ArtSim: Improved estimation of current impact for recent articles

Serafeim Chatzopoulos[1,2], Thanasis Vergoulis[2], Ilias Kanellos[2], Theodore Dalamagas[2], and Christos Tryfonopoulos[1]

[1] Univ. of the Peloponnese, Dep. of Informatics & Tel/tions, Tripoli, 22100, Greece
{schatzop,trifon}@uop.gr
[2] IMSI, "Athena" Research Center, Athens, 15125, Greece
{vergoulis,ilias.kanellos,dalamag}@athenarc.gr

Abstract. As the number of published scientific papers continuously increases, the need to assess paper impact becomes more valuable than ever. In this work, we focus on citation-based measures that try to estimate the popularity (current impact) of an article. State-of-the-art methods in this category calculate estimates of popularity based on paper citation data. However, with respect to recent publications, only limited data of this type are available, rendering these measures prone to inaccuracies. In this work, we present `ArtSim`, an approach that exploits paper similarity, calculated using scholarly knowledge graphs, to better estimate paper popularity for recently published papers. Our approach is designed to be applied on top of existing popularity measures, to improve their accuracy. We apply `ArtSim` on top of four well-known popularity measures and demonstrate through experiments its potential in improving their popularity estimates.

Keywords: Scientific Impact Assessment · Scholarly Knowledge Graphs.

1 Introduction

With the growth rate of scientific articles continuously increasing [8], the reliable assessment of their scientific impact is now more valuable than ever. As a result, a variety of *impact measures* have been proposed in the literature, aiming to quantify scientific impact at the article level. Such measures have various practical applications: for instance, they can be used to rank the results of keyword-based searches (e.g., [18]), facilitating literature exploration and reading prioritisation, or to compare and monitor the impact of different articles, research projects, institutions, or researchers.

Since scientific impact can be defined in many, distinct ways [3], the proposed measures vary in terms of the approach they follow (e.g., citation-based, altmetrics), as well as in the aspect of scientific impact they attempt to capture (e.g., impact in academia, social media attention). In this work, we focus on citation-based measures, that attempt to estimate the current impact of each article, i.e., its current *popularity* or hype. Providing accurate estimations of article popularity is an open problem, as has been shown by a recent extensive experimental evaluation [6]. Furthermore, popularity distinctly differs from the overall (long-term) impact of an article that is usually captured by traditional citation-based measures (e.g., citation count).

One important issue in estimating article popularity is to provide accurate estimations for the most recently published articles. The estimations of most popularity measures rely on the existing citation history of each article. However, since very limited citation history data are available for recent articles, their impact estimation based on these data is prone to inaccuracies. Hence, these measures fail to provide correct estimations for recent articles. To alleviate this issue, in this work we introduce `ArtSim`, a new approach to assess article popularity. Our approach does not only rely on each article's historical data, but also considers the history of older, similar papers, for which these data are more complete. The intuition behind our method is that similar papers are likely to follow a similar trajectory in terms of popularity. To quantify article similarity, we exploit the corresponding author lists and the involved topics[3]. This information is available in *scholarly knowledge graphs*, a large variety of which has been made available in recent years (e.g., AMiner's DBPL-based datasets [17], the Open Research Knowledge Graph [5], the OpenAIRE Research Graph [9, 10].)

The real power of `ArtSim` comes from the fact that it can be applied on top of any existing popularity measure to improve its accuracy. To demonstrate this, we first apply `ArtSim` on top of the best performing popularity measures (according to [6]) to produce a set of improved measures. Then, we examine the achieved benefits (by replicating the experimental process in [6]). Our experiments indicate that `ArtSim` effectively enhances the performance of common measures in estimating article popularity.

2 Our approach

2.1 Background

Our proposed method aims at transforming popularity scores based on any popularity measure, in order to increase the accuracy of its estimations. To achieve this, it attempts to improve the estimation for all recent articles by exploiting path-based article similarities in scholarly *knowledge graphs*.

Knowledge graphs, also known as *heterogeneous information networks* [15], are graphs that encode rich domain-specific information about various types of entities, represented as nodes, and the relationships between them, represented as edges. Figure 1 presents an example of such a knowledge graph, consisting of nodes representing papers, authors and topics (i.e., three different node types). Two types of (bidirectional) edges are present in this example network: edges between authors and papers, denoted as `Author - Paper` (`AP` or `PA`, for brevity), and edges between papers and topics, denoted as `Paper - Topic` (`PT`

[3] Here we use similarity based on authors and topics as a proof of concept. However, our approach can be generalized using any other definition of article similarity.

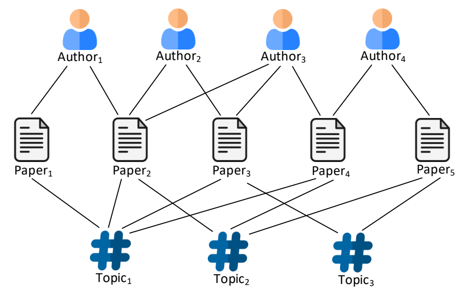

Fig. 1: A scholarly knowledge graph including papers, authors and topics.

or `TP`). The former represent the authorship of papers, while the latter encode the information that a particular paper is written on a particular topic.

Various semantics are implicitly encoded in the paths of knowledge graphs. For example, in the graph of Figure 1, a path from an author node to another one of the same type that involves an `AP` edge followed by a `PT`, a `TP`, and a `PA` edge relates two authors that have published works in the same topic (e.g., both $Author_1$ and $Author_4$ have papers about $Topic_1$). In fact, all paths that correspond to the same sequence of node and edge types encode latent, "multi-hop" relationships having the same interpretation. In the literature, such a sequence of node and edge types (e.g., the `APTPA` of the previous example) is known as a *metapath*. Metapaths are useful for many graph analysis and exploration applications. For example, in our approach, we use them to calculate *metapath-based similarities*: the similarity between two nodes of the same type, based on the semantics of a given metapath, can be captured by counting the number of instances of this metapath connecting these nodes (e.g., [16, 20]).

2.2 ArtSim

Our proposed method, called `ArtSim`, can be applied on top of any popularity measure to increase the accuracy of its estimations. As such, `ArtSim` takes the scores calculated by any popularity measure as input, applies transformations on them, and produces a new set of improved popularity scores. This process is presented in Figure 2.

The transformations applied on popularity scores by `ArtSim` rely on the intuition that similar articles are expected to share similar popularity dynamics. To calculate the similarity between different papers, `ArtSim` relies on the JoinSim [20] similarity measure calculated on `PAP` and `PTP` metapaths. Evidently, the similarity between papers is not uniquely defined, hence different metapaths encode different similarity semantics. For example, while `PAP` metapaths define paper similarity based on their common authors, `PTP` metapaths define paper similarity based on their common topics. `ArtSim` uses the calculated similarity scores to provide improved popularity estimates (scores), focusing in particular

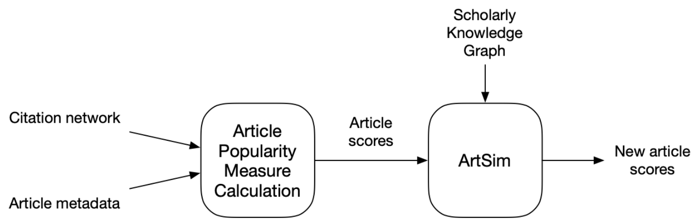

Fig. 2: Our proposed approach.

on recent papers that have a limited citation history. The calculation of `ArtSim` scores is based on the following formula:

$$S(p) = \begin{cases} \alpha * S_{PAP}(p) + \beta * S_{PTP}(p) + \gamma * S_{initial}(p), & p.year \geq t_c - y \\ S_{initial}(p), & \text{otherwise} \end{cases}$$

where S_{PAP} and S_{PTP} are the average popularity scores of all the articles that are similar to p, based on metapaths `PAP` and `PTP` respectively. $S_{initial}$ is the popularity score of paper p, based on the original popularity measure and t_c denotes the current time. Finally, our method applies transformations on popularity scores for those papers published in years, which range in the time span $[t_c - y, t_c]$, where $y \geq 0$.

Note, that parameters $\alpha, \beta, \gamma \in [0, 1]$. Furthermore, we set α, β, γ so that $\alpha + \beta + \gamma = 1$. Varying these parameters in the range $[0 - 1]$ has the following effects: as α increases, article similarity is mostly calculated based on common authors. As β increases, article similarity depends mainly on common topics. Finally, as γ approaches 1 the popularity scores remain identical to those calculated by the initial popularity measure.

3 Evaluation

In this section, we discuss the experiments conducted to assess the effectiveness of our method. Section 3.1 discusses the experimental setup of our evaluation approach i.e., the datasets, methodology and popularity measures used, and Section 3.2 showcases the improvements that `ArtSim` brings to popularity measures.

3.1 Setup

Datasets. For our experiments, we used the following datasets:

- *DBLP Scholarly Knowledge Graph (DSKG) dataset.* It contains data for $3,079,008$ papers, $1,766,548$ authors, $5,079$ venues and $3,294$ topics from DBLP. It is based on AMiner's citation network dataset [17] enriched with topics from the CSO Ontology [13] using the CSO Classifier [12] on paper abstracts.

- *DBLP Article Similarities (DBLP-ArtSim) dataset.* It contains similarities among articles in the previous network based on different metapaths. In particular, we calculated paper similarities based on their author list using metapath `Paper - Author - Paper` and on common topics, captured by metapath `Paper - Topic - Paper`. This dataset is openly available on Zenodo[4] under CC BY 4.0 license.

Methodology. To assess paper popularity we follow the experimental framework proposed in [6], which is based on the evaluation of total orderings (rankings) of papers based on their short-term future citations. As explained in the referenced work, the number of citations a paper receives in the near future, is a good a-posteriori indicator of its current popularity. Thus, the aforementioned rankings can be used as a ground truth for experiments to evaluate the effectiveness of measures in ordering papers based on their popularity. This approach is also suitable for our needs, since an overall ordering of papers can be used as a basis for the comparison of any pair of papers based on their relative impact.

Based on the above, we define a split time t_s that splits our dataset in half, into two equally sized sets. The first half, denoted as $S(t_s)$ contains papers published before t_s and is considered known to the examined popularity measures. We also consider a future state of the network in the time $t_s + \tau$ which we use to construct the ground truth. In our case, set $S(t_s + \tau)$ contains 30% more articles than $S(t_s)$. We finally rank each paper in the future state of the network based on the number of its citations (i.e., the citations it received in the time span from t_s to $t_s + \tau$). This ranked list acts as the ground truth and is used to evaluate the effectiveness of popularity measures.

We measure the effectiveness of any approach compared to the ground truth using the following two measures:

- *Kendall's* τ [7], is a non-parametric measure for the similarity in the ordering of two ranked lists, based on the number of concordantly ordered pairs of items between them. Its values range from -1 to 1, with 1 denoting perfect agreement, -1 denoting perfect inversion, and 0 denoting no correlation.
- *Normalised Discounted Cumulative Gain at rank k (nDCG@k)* is a measure of ranking quality using the graded relevance scale of documents in the ranking list. It is a normalized version of the *Discounted Cumulative Gain (DCG) at rank* k in the range $[0, 1]$. The value 1 corresponds to the ideal *DCG*, achieved when the ranking perfectly agrees with the ground truth.

We use Kendall's τ and nDCG@k to capture the overall, and top-k similarity of the ranked lists to the ground truth, respectively.

Popularity measures. To evaluate our method, we chose the four overall best performing popularity measures in terms of correlation in the DBLP dataset for the scenario of popularity according to a recent experimental study [6]. The

[4] `https://doi.org/10.5281/zenodo.3778916`

Table 1: Parameter configuration for each popularity measure.

Method	Configuration
ECM	$\alpha = 0.2, \gamma = 0.4$
RAM	$\gamma = 0.4$
CR	$\alpha = 0.4, \tau_{dir} = 10$
FR	$\alpha = 0.5, \beta = 0.2, \gamma = 0.3, \rho = -0.42$

optimal parameter settings per measure were selected after running each one for various parameterisations and calculating the correlation of the ranked list produced by each one to the ground truth ranking. Table 1 presents the selected parameter setting per popularity measure. We briefly describe the intuition behind each measure below:

- *Retained Adjacency Matrix (RAM)* [4] estimates popularity using a time-aware adjacency matrix to capture the recency of cited papers. The parameter $\gamma \in (0, 1)$ is used as a basis of an exponential function to scale down the value of a citation link according to its age.
- *Effective Contagion Matrix (ECM)* [4] is an extension of RAM that also considers the temporal order of citation chains apart from direct links. It uses two parameters $\alpha, \gamma \in (0, 1)$ where α is used to adjust the weight of citation chains based on their length and γ is the same as in RAM.
- *CiteRank (CR)* [19] measures popularity by simulating the behaviour of researchers searching for new articles. It uses two parameters $\alpha \in (0, 1)$ and $\tau_{dir} \in (0, \infty)$ to model the traffic to a given paper. A paper is randomly selected with an exponentially discounted probability according to its age with τ_{dir} being the decay factor. Parameter α is the probability that a researcher stops its search, with $1 - \alpha$ being the probability that he continues with a reference of the paper he just read.
- *FutureRank (FR)* [14] scores are calculated combining PageRank ~~scores~~ with calculations on a bipartite graph with authors and papers, while also promoting recently published articles with time-based weights. It uses parameters $\alpha, \beta, \gamma \in (0, 1)$ and $\rho \in (-\infty, 0)$; α is the coefficient of the PageRank scores, β is the coefficient of the authorship scores and γ is the coefficient of time-based weights which exponentially decrease based on the exponent ρ.

3.2 Effectiveness of our approach

Improvements in correlation. In this experiment, we examine the gains of `ArtSim` in terms of Kendall's τ correlation. For each examined popularity measure (ECM, RAM, CR and FR) we vary parameters α, β, γ of our method, as well as parameter y, which sets the number of past years for which we consider papers in the cold start phase. We visualise, in the form of heatmaps, the correlation achieved for each method for different configurations when $\alpha, \beta \in [0, 0.6]$ and $y \in [1, 4]$ (Figure 3). Parameter γ is implied, since $\alpha + \beta + \gamma = 1$.

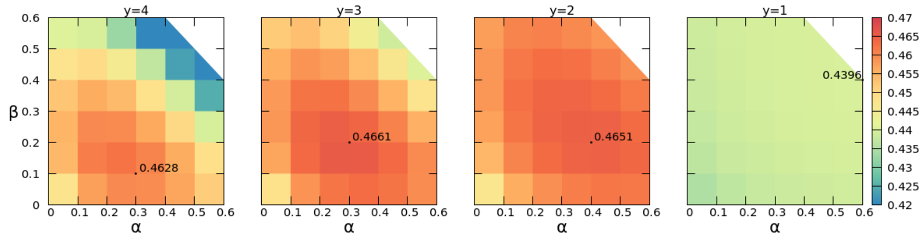

(a) ECM

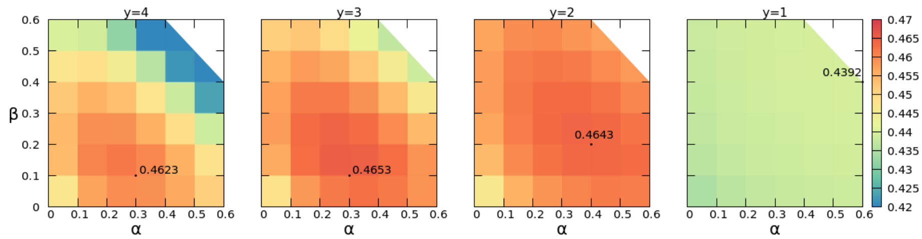

(b) RAM

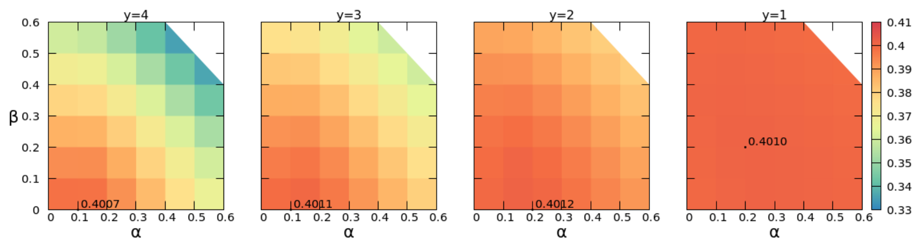

(c) CR

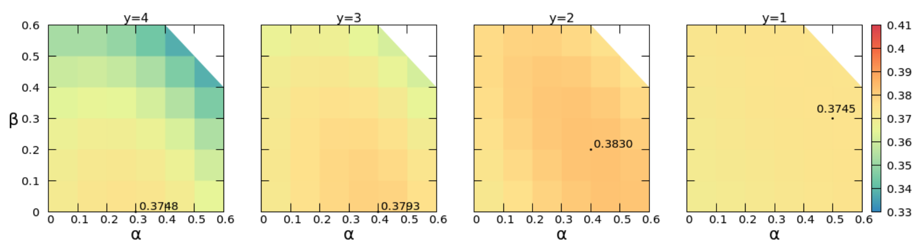

(d) FR

Fig. 3: Heatmaps depicting the effectiveness of our approach for different parameters in terms of Kendall's τ correlation for each popularity measure.

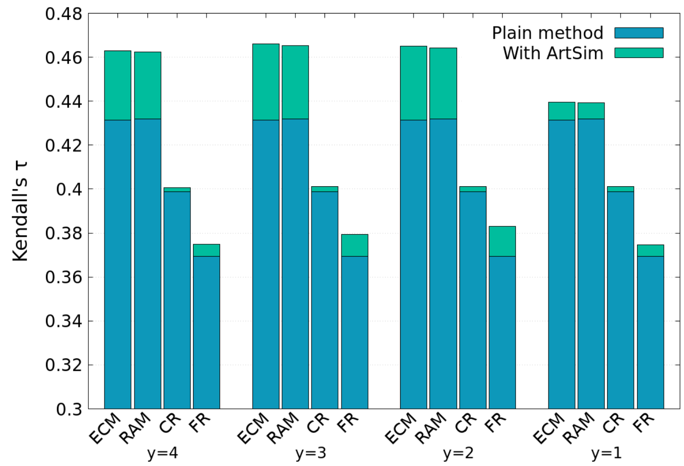

Fig. 4: Effectiveness of our approach in terms of Kendall's τ for the best parameterisation per year for each popularity measure.

In general, we observe that our approach achieves the maximum gains for $y \in \{2, 3\}$. For $y = 1$, we see that for all methods, the scores have a small deviation as expected, as our method adjusts the popularity scores of a small fraction of the overall articles, i.e., only those published in the last year. The heatmaps also validate that both scores based on similarity of authors and topics are important, since the correlation observed decreases when both parameters α and β approach zero.

Based on the experiments, `ArtSim` achieves the best correlation, $\tau = 0.4661$ using the ECM method when $\{\alpha = 0.3, \beta = 0.2, \gamma = 0.6, y = 3\}$. The best scores for `ArtSim` using the other methods are $\tau = 0.4653$ at $\{\alpha = 0.3, \beta = 0.1, \gamma = 0.7, y = 3\}$ for RAM, $\tau = 0.4012$ at $\{\alpha = 0.2, \beta = 0, \gamma = 0.8, y = 2\}$ for CR, and $\tau = 0.3830$ at $\{\alpha = 0.4, \beta = 0.2, \gamma = 0.4, y = 2\}$ for FR.

We further examined the gains of `ArtSim` in terms of Kendall's τ correlation compared to the plain popularity measures. The best parameter configuration for each method is selected for each year. The results are illustrated in Figure 4. Overall, significant improvements in correlation are observed when `ArtSim` is applied on the ECM and RAM measures. In particular, ECM and RAM are improved by 8% for the best parameter configuration for $y \in [2, 4]$. As expected, smaller gains for all methods are achieved for $y = 1$. In that case, as previously mentioned, our approach affects the popularity score of the papers published only in the last year, affecting only a small fraction of the overall papers.

Improvements in nDCG. In this experiment, we examine the effectiveness of `ArtSim` in terms of $nDCG@k$ for all considered popularity measures compared to the ground truth. We performed two sets of experiments: (a) we measure

Table 2: Effectiveness of our approach for $y = 3$ in terms of $nDCG@k$.

	Small values of k			**Large values of k**		
	5	**50**	**500**	**400,000**	**500,000**	**600,000**
ECM	0.8323	0.8634	0.8953	0.8780	0.8833	0.8884
ArtSim-ECM	0.8323	0.8634	0.8953	**0.8837**	**0.8912**	**0.9003**
RAM	0.8588	0.8521	0.8943	0.8774	0.8842	0.8881
ArtSim-RAM	0.8588	0.8521	0.8943	**0.8836**	**0.8904**	**0.9008**
CR	0.3530	0.5263	0.6060	0.7904	0.8149	0.8272
ArtSim-CR	0.3530	0.5263	0.6060	**0.7983**	**0.8199**	**0.8307**
FR	0.3403	0.5018	0.5526	0.7586	0.7934	0.8101
ArtSim-FR	0.3403	0.5018	0.5526	**0.7731**	**0.7961**	**0.8152**

the nDCG@k achieved by `ArtSim`, varying k, and (b) we examine how `ArtSim` affects top-k results in two indicative keyword search scenarios.

Table 2 presents the $nDCG@k$ values, per popularity measure, both when plainly run, as well as when `ArtSim` is applied on them. In this experiment we select $y = 3$, which produces the best correlation according to the previously presented results. In particular, we separately examined `ArtSim`'s behaviour for small and for large values of parameter k. In particular we examine nDCG@k for $k \in \{5, 50, 500\}$, as well as for $k \in \{400.000, 500.000, 600.000\}$. Interestingly, for small values of k, our approach performs equally to the initial popularity measures, at its best configuration. This behaviour indicates that existing state-of-the-art popularity measures accurately identify the top papers in terms of popularity. Another apparent explanation is that the set of most popular papers, at the global level, mainly includes those that already have a more extended citation history, i.e., they have become known by the scientific community and maintain their status. `ArtSim`'s performance gain becomes apparent for larger values of k. In relative terms, our method improves upon the nDCG achieved by the popularity measures, starting at the top 7% of the most popular papers and beyond. In other words, our method does provide gains in terms of nDCG for the large majority of papers, while maintaining the nDCG values achieved for the overall most popular papers. Likely, these larger sets of top popular papers also include recently published ones for which the popularity estimations are improved by `ArtSim`. This is further supported by the observation that the nDCG values achieved increase with k, i.e., the more recent papers are included, the more noticeable `ArtSim`'s effect.

In our second set of experiments we illustrate that the performance gains, which are observed at the global level only for large values of k, are not negligible in practical applications. For example, in a real scenario of literature exploration, academic search engine users usually refine their searches using multiple keywords and by applying filters (e.g., based on the venues of interest or the publication years). Their intention is to reduce the number of papers they have to examine, however even in this case usually at least hundreds of papers are contained in the results. Hence, effective ranking is crucial to facilitate the reading

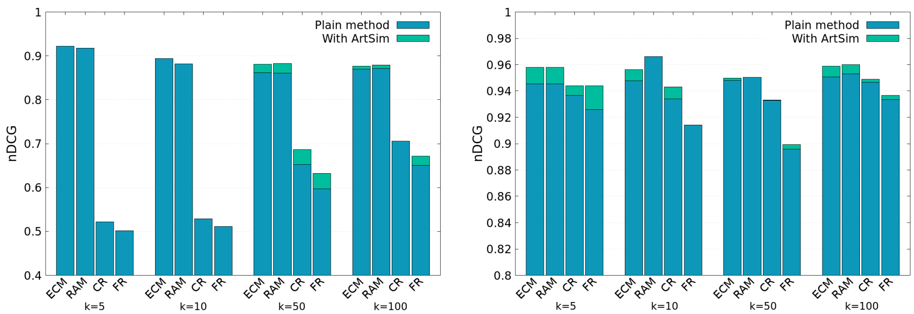

(a) Expert Finding (b) Recommender Systems

Fig. 5: Effectiveness of our approach in terms of nDCG for different keyword search scenarios with $y = 3$ and varying k.

prioritization. Furthermore, the resulting article lists usually contain only a small subset of all the articles in a dataset. We do not expect to only find the overall most popular papers in such subsets, owing to the different publication dynamics in each subdomain (i.e., different research communities have different sized and publish and/or cite at different speeds). Hence, in our second set of experiments we measure the nDCG@k achieved by `ArtSim` for $k \in \{5, 10, 50, 100, 500\}$ on the results of two indicative keyword search scenarios.

In the first search scenario we used the query "expert finding". This keyword search resulted in a set of 549 articles. Figure 5a presents the nDCG values for this search, per popularity measure, along with the gains of `ArtSim` for $y = 3$. We observe that `ArtSim` improves the nDCG values for $k = 50$ and $k = 100$. In our second scenario, we tried a conditioned query. Particularly, we used "recommender systems" as the search keywords keeping only papers published in well-known venues of data management and recommender systems, namely VLDB, SIGMOD, TKDE, ICDE, EDBT, RecSym and ICDM. The result set includes 525 articles. Figure 5b presents the nDCG results. We observe that `ArtSim` boosts nDCG scores for all measures, starting from the smallest value of $k = 5$. These results indicate that in addition to improving the overall correlation, our approach also offers improvements in the case of practical, keyword-search based queries with regards to the top returned results.

4 Related Work

There is a lot of work in the areas of bibliometrics and scientometrics to quantify the impact of scientific articles. In particular, much focus has been put on quantifying current or recent impact of scientific publications [4, 14, 19], in contrast to the overall impact traditionally estimated by bibliometric measures, such as the citation counts. In depth examinations of various impact measures that have been proposed in the literature can be found in [1, 6]. In contrast to the above, our own approach does not aim to introduce a new popularity measure, but rather aims at improving the accuracy of existing ones. To the best of our knowledge, this is the first approach of this type to be introduced.

Our approach is built upon recent work on entity similarity in the area of heterogeneous information networks. Some of the first entity similarity approaches for such networks (e.g., PopRank [11] and ObjectRank [2]) are based on random walks. Later works, like PathSim [16], focus on providing more meaningful results by calculating node similarity measures based on user-defined semantics. Our own work is based on JoinSim [20], which is more efficient compared to PathSim, making it more suitable for analyses on large scale networks.

5 Conclusions

We presented `ArtSim`, an approach that can be applied on top of existing popularity measures to increase the accuracy of their results. The main idea of our approach is that the popularity of papers in their cold start period can be better estimated based on the characteristics of other, similar papers. We calculate the similarity of papers using metapath analyses on the underlying scholarly knowledge graphs. Our experimental evaluation showcases the effectiveness of `ArtSim`, yielding noteworthy improvements in terms of Kendall's *tau* correlation and nDCG when applied on four state-of-the-art popularity measures.

Acknowledgments

We acknowledge support of this work by the project "Moving from Big Data Management to Data Science" (MIS 5002437/3) which is implemented under the Action "Re-inforcement of the Research and Innovation Infrastructure", funded by the Operational Programme "Competitiveness, Entrepreneurship and Innovation" (NSRF 2014-2020) and co-financed by Greece and the European Union (European Regional Development Fund). Icons in Figure 1 were collected from www.flaticon.com and were made by Freepik, Good Ware and Pixel perfect.

References

1. Bai, X., Liu, H., Zhang, F., Ning, Z., Kong, X., Lee, I., Xia, F.: An overview on evaluating and predicting scholarly article impact. Information (2017)
2. Balmin, A., Hristidis, V., Papakonstantinou, Y.: Objectrank: Authority-based keyword search in databases. In: VLDB (2004)
3. Bollen, J., Van de Sompel, H., Hagberg, A., Chute, R.: A principal component analysis of 39 scientific impact measures. PloS one (2009)
4. Ghosh, R., Kuo, T., Hsu, C., Lin, S., Lerman, K.: Time-aware ranking in dynamic citation networks. In: International Conference on Data Mining Workshops. pp. 373–380 (2011)

5. Jaradeh, M.Y., Oelen, A., Farfar, K.E., Prinz, M., D'Souza, J., Kismihók, G., Stocker, M., Auer, S.: Open research knowledge graph: Next generation infrastructure for semantic scholarly knowledge. In: International Conference on Knowledge Capture (2019)
6. Kanellos, I., Vergoulis, T., Sacharidis, D., Dalamagas, T., Vassiliou, Y.: Impact-based ranking of scientific publications: A survey and experimental evaluation. IEEE Transactions on Knowledge and Data Engineering (2019)
7. Kendall, M.G.: Rank correlation methods. (1948)
8. Larsen, P.O., von Ins, M.: The rate of growth in scientific publication and the decline in coverage provided by science citation index. Scientometrics **84**(3), 575–603 (2010)
9. Manghi, P., Atzori, C., Bardi, A., Shirrwagen, J., Dimitropoulos, H., La Bruzzo, S., Foufoulas, I., Löhden, A., Bäcker, A., Mannocci, A., Horst, M., Baglioni, M., Czerniak, A., Kiatropoulou, K., Kokogiannaki, A., De Bonis, M., Artini, M., Ottonello, E., Lempesis, A., Nielsen, L.H., Ioannidis, A., Bigarella, C., Summan, F.: Openaire research graph dump (2019), `https://doi.org/10.5281/zenodo.3516918`
10. Manghi, P., Bardi, A., Atzori, C., Baglioni, M., Manola, N., Schirrwagen, J., Principe, P.: The openaire research graph data model (Apr 2019), `https://doi.org/10.5281/zenodo.2643199`
11. Nie, Z., Zhang, Y., Wen, J.R., Ma, W.Y.: Object-level ranking: bringing order to web objects. In: WWW (2005)
12. Salatino, A., Osborne, F., Thanapalasingam, T., Motta, E.: The CSO Classifier: Ontology-Driven Detection of Research Topics in Scholarly Articles, pp. 296–311 (2019)
13. Salatino, A.A., Thanapalasingam, T., Mannocci, A., Osborne, F., Motta, E.: The computer science ontology: a large-scale taxonomy of research areas. In: International Semantic Web Conference (2018)
14. Sayyadi, H., Getoor, L.: Futurerank: Ranking scientific articles by predicting their future pagerank. In: Proceedings of the 2009 SIAM International Conference on Data Mining. SIAM (2009)
15. Shi, C., Li, Y., Zhang, J., Sun, Y., Yu, P.S.: A survey of heterogeneous information network analysis. IEEE Trans. Knowl. Data Eng. (2017). https://doi.org/10.1109/TKDE.2016.2598561
16. Sun, Y., Han, J., Yan, X., Yu, P.S., Wu, T.: Pathsim: Meta path-based top-k similarity search in heterogeneous information networks. VLDB (2011)
17. Tang, J., Zhang, J., Yao, L., Li, J., Zhang, L., Su, Z.: ArnetMiner: Extraction and Mining of Academic Social Networks. In: ACM SIGKDD. ACM (2008)
18. Vergoulis, T., Chatzopoulos, S., Kanellos, I., Deligiannis, P., Tryfonopoulos, C., Dalamagas, T.: Bip! finder: Facilitating scientific literature search by exploiting impact-based ranking. In: CIKM (2019)
19. Walker, D., Xie, H., Yan, K., Maslov, S.: Ranking scientific publications using a model of network traffic. JSTAT (2007)
20. Xiong, Y., Zhu, Y., Yu, P.S.: Top-k similarity join in heterogeneous information networks. IEEE Transactions on Knowledge and Data Engineering (2015)