3, 5, 7, 9, 11$ years after publication.

previous work [8] by including crowdsource based textual features like countX and citeWords.

9 CONCLUSION AND FUTURE WORK

This paper has investigated influence of early citers (EC) on long-term scientific impact. We have been successfully able to provide empirical evidence that early citers play a significant role in determining the long-term scientific impact. More specifically, we find that influential EC have a negative impact while non-influential EC have a positive impact on a paper's *LTSI*. We have provided further evidence that the negative impact is more intense when EC is closer to the authors of the candidate article in the collaboration network. Drawing from these observations, we incorporate the EC properties in a state-of-the-art supervised prediction model obtaining high performance gains. We believe that the identification of this social process actually leads to a new paradigm in citation behavior analysis.

In future, we believe that our work can be easily generalized for other scientific research fields. This study is the first step towards enhancing our understanding of influence of EC. To further our research we plan to analyze effects of EC in the patent datasets as well. Future work will concentrate on mathematical modeling of EC influence.

REFERENCES

- [1] Jonathan Adams. 2005. Early citation counts correlate with accumulated impact. *Scientometrics* 63, 3 (2005), 567–581. DOI: <http://dx.doi.org/10.1007/s11192-005-0228-9>
- [2] Carl T Bergstrom, Jevin D West, and Marc A Wiseman. 2008. The Eigenfactor? metrics. *The Journal of Neuroscience* 28, 45 (2008), 11433–11434.
- [3] Lutz Bornmann, Loet Leydesdorff, and Jian Wang. 2013. Which percentile-based approach should be preferred for calculating normalized citation impact values? An empirical comparison of five approaches including a newly developed citation-rank approach (P100). *Journal of Informetrics* 7, 4 (2013), 933–944.
- [4] Leo Breiman, Jerome Friedman, Charles J Stone, and Richard A Olshen. 1984. *Classification and regression trees*. CRC press.
- [5] Tim Brody, Stevan Harnad, and Leslie Carr. 2006. Earlier web usage statistics as predictors of later citation impact. *Journal of the American Society for Information Science and Technology* 57, 8 (2006), 1060–1072.
- [6] Michael Callahan, Robert L Wears, and Ellen Weber. 2002. Journal prestige, publication bias, and other characteristics associated with citation of published

- studies in peer-reviewed journals. *Jama* 287, 21 (2002), 2847–2850.
- [7] A Colin Cameron and Frank AG Windmeijer. 1997. An R-squared measure of goodness of fit for some common nonlinear regression models. *Journal of Econometrics* 77, 2 (1997), 329–342.
- [8] Tanmoy Chakraborty, Suhansanu Kumar, Pawan Goyal, Niloy Ganguly, and Animesh Mukherjee. 2014. Towards a Stratified Learning Approach to Predict Future Citation Counts. In *Proceedings of the 14th ACM/IEEE-CS Joint Conference on Digital Libraries (JCDL '14)*. IEEE Press, 351–360.
- [9] Fereshteh Didegah and Mike Thelwall. 2013. Which factors help authors produce the highest impact research? Collaboration, journal and document properties. *Journal of Informetrics* 7, 4 (2013), 861–873.
- [10] Leo Egghe. 2006. Theory and practise of the g-index. *Scientometrics* 69, 1 (2006), 131–152.
- [11] Lawrence D. Fu and Constantin Aliferis. 2008. Models for Predicting and Explaining Citation Count of Biomedical Articles. *PMC* 2008 (2008), 222–226.
- [12] Eugene Garfield. 1999. Journal impact factor: a brief review. *Canadian Medical Association Journal* 161, 8 (1999), 979–980.
- [13] Mark Hall, Eibe Frank, Geoffrey Holmes, Bernhard Pfahringer, Peter Reutemann, and Ian H Witten. 2009. The WEKA data mining software: an update. *ACM SIGKDD explorations newsletter* 11, 1 (2009), 10–18.
- [14] Jorge E Hirsch. 2005. An index to quantify an individual's scientific research output. *Proceedings of the National academy of Sciences of the United States of America* (2005), 16569–16572.
- [15] Jorge E Hirsch and Gualberto Buena-Casal. 2014. The meaning of the h-index. *International Journal of Clinical and Health Psychology* 14, 2 (2014), 161–164.
- [16] Abhaya V Kulkarni, Jason W Busse, and Ifat Shams. 2007. Characteristics associated with citation rate of the medical literature. *PLoS one* 2, 5 (2007), e403.
- [17] Cyril Labbé. 2010. Ike Antkare one of the great stars in the scientific firmament. *ISSI newsletter* 6, 2 (2010), 48–52.
- [18] Joseph Lee Rodgers and W Alan Nicewander. 1988. Thirteen ways to look at the correlation coefficient. *The American Statistician* 42, 1 (1988), 59–66.
- [19] Avishay Livne, Eytan Adar, Jaime Teevan, and Susan Dumais. 2013. Predicting citation counts using text and graph mining. In *Proc. the iConference 2013 Workshop on Computational Scientometrics: Theory and Applications*.
- [20] Cynthia Lokker, K Ann McKibbin, R James McKinlay, Nancy L Wilczynski, and R Brian Haynes. 2008. Prediction of citation counts for clinical articles at two years using data available within three weeks of publication: retrospective cohort study. *BMJ* 336, 7645 (2008), 655–657.
- [21] John Mingers. 2008. Exploring the dynamics of journal citations: modelling with S-curves. *Journal of the Operational Research Society* 59, 8 (2008), 1013–1025.
- [22] Carl Edward Rasmussen. 2006. Gaussian processes for machine learning. (2006).
- [23] Mayank Singh, Vikas Patidar, Suhansanu Kumar, Tanmoy Chakraborty, Animesh Mukherjee, and Pawan Goyal. 2015. The Role Of Citation Context In Predicting Long-Term Citation Profiles: An Experimental Study Based On A Massive Bibliographic Text Dataset. In *Proceedings of the 24th ACM International Conference on Information and Knowledge Management*. ACM, 1271–1280.
- [24] Alex Smola and Vladimir Vapnik. 1997. Support vector regression machines. *Advances in neural information processing systems* 9 (1997), 155–161.
- [25] Clara Stegehuis, Nelly Litvak, and Ludo Waltman. 2015. Predicting the long-term citation impact of recent publications. *Journal of informetrics* 9, 3 (2015), 642–657.
- [26] David I. Stern. 2014. High-Ranked Social Science Journal Articles Can Be Identified from Early Citation Information. *PLOS ONE* 9 (11 2014), 1–11. <http://dx.doi.org/10.1371/journal.pone.0112520>
- [27] Dashun Wang, Chaoming Song, and Albert-László Barabási. 2013. Quantifying long-term scientific impact. *Science* 342, 6154 (2013), 127–132.
- [28] Jian Wang. 2013. Citation time window choice for research impact evaluation. *Scientometrics* 94, 3 (2013), 851–872. DOI: <http://dx.doi.org/10.1007/s11192-012-0775-9>
- [29] Mingyang Wang, Guang Yu, and Daren Yu. 2009. Effect of the age of papers on the preferential attachment in citation networks. *Physica A: Statistical Mechanics and its Applications* 388, 19 (2009), 4273–4276. DOI: <http://dx.doi.org/10.1016/j.physa.2009.05.008>
- [30] Michafll Charles Waumans and Hugues Bersini. 2016. Genealogical Trees of Scientific Papers. *PLOS ONE* 11, 3 (03 2016), 1–15. DOI: <http://dx.doi.org/10.1371/journal.pone.0150588>
- [31] Shuai Xiao, Junchi Yan, Changsheng Li, Bo Jin, Xiangfeng Wang, Xiaokang Yang, Stephen M. Chu, and Hongyuan Zha. 2016. On Modeling and Predicting Individual Paper Citation Count over Time. In *Proceedings of the Twenty-Fifth International Joint Conference on Artificial Intelligence, IJCAI 2016, New York, NY, USA, 9-15 July 2016*. 2676–2682. <http://www.ijcai.org/Abstract/16/380>
- [32] Rui Yan, Congrui Huang, Jie Tang, Yan Zhang, and Xiaoming Li. 2012. To better stand on the shoulder of giants. In *Proceedings of the 12th ACM/IEEE-CS joint conference on Digital Libraries*. ACM, 51–60.
- [33] Rui Yan, Jie Tang, Xiaobing Liu, Dongdong Shan, and Xiaoming Li. 2011. Citation count prediction: learning to estimate future citations for literature. In *Proceedings of the 20th ACM international conference on Information and knowledge management*. ACM, 1247–1252.