

Efficient anomaly detection from medical signals and images with convolutional neural networks for Internet of medical things (IoMT) systems

G.Srinivasa Rao, K.Upendrarao, N.Mownika

Associate Professor1, Assistant Professor2,3

Dept. of CSE,

mail-id:g.sr.anurag@gmail.com,kakanaboyina57@gmail.com, mounika.nekkanti@gmail.com

Anurag Engineering College,Anatagiri(V&M),Suryapet(Dt),Telangana-508206

Abstract

The question of which type of Computer Science (CS) publication— conference or journal – is likely to result in more citations for a published paper is addressed. A series of datasets are examined and joined in order to analyse the citations of over 195,000 conference papers and 108,000 journal papers. Two means of evaluating the citations of journals and conferences are explored, h5 and average citations per paper; it is found that h5 has certain biases that make it a difficult measure to use (despite it being the main measure used by Google Scholar). Results from the analysis show that CS, as a discipline, values conferences as a publication venue more highly than any other academic field of study. The analysis also shows that a small number of elite CS conferences have the highest average paper citation rate of any publication type, although overall, citation rates in conferences are no higher than in journals. It is also shown that the length of a paper is correlated with citation rate.

Introduction

In many fields of academic study, a journal paper is the established destination for reporting research. In Computer Science (CS), however, conferences are also valued as a dissemination route. Unlike most other fields of academia, CS academics often view a conference publication as the final venue to describe a piece of work. Within CS, there are around fifty conferences, where having a paper accepted is a notable mark of academic success. The papers are typically between 4,000-7,000 words in length, which enables sufficient detail of the work to be reported. Although CS journals are often willing to publish extensions of conference papers, often CS academics choose not to do this. As a consequence, a number of the CS conferences have become the sole record of research advances.

Compared to other academic fields, this is an unusual publication practise, which has for many years prompted a debate on which is the better publication forum: conferences or journals? In past work, Vardi argued that despite the popularity of conferences for publication, the reviewing quality was not good enough (Vardi, 2009). Fortnow agreed stating that the discipline of CS should “*grow up*” and start publishing in journals where submission deadlines and page limits would not impact on the quality of what was published (Fortnow, 2009). Van Hemert claimed that the reviewing time of CS Journals was far slower than conferences, which resulted in conference being preferred by academics (van Hemert, 2009). Vardi, in reply, called for changes to the journal reviewing process (Vardi, 2010).

One reason why the debate has not been resolved has been a lack of data to measure and compare the relative worth of each publication type. There are well established citation based measures for journals such as two and five year impact factors, as well as the more recent Eigenfactor (Bergstrom, 2007). The custodians of these measures have not focussed attention on applying these measures to conferences. The reasons for this appear to be that in many scientific disciplines, conferences are not valued, and the quality of conference bibliographic metadata is often poorer than it is for journals, making it harder to measure conference impact.

Relatively recently, services, such as Google Scholar¹ (GS) and Microsoft Academic Search² (MAS) have overcome some of these metadata problems enabling a more detailed comparison of the citation rates from journals and conferences. While, there have been past studies examining such questions (Rahm & Thor, 2005), there have been shortcomings of that work, which make comparisons less valid. Google Scholar recently introduced its Metrics (GSM) service, which enables comparison of journal and conference citations. However, the service uses a measure that while well-established, has not been used to assess the citation impact of conferences or journals. It was decided therefore to undertake a study to compare the citation rates of journals and conferences, exploring the following research questions:

- Are the publication practices of CS academics different from other academics?
- Do CS conferences have a stronger citation impact than journals?
- How does such a comparison vary between the two across different quality journals and conferences?
- Do different lengths of paper get different levels of citation?

The paper starts with a summary of past work in this area, followed a description of the data sets used in the study as well as the methodology and measures of our work. The paper next details the results of our study before discussing the results and finally concluding.

Past work

Measuring citation has long been used as a means to evaluate publications. Garfield's *impact factor* is probably the best known (Garfield, 1972). Citations to the papers in a publication are measured over a particular period of time (typically two or five years), the higher the aggregate citation count, the higher the publication's impact. There is an assumption that a paper appearing in a high impact publication will over time receive more citations, therefore, such publications tend to be the most desirable venues for academic work. Larivière and Gingras, using a novel methodology, confirmed this assumption showing that the impact factor of a journal significantly affects the number of citations a paper receives (Larivière & Gingras, 2010). See (Leydesdorff, Radicchi, Bornmann, Castellano, & Nooy, 2013) for a review, test, and comparison of proposed alternatives to impact factor.

Impact is commonly calculated for journals only. Though there is no methodological reason for the measure not to be computed for conferences, a lack of coverage of these publications in the traditional citation databases, such as Web of Science³ and Scopus⁴, prevents their impact from being calculated. This lack of coverage has been noted several times (Meho & Yang, 2007; Sanderson, 2008); it presents CS academics with a challenge. Said academics assert that CS conferences have a strong impact – indeed there is empirical evidence to support this (Bar-Ilan, 2010) – however, a lack of metrics supported by the citation databases makes it hard to persuade non-CS academics that conference papers have a value. Measures such as the acceptance rate of a conference are sometimes used as a proxy of quality, but cannot be compared with impact.

Researchers have tried to compare the two types of publications. Rahm and Thor examined two highly regarded database conferences and three equally important journals (Rahm & Thor, 2005). The researchers found that the conference papers had, on average, higher citations than journal papers. In contrast, Franceschet conducted an analysis showing that while most academics publish the majority of their work in

conferences, the majority of citations come from their journal publications (Franceschet, 2010). The reason for the opposing conclusions between these two works appears to be the data source used for citations: Rahm and Thor used Google Scholar, Franceschet used Web of Science, which has strong journal citation coverage.

Another citation and publication source is the ACM Digital Library⁵ holding both journal and conference papers with broad citation information. The library details (self-reported) acceptance rates for many CS conferences. Chen and Konstan (Chen & Konstan, 2010) examined 24,294 papers appearing in 2,108 conferences and journals. They showed that full papers⁶ published in low acceptance rate conferences (1 in 3 or lower) were cited more than papers published in a basket of CS journals. Comparing a small sub-set of good quality conferences with all the journals in the sample

is only valid if the journals in the sample are of a similar high quality. Such an assumption was not tested.

Taking both journal and conference publication quality into account was the focus of an examination of citation rates of 8,000 papers published in 32 CS conferences and journals covering artificial intelligence, machine learning, information retrieval, and computational linguistics (Freyne, Coyle, Smyth, & Cunningham, 2010). Unlike Chen et al, Freyne et al stratified the publications by quality based on data from the Australian government's research assessment exercise *Excellence in Research for Australia* (ERA). Through consultation with academic associations and rounds of wider feedback, ERA employed subject specialists to label thousands of journals and conferences with one of four quality grades (A*, A, B, or C). The researchers found that the most cited conferences were equal in citation rates to journals rated as A in ERA. The journals rated A* were found to have higher average citation rates than any conference in the group studied.

This brief survey of past work shows contrasting conclusions on the question of which produces more citations: journals or conferences. The reasons for the difference appear to be due to limited data sources for each study. It was concluded that a larger study would provide more reliable results.

Is CS publication practise different to other fields?

To answer the first research question, the ERA (2010) list of 20,712 journal and 1,952 conference titles – created to assess the research output of Australian academics – was examined. The list could be reasonably inferred to be a good representation of publication venues of each field of academic research. Every publication in the list was assigned at least one Field of Research (FoR) codes: six digit numbers that detail a three level hierarchy of research areas. As described above, the organisers of ERA, after consulting with a wide range of Australian academic groups, manually graded the journals and conferences for quality. Four grades were used for journals (A*, A, B, C) and three for conferences (A, B, C)⁷.

A comparison was made between the practises of CS (publications coded '08') and engineering (coded '09'). The percentage of all journals and conferences that were assigned a primary CS or engineering FoR code was measured and tabulated (Table 1 and Table 2). The CS journals represented 4% of all the journals ERA considered. The number of engineering journals was larger (7%), but this was to be expected as the field is broader than computer science. For both CS and engineering, there were fewer high quality journals than low. This trend was largely in line with the distribution of grades in the rest of the ERA list.

Data, methodology and measures

In this section, the data sets and the methodologies used in the paper are described.

Data

The aim of this study was to examine the widest possible list of CS journals and conferences, publication quality information, lists of papers published in the journals and conferences, paper citation counts, and page length information. The data sets chosen were as follows:

ERA 2010 conference and journal list. As described in Section 3, ERA compiled a list of 1,952 conferences and 20,712 journals, each of which was manually assigned a primary FoR code. The 1,489 CS conferences and 744 journals from the ERA 2010 list were extracted.

Microsoft Academic Search. Lists of papers and citation data were taken from Microsoft Academic Search (MAS). According to (Roy et al., 2013), in 2013, MAS held 50 million papers, written by 19 million authors. Rahm showed that an early version of MAS (called Microsoft Libra) had good coverage of citations to CS conferences compared to CS journals (Rahm, 2008). Jacsó (Jacsó, 2011) also states that computer science coverage (including conferences) is the most developed in MAS. An API is provided by MAS to its data and search functions. There are some challenges in joining the ERA list with the data in MAS due to variations in the names of the conferences and journals. Means of dealing with this are described in Section 5.2.1. Several hundreds of thousands of papers were collected from MAS covering the years 2007-2011.

The DBLP Computer Science Bibliography. Unfortunately, page length information was not easily available from MAS. The data set from the DBLP Computer Science Bibliography (Ley, 2009) is accurate, structured, and free to download. Therefore, the title of each listed paper was looked up in DBLP to obtain page length information.

ERA 2010/CORE 2008 conference and journal ratings. Each publication in the ERA 2010 list was manually assigned a rating. The journals were given one of four categories: A*, A, B, and C⁸. Conferences in the 2010 list were placed into three categories: A, B, and C. Unfortunately, these ratings fail to capture information on a small group of elite conferences that are often viewed by the CS community as exceptionally strong. In 2008, a list of elite conferences was compiled by CORE (an Australian academic CS association) using similar criteria to ERA and using all four of the journal grades. Around 5% of conferences were assigned an A*. Therefore, for those conferences rated A* in the CORE ratings, the CORE 2008 rating subsumed the ERA 2010 rating.

Methodologies to extract paper, citation, and page length information

To extract all of the required data, it was necessary to join the different data sets, which is described here.

Joining the ERA and MAS data

An algorithm was developed to match the conference and journal names from the ERA list to those within MAS. Broadly, the algorithm worked by submitting a query, via the MAS API, based on the name of the conference or journal, and examining the returned results. The name of each returned conference or journal was compared to the original query using a string matching metric. The closest match was considered the MAS conference or journal entry that matched to the original ERA listed name. In a pilot study, it was found that matching on journal names was 42% more accurate than matching on conferences. Therefore, before running the matching algorithm against our entire data set, a random set of 250 conference names were chosen to manually test and tune the matching process.

The first version of our algorithm involved submitting, as a query, the full ERA name of the 250 test conferences to MAS. If more than a single result was returned, the one with the lowest Edit Distance score (compared to the query) was considered to be the matching conference. Testing this version of the algorithm 127 (50.8%) of the conference names matched correctly; 106 (42.4%) did not match to anything; and 17 (6.8%) matched the wrong conference.

The ERA list also provided a short name, typically a unique acronym, for each conference, which was tested as a query to MAS, again comparing the short names of the results using Edit Distance. Now, 134 (53.6%) names were found to be correct matches; 87 (34.8%) no matches; and 29 (11.6%) mismatches. Combining the full and short name search results for a conference name and selecting the match with the shortest Edit Distance gave 152 (60.8%) correct matches; 73 (29.2%) no matches; and 25 (10%) mismatches.

An examination of search results revealed two problems with the algorithm. First, providing the full name of a conference proved restrictive in that commonly occurring words within the name of a conference sometimes resulted in the MAS not finding a match. For example, the conference listed in ERA as “*International Conference on Intelligent Tutoring Systems*” existed within the MAS database only as “*Intelligent Tutoring Systems*”; providing the full ERA name resulted in the MAS not finding a match. Second, Edit Distance takes into account the order of the words and letters in a name, but word order is perhaps less important in conference or journal name matching.

To solve the first problem, the most frequent terms from a conference name were removed before searching MAS. The full list of words that were removed from the conference and journal names was: a, ACM, an, and, annual, conference, for, IEEE, IFIP, international, journal, of, on, SPIE, symposium, the, to, with⁹.

The string comparison metric was changed to the Jaccard Measure, (Jaccard, 1912) which computes a simple word overlap, ignoring word order. Based on an analysis of the sample, a Jaccard score threshold of ≥ 0.2 was required before a match was judged to have occurred. The results of the keyword and short name search algorithm including this threshold were: 167 (66.8%) correct matches; 76 (30.4%) no matches; and 7 (2.8%) mismatches.

The results produced by this final combination of keyword and short name search were considered acceptable and this version of the algorithm was used to search for all the CS conferences and journals from the ERA list. The number of identified conferences was 1,021 and the number of journals was 568, which was manually checked for further errors. In addition, any conferences or journals for which MAS returned fewer than ten papers were discarded to try to avoid citation measures being overly influenced by outliers. This left 630 conferences and 290 journals, within which, there were 195,513 and 108,600 papers respectively.

While undoubtedly, some mismatches exist in the final data set, no evidence was found during the checking of mismatches that would affect the comparison of journals against conferences, the core research question of this paper. Note, this data will be made available for download for other to work with.

Determining paper length

Publications that accept short papers typically reduce reviewing criteria for such papers. As conferences, in particular, can accept a large number of such papers, it was decided to separate such short papers out. An algorithm was developed to look up the papers in DBLP based on their titles and year of publication. First a case-normalisation on the titles in both the candidate and DBLP papers was performed. The DBLP paper with the shortest Edit Distance, published in the same year as a query paper, was selected. Edit Distance was used over Jaccard as here word order was important. Matches with Edit Distance ≤ 5 were selected. The algorithm was applied to a set of 200 randomly chosen papers and the matches were manually inspected. The algorithm was found to correctly match 183 (91.5%) of the papers. The algorithm also ‘correctly missed’ 12 (6%) of the papers. For 5 (2.5%) of the papers, the algorithm matched the wrong paper. Such results were considered acceptable and no further refinement was made.

With the lengths of each paper established, the papers were assigned to one of three groups: short papers (with between one and four pages); long papers (page length greater than or equal to five); and unknown (papers for which page length could not be determined). A lack of page length information was either due to DBLP not having that information, or the paper title matching algorithm failing to find a match. The number of papers in each group is

shown in Table 3. The number of short papers for conferences was notably higher than the number of short papers for journals.

Discussion

This work studied the citation impact of CS conferences vs. journals. Such research questions were examined before, however, past work did not explore this topic on the scale or breadth conducted here. Rahm and Thor (Rahm & Thor, 2005) and Freyne et al (Freyne et al., 2010) each used fewer than fifty publications from particular sub-areas of CS. The wide standard deviations measured over the MAS dataset in this paper, indicate that using a large data set was necessary to obtain reliable statistics. Chen and Konstan (Chen & Konstan, 2010) showed that conferences with low acceptance rates had higher average citation rates than a wide basket of journals, but the lack of stratification of the journals made comparison between their work and that presented here hard.

Past work does not appear to have studied page length in the way that was conducted here. As can be seen, the graph (see Figure 10 and Figure 11) showed a link between page length and citation. The curves in the two graphs also appeared to be somewhat different with short conference papers receiving consistently low citations and only longer papers received more notable levels of citation. Causes for this apparent pattern are likely to be multiple. Longer papers on average will report more results than shorter papers; longer papers tend to be reviewed more carefully and so should represent more substantial (more citable) work; shorter papers might cover published materials that are less likely to be cited, such as abstract of keynotes; or letters to the editor.

The examination of the balance between conferences and journals in the CS ERA lists (described in Section 4) showed that the venues that a group of CS academics wished to be considered in a research evaluation exercise was dominated by conferences. However, a great many of the venues were weak. It can be seen in Table 2 that there are five times more B and C conferences than there are A. If one were to measure citations per paper across the ERA listed CS conferences regardless of grade and compared that to citations per paper across all ERA journals, journals have a higher average citation rate.

It was only when the quality of publication types was considered, that some citation rates in conferences were measured to be higher than journals. Here, A* conferences had a higher average citation rate (36% higher) compared to A* journals. The difference was statistically significant, though due to high variability of citation rates for conferences, the level of significance was not high. For A journals and conferences there appeared to be little difference and for B and C grade publications, journals had a higher citation rate.

Conclusions

Research was described in this paper that addressed the following research questions:

- Are the publication practises of CS academics different from other academics?
- Do computer science conferences have a stronger citation impact than journals?
- How does such a comparison vary between the two across different quality journals and conferences?
- Do different lengths of paper get different levels of citation?

Based on the data analysed in this study, the overall answer to the first question was yes, CS publication practises are very different from other academics. The answer to the second research question is that there is little difference between journals and conferences in terms of citation. However, once journals and conferences were stratified based on a manual assessment of the quality of each publication (the third research question), it was found that for a small number of elite conferences, average citation rates were notably higher than for elite journals. Papers in lower quality journals have more citations than papers in lower quality conferences. The length of a paper was found to have an impact on citation rate. The relationship between length and citation exists across all ranges of page length.

The fact that there continues to be a debate in CS on publication in conference vs. journals perhaps indicates that the rest of academia (as represented on tenure, promotion, or grant award committees) is not yet convinced of the value of conference publication in CS. The uniqueness of the publication profile of CS academics (see Section 4) is perhaps not fully understood by the rest of academia. However, the apparent value that CS academics place on publication in a large number of low grade conferences perhaps weakens the arguments of CS academics on the clear importance of publication in the elite CS conferences.

Perhaps a way to resolve the debate with others on whether journals or conferences are better is for CS academics to not try to suggest that all CS conferences are good publication venues, and instead the field should focus on promoting the virtues of publication in only the very best such forums. The remaining lower grade venues should be viewed like conferences in any other academic field of study: valuable meeting places for interaction and exchange of ideas, but as a suitable place for publishing the final word on a piece of work, their value is questionable.

A note of caution

While the length of the paper, the type of publication (conference or journal) or the quality of the publication it is published in does appear to influence the number of citations, behind the averages computed in this study, was substantial variation. This has been pointed out in the past, Seglen stated that many papers published in high impact journals are poorly cited (Seglen, 1997). Lozano et al presented results indicating that journal impact factors have a weaker influence on citation than they did in the past (Lozano, Larivière, & Gingras, 2012). While the nature of this variation was not examined in this paper, we speculate that it is simply due to the influence of the quality of the work reported in each individual paper. While attempting to get a paper accepted in the very best type of publication has been shown to influence subsequent citation rates, the principle means of influencing citations is the quality and clarity of the paper's content.

Acknowledgements

Many thanks to Dr. Tim Cahill and the paper's reviewers for reading and providing valuable comments on early drafts of the paper.

References

- John Bar-Ilan (2010). Computer science as a test scenario for combining Web of Science with citation indexes from conference proceedings *Scientific Analysis*, 83(3), 809.
- C. Bergstrom (2007). The Eigenfactor is a quantitative indicator of a journal's impact and citation impact. *Journal of Collegiate & Research Libraries* 68(5):314–316.
- Chen, J., and J. A. Konstan (2010). Selectivity and significance of conference papers. *Doi:10.1145/1743546.1743569 Communications of the ACM* 53(6), 79–83
- L. Fortnow (2009). Viewpoint It's high time for the field of computing to enter adulthood. *The ACM Companion*, 52(8), 33–35.
- A. Franceschet and M. (2010). *Commun. ACM*. 53(12), 129-132. *doi:10.1145/1859204.1859234*. The Importance of Conference Publications in Computer Science.
- John Freyne; Laura Coyle; Brian Smyth; and Patricia Cunningham (2010). In computer science, the significance of articles published in journals vs those presented at conferences is debatable. *The Communications of the ACM*, 53(11), pages 124-132.
- As cited in Garfield (E. (1972). Periodicals in the field of science policy may be evaluated using citation analysis by determining how often and how significantly each journal is cited. 178(4060), 471-479 in *Science*.
- It was written by A.-W. Harzing and R. van der Wal (2009). Journal influence in the fields of economics and business may also be measured using a Google Scholar h-index. 60(1), 41-46. *Journal of the American Society for Information Science and Technology*.
- J. E. Hirsch (2005). A measure of one's contribution to the field of scientific inquiry. The results of this study were published in *PNAS* 102(46), pages 16569–16572.
- According to Jaccard, P. (1912). Where plants grow in the alpine environment. *The New Phytologist*, 11(2) Jacsó, P. (2011). The bibliometric merits and flaws of Microsoft Academic Search. *Digital Library World*, 35(6), 987-997.
- V. Larivière & Y. Gingras (2010). Matthew Effect in bibliometrics, is the correlation between citations and impact factors.
- It was published in the 61(2) issue of the *Journal of the American Society for Information Science and Technology*.
- Ley, M. (2009). (2009). Insights gained by using the DBLP. Two-and-a-half volume of the VLDB Endowment Proceedings, pages 1493-1500.
- Published in: Leydesdorff, L., Radicchi, F., Bornmann, L., Castellano, C., and Nooy, W. (2013). Impact factors (IFs) standardized to a field: a comparison of rescaling and fractional counts. *JASIS&T*, 64(11), 22992309. *Journal of the American Society for Information Science and Technology*.
- Gingras, Y., Lozano, G. A., and Larivière, V. (2012). The decline in significance of the impact factor as a measure of the importance of a publication as measured by the number of times its articles are cited. 63(11), pp. 2140-2145 in *JASIST: Journal of the American Society for Information Science and Technology*. (Meho, L. I., & Yang, K. (2007). Differences in LIS professor citation statistics across Web of Science, Scopus, and Google Scholar. The 58(13) issue of the *Journal of the American Society for Information Science and Technology* has articles with page numbers of 2105 and 2125.
- Rahm, E. (2008). Examining the significance of computer science conference papers against those published in academic journals.

28(2), 127-128 in *Information Services and Usage*.

E. Rahm and A. Thor (2005). Documentation analysis based on citation counts in a database. Specifically, the page numbers for this issue of the *ACM Sigmod Record* are: 34(4), 48-53. De Cock, M., Roy, S. B., Mandava, V.,

Savanna, S., Dalessandro, B., Perlich, C.,.... There is a new book out by Hamner, B. (2013). Both the 2013 KDD Cup and the Microsoft Academic Search Dataset are available. This work was published as part of the 2013 KDD Cup Workshop Proceedings (p. 1).

Researchers Sanderson, M. (2008). Measuring h again among UK scholars in LIS and IR. *J. Amer. Soc. Inf. Sci. Technol.*, 59(7), 1184-1190. doi:<http://dx.doi.org/10.1002/asi.v59:7>

The authorship of Seglen, P. O. (1997). There are several reasons why journal impact factors should not be used to rank research. *British Medical Journal*, 314, no. 7079, 498.

Journal of Law and Public Policy [by] Van Hemert, J. (2009). Communicating with the editor via letters. This article may be found in the 52(11) issue of *Communications of the ACM* on pages 6-7.

This is the work of Vardi, M. (2009). Research in the field of computers: conferences vs. journals. *ACM Transactions on Computers*, 52(54), 5.

Authors: Vardi, M. (2010). Taking a look back at the norms of publishing in the field of computer science. *Proceedings of the ACM*, Volume 53, Issue 4 (3),5.