Evolution of Scientific Interests: A Time-Aware Recommender System

Bruno S. B. Contreras[1], **Luciano A. Digiampietri**[1]

[1]School of Arts, Sciences and Humanities – University of São Paulo (EACH/USP)
São Paulo – SP – Brazil

{bsbcontreras,digiampietri}@usp.br

Abstract. *The increasing volume of scientific publications makes it challenging for researchers to stay updated. Traditional recommendation systems often overlook the evolving nature of academic interests, leading to outdated suggestions. This article introduces a time-aware recommendation system (TARS) that incorporates temporal dynamics to improve the relevance of scientific article recommendations. In experiments with OpenAlex data, TARS outperformed static models, achieving up to 12% better precision and recall, particularly for active researchers. The study highlights the importance of temporal modeling in capturing evolving preferences and enhancing engagement in academic information systems.*

1. Introduction

In recent years, the volume of scientific publications has greatly increased, posing a significant challenge for researchers who seek to stay updated on advancements in their fields. In this context, scientific article recommendation systems have emerged as essential tools, helping to filter relevant information and connect researchers to content aligned with their academic interests. However, most traditional recommendation systems assume that researchers' interests remain static, overlooking that these interests may evolve due to the acquisition of new skills, collaboration across interdisciplinary fields, or adaptation to new trends in the scientific domain [Perera and Zimmermann 2020].

This work aims to explore how incorporating a temporal dimension into scientific article recommendation systems can lead to more accurate and tailored recommendations that meet the dynamic needs of researchers. This paper aims to assess how recommendation models with a temporal contextualization stage compare to static models regarding recommendation effectiveness and relevance. To achieve this, we will conduct a specific analysis of the evolution of interests among Brazilian researchers in Computer Science, examining how these interests have transformed over time [Campos et al. 2013].

Ignoring the evolving nature of researchers' interests can cause recommendation systems to suggest outdated or misaligned content with users' current preferences. Recommendation models that incorporate a temporal perspective have the potential to provide more accurate and adaptive recommendations, adjusting to each researcher's academic trajectory.

In this paper, we evaluate a scientific recommendation approach that considers the temporal evolution of researchers' interests. Using a database of scientific publications that documents academic output over the years, we explore methods to capture and model these changes and investigate the impact of temporal inclusion on recommendation quality. Furthermore, we compare the performance of a traditional (static) model with a time-sensitive model, analyzing the effectiveness of each approach in terms of precision and relevance.

2. Basic Concepts

2.1. Time-Aware Approach in Recommender Systems (TARS)

A temporal approach in recommendation systems refers to the integration of time as a central element in modeling researchers' interests and preferences. Instead of treating interests as fixed and unchanging, this approach recognizes that they may evolve due to various factors, such as the acquisition of new knowledge, shifts in research focus, or the emergence of new trends in scientific fields [Perera and Zimmermann 2020].

TARS models can use various techniques, such as decay functions, time-aware collaborative filtering, and time-series analysis, to capture changes in user interests over time [Chen et al. 2020, Gómez-Losada and Duch-Brown 2019]. These techniques allow recommendation systems to adjust researchers' preferences based on recent publications, co-authorships, and citations, resulting in a more personalized and dynamic recommendation experience.

2.2. Cold Start Initialization

The cold start problem in scientific article recommendation algorithms refers to the difficulty of recommending articles to researchers with little or no recent activity data, due to the lack of historical information, such as publication records, co-authorships, or citations for that researcher. This issue often arises when a new researcher joins the platform without sufficient publication history [Bai et al. 2019].

The cold start for a new researcher presents significant challenges and is typically addressed in two ways. When no usage history is available and no researcher profile information is provided (such as field, role, academic background, etc.), systems usually suggest articles that have been popular over recent days, weeks, or months. These articles are selected based on popularity, citation count, or relevance in specific contexts, offering global or personalized suggestions.

Alternatively, if there is associated profile information, systems tend to use a collaborative filtering approach. In this case, articles read by researchers with similar profiles are suggested, using profile similarity through resources like *Knowledge Graphs* [Kuznetsov and Kordík 2023].

2.3. Doc2Vec and Node2Vec in Recommender Systems

Doc2Vec is a natural language processing (NLP) technique that extends the *Word2Vec* [Mikolov et al. 2013] model to represent entire documents as continuous vectors in a vector space. By capturing the semantic relationships between words and their contexts, *Doc2Vec* enables the representation of long and complex documents, such as scientific articles, in a way that preserves their content and meaning. This makes it particularly suitable for recommendation systems where understanding document content is critical. For example, *Doc2Vec* can be applied to model research papers based on their abstracts, titles, and keywords, enabling the identification of relevant articles for a given researcher based on semantic similarity [Le and Mikolov 2014].

Node2Vec, on the other hand, is an algorithm for learning continuous representations (embeddings) of nodes in a graph. Inspired by the concept of word embeddings, Node2Vec generates high-dimensional vector representations of graph nodes, preserving the graph's structural and relational information. This is achieved through biased random walks, which produce sequences of nodes that are subsequently used to train a Skip-gram model, similar to the training process in Word2Vec. In the context of recommendation systems, Node2Vec is particularly effective for modeling relationships in knowledge graphs, such as those constructed from citation networks, or co-authorship networks [Grover and Leskovec 2016].

Combining *Doc2Vec* and *Node2Vec* allows the integration of semantic document content and structural graph relationships into a unified recommendation framework. For instance, while *Doc2Vec* embeddings capture the thematic alignment of articles, *Node2Vec* embeddings provide insights into their relational context, such as citation or co-authorship patterns. Together, these methods enable the development of hybrid recommendation systems that leverage both content-based and graph-based approaches for improved recommendation quality.

2.4. BERT Embeddings

BERT is a large language model for NLP, based on the *Transformer* architecture [Devlin et al. 2019]. The defining feature of BERT is its ability to understand the context of words bidirectionally (analyzing both left-to-right and right-to-left), allowing for a more precise understanding of the meaning of words concerning others within a sentence. This bidirectional processing makes BERT especially effective in tasks requiring linguistic comprehension, such as entity recognition, sentiment analysis, and, particularly, question answering. By applying embeddings generated by BERT, the recommendation system can compare these article vectors with embeddings that represent each researcher's interest profile, producing more accurate reading suggestions [Devlin et al. 2019].

2.5. Recommender System Evaluation

In addition to traditional classification metrics like precision, recall, and F1-score, recommendation systems are also evaluated using metrics such as NDCG and MRR, which specifically measure the quality and relevance of the generated recommendations.

2.5.1. *Normalized Discounted Cumulative Gain*

NDCG@K (Normalized Discounted Cumulative Gain) is used to evaluate the quality of ranked lists of items. There are a few steps necessary to understand how $NDCG$ works. The first step is calculating the *Cumulative Gain* (CG), this metric sums the relevance scores of items in an ordered list, where relevance is an indicator function equal to 1 if the item in position k is relevant and 0 otherwise. To improve the result, the concept of *Discounted Cumulative Gain* (DCG) is introduced, which incorporates a discount factor to assign greater importance to items at the top of the list. Finally, the $NDCG$ metric is a normalized form of DCG. Here, $IDCG$ represents the ideal DCG, indicating the maximum DCG achievable for the set of items. It is calculated using the same formula as DCG, but with items arranged in descending order of their true relevance scores.

$$\text{DCG@k} = \sum_{i=1}^{k} \frac{\text{rel}_i}{\log_2(i+1)}; \;\; NDCG@k = \frac{DCG@k}{IDCG@k} \tag{1}$$

2.5.2. Mean Reciprocal Rank

MRR@K (*Mean Reciprocal Rank*) is a metric used to evaluate the relevance and quality of rankings in recommended items. It considers the position of relevant items within the recommendation list, giving more weight to items appearing at the top. For each user, the *Reciprocal Rank* (RR) is calculated as the inverse of the position of the first relevant item in the recommendation list. For example, if the first relevant item is in the 2nd position on the list, then RR will be $1/2$. The $MRR@K$ is the average of the RRs for a group of users, considering only the top K items in each user's recommendation list. Thus, it assesses the relevance of items in the most important positions on the list. The calculation of this metric is presented in Equation 2.

$$MRR@k = \frac{1}{K} \sum_{i=1}^{k} \frac{1}{\text{rank}_i} \tag{2}$$

3. Literature Review

This section explores relevant studies on various approaches to scientific papers' recommendation, including the use of knowledge graphs, contextual embeddings, and time-aware modeling. It focuses on how each technique can capture and represent the evolution of academic interests over time.

Heterogeneous graphs have emerged as a promising approach for capturing the complexity of academic interactions, surpassing homogeneous graphs by encoding different relationships, such as authorship, and co-authorship. Recent studies show that these graphs can represent real information networks with greater semantic richness, enabling the modeling of complex structures and relationships in a more representative way [Ly et al. 2024]. Heterogeneous graphs have proven to be more effective in representing the multifaceted relationships between scientific articles, authors, and topics, especially when used to recommend publications in co-authorship networks and related themes.

The work from Ikoma and Matsubara [Ikoma and Matsubara 2023] introduces an article recommendation method that leverages citation contexts in academic documents. Citation contexts are portions of text that describe the purpose and relationship between the cited article and the one making the citation. This method consists of two main stages: the first step involves a candidate selection model that maps the input article and existing articles into a vector space, while the second stage is a decision model that estimates the likelihood of each candidate being a relevant citation. This method suggests that using context in recommendations allows for more accurate and content-aligned suggestions, especially in highly specialized domains.

Furthermore, the time-aware approach in recommendation systems has been extensively explored to improve recommendation accuracy over time. Campus et al. [Campos et al. 2013] provide a comprehensive review of time-aware recommendation systems, analyzing existing evaluation protocols and offering methodological guidelines for robust evaluation. They emphasize the importance of considering temporality to capture changes in user interests over time, utilizing techniques such as decay functions, time-aware collaborative filtering, and time-series analysis.

Jiang et al. [Jiang et al. 2023] developed a model called TAPRec (Time-Aware Paper Recommendation via the Modeling of Researchers' Dynamic Preferences). It was proposed to address scientific article recommendations by considering researchers' dynamic preferences. This model uses *Self-Attention* methods to aggregate researchers' long-term research interests and Temporal Convolutional Networks (TCN) to capture short-term research focuses. Additionally, for researchers with few publications, TAPRec combines the dynamic preferences of co-authors to address the cold-start problem. TAPRec highlights the importance of incorporating temporality in scientific recommendation modeling, allowing the system to evolve alongside researchers' interests and provide more accurate and relevant recommendations.

These techniques enable recommendation systems to continuously adjust user preferences based on recent and historical behavior, resulting in a more personalized and dynamic recommendation experience. Furthermore, temporality can help identify seasonal patterns and specific events that influence users' interests, allowing for more contextually relevant recommendations.

4. Materials and Methods

This section presents the procedures adopted for building and evaluating the recommendation system. We describe the data collection, model architecture, and metrics used to measure the results.

4.1. Data Collection

Data collection was conducted directly from the OpenAlex source[1], a comprehensive database containing over 250 million indexed academic works across a vast array of subjects and fields. The purpose of this data collection is to establish a graph-based dataset. A heterogeneous network was created connecting authors and works, where each work is linked to others through reference relationships. On the other hand, the relationship between authors and works is defined through authorship connections. Additionally, there exists an implicit co-authorship relationship between authors, which occurs when two different authors contribute to the same work.

It is important to highlight that OpenAlex is a dynamic system continuously updated with new academic contributions. It implies that, depending on the area of interest and the date of access, it is possible to specify which topics of the articles will be analyzed. This dynamic characteristic provides a contemporary perspective and the necessary flexibility to tailor the research to the specific demands of the project, offering more contextualized results.

[1] `https://openalex.org/`, accessed on 03/03/2025

4.1.1. Data Collection Strategy

Recognizing the importance of a strategic approach, this project focuses on collecting data from a subset of the 250 million articles rather than using the entire database. This decision aims to optimize computational efficiency and enables a more focused and specialized analysis tailored to the project's scope.

When narrowing the OpenAlex database to create a scientific article dataset, it is crucial to maintain all publication relationships for each researcher. This approach aims to enhance the development and training of the recommendation system. By preserving all published works, we ensure that the resulting dataset represents the actual academic relationships, thus contributing to the model's fidelity to real-world scenarios. When applying a temporal filter, for instance, by focusing on the most recently published works, certain relationships formerly present in the dataset are excluded. This approach results in the loss of significant connections between authors and works, thus reducing the dataset's alignment with the complexity of interactions from real-world academic networks.

To address this, we used a data collection technique known as *snowball sampling*. In this method, an initial reference work [Mena-Chalco et al. 2014][2] was chosen as the starting point. First, all authors of this article were identified and included in the author network. Then, all other works published by these authors were collected, progressively expanding the co-authorship network. This expansion continues with each new set of authors and their respective works, forming a comprehensive network that captures the connections between authors through their publications.

However, it is important to acknowledge that the snowball sampling strategy may introduce biases into the collected data. To mitigate this limitation in future work, we propose expanding the dataset by starting points that could be selected at random or based on different seed articles from various regions to ensure a more balanced and representative author network.

4.1.2. Collection Results

A total of 102,716 papers were collected, with 1.031 different authors.

Figure 1 illustrates the number of papers published over the years, grouped into five-year intervals, with publications before 1928 consolidated for better visualization. The number of papers is transformed using the natural logarithm to adjust the scale. Table 1 presents the distribution of primary and secondary publication languages across the countries with the highest number of publications in this dataset.

To construct the Table 1, each publication is assigned a single language. However, since a publication can have multiple authors from different countries, the same work may be attributed to multiple nationalities, thereby increasing the total number of publications counted across different countries.

The table highlights that English is the predominant language for publications in most countries. Due to the snowball sampling strategy, nearly 30% of the collected

[2]Work available on OpenAlex: `https://openalex.org/works/W1571611271`

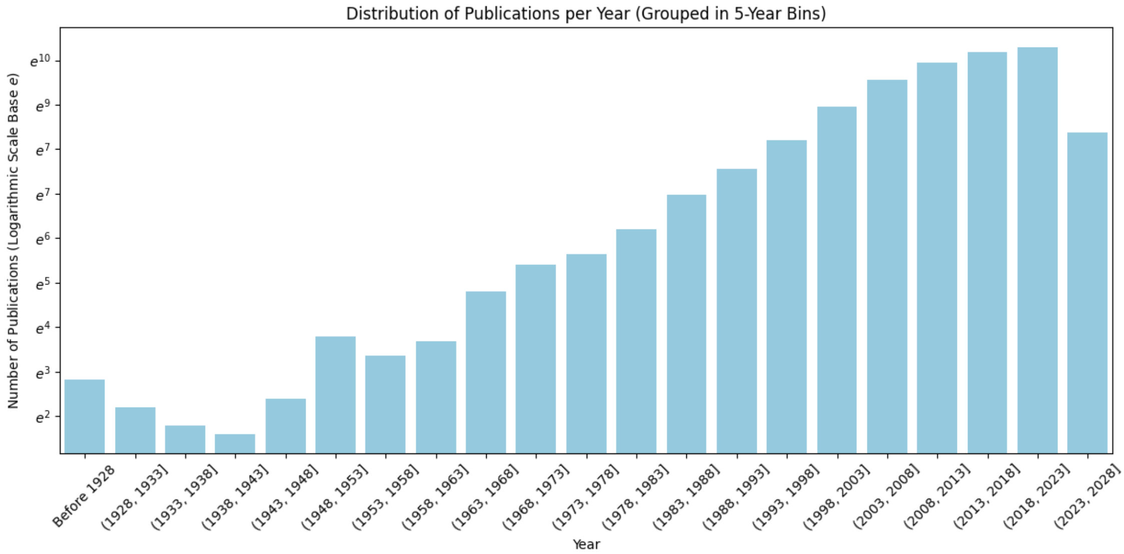

Figure 1. Number of Publications per Year

Table 1. Countries and their primary and secondary publications languages.

Country	Publications	Primary Language	Secondary Language
Brazil	29.609 (27.82%)	**English (85.97%)**	**Portuguese (12.34%)**
United States	10.801 (10.15%)	English (96.86%)	Others (1.14%)
Germany	3.830 (3.60%)	English (94.86%)	German (3.84%)
Spain	3.087 (2.90%)	English (97.83%)	Spanish (0.58%)
Canada	3.059 (2.87%)	English (98.73%)	Others (0.56%)
Great Britain	2.948 (2.77%)	English (98.41%)	Others (0.64%)
Italy	1.844 (1.73%)	English (96.15%)	Portuguese (1.68%)
France	1.595 (1.50%)	**English (78.75%)**	**French (19.00%)**
Portugal	1.248 (1.17%)	English (93.99%)	Portuguese (4.65%)
Others	48.418 (45.49%)	English (84.38%)	Portuguese (9.91%)

articles are from Brazilian authors. However, even within this subset, 85% of the Brazilian publications are written in English. This reflects the strong influence of English as the global language of scientific communication, particularly in the field of computer science. The implementation of the data collection process is publicly available in the GitHub repository[3].

4.2. Recommender System

In this work, we present an approach for modeling researchers and publications aiming at building a scientific article recommendation system. Our methodology consists of three main parts: modeling publications, modeling researchers, and applying temporal context.

[3]Repository available on *GitHub*: `https://github.com/BSBContreras/csv-openalex`

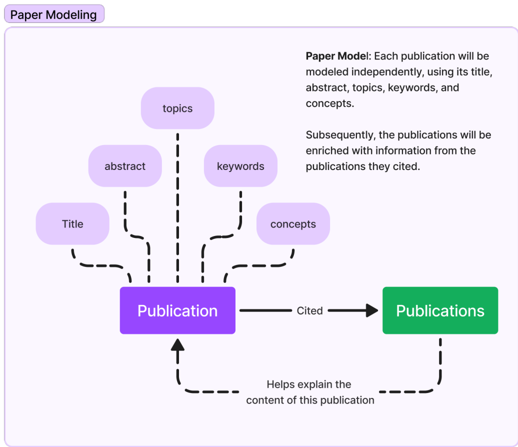

Figure 2. Papers Modeling

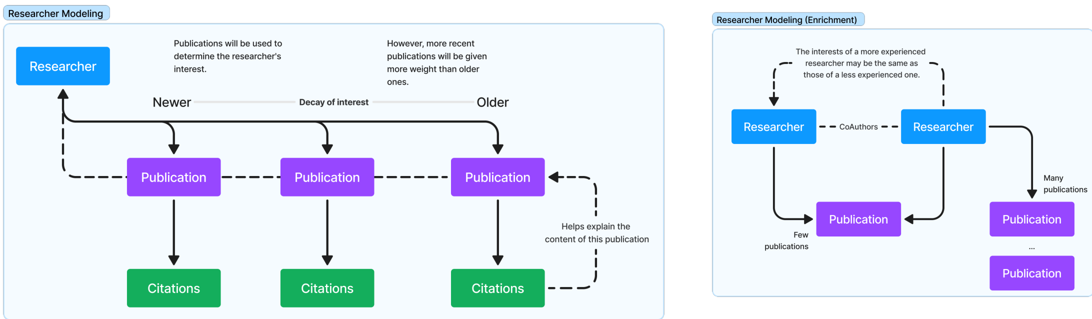

(a) Researchers Modeling

(b) Researcher Modeling Enrichment

Figure 3. Researcher Modeling

4.2.1. Papers Modeling

Each publication is modeled individually, using their characteristics such as title, abstract, topics, keywords, and concepts. This modeling provides a rich representation of the content of each article, facilitating the understanding of its specifics and context within the research field.

In addition, the publications are enriched with information from citations made to other works, as shown in Figure 2. *Node2Vec* embeddings are generated using the citation graph, where nodes represent publications and directed edges denote citation relationships. These embeddings capture the structural context of each publication within the citation network, encoding patterns such as influential citations, dense citation clusters, or peripheral works. By integrating these *Node2Vec* embeddings into the publication modeling process, we create a representation that combines both the semantic content of the publication and its relational context within the academic literature, enabling more accurate and context-aware recommendations.

4.2.2. Researchers Modeling

The modeling of researchers, illustrated in Figure 3a, is performed based on the publications associated with each researcher, as scientific output reflects research interests over time. To capture the temporal relevance of researchers' interests, we use a hyperbolic decay function, which assigns less weight to older interactions and places more emphasis on recent publications and interactions. This decay function (see Equation 3) allows the system to adjust the influence of each item over time, reflecting the natural changes in users' interests.

$$f(t) = \frac{1}{1 + \alpha(T - t)} \tag{3}$$

Where:

- $f(t)$ represents the weight of an item or interest at time t, serving as a measure of its current relevance,
- T is the current time at which the recommendation is being made in years,
- α is the decay parameter, which determines the rate at which relevance decreases over time.

In addition to publications, we use the references made by the authors to contextualize what the researcher was individually reading, enriching the interests of the researchers as a whole. Finally, the researcher modeling is further enriched by co-authorship information. Researchers with more experience may have more complex or diverse interests than those with fewer publications. Thus, we introduce a model that allows for sharing interests among co-authors, promoting a more comprehensive representation for researchers with few publications, addressing the cold start problem for authors, as showed in Figure 3b.

4.2.3. Recommendation Generation

$$\text{cosine similarity (A, B)} = \frac{A \cdot B}{\|A\|\|B\|} = \frac{\sum_i^n A_i B_i}{\sqrt{\sum_i^n A_i^2 B_i^2}} \tag{4}$$

The representation of a researcher is constructed from the condensed aggregate of their publications, generating a vector that synthesizes their research interests. To recommend new articles, we calculate the similarity between the researcher's representation and the representations of the publications in the set of recommendable articles (publications that the researcher has not referenced in any of their works, or their work). Similarity is calculated using the cosine similarity metric, as shown in Equation 4, which allows for identifying the publications most aligned with the researcher's current interests. In this way, the system provides relevant article recommendations that align with the researcher's academic trajectory and temporal interests.

4.3. Model Evaluation

The recommendation model was evaluated using a dataset of interactions between authors and articles. Initially, the data was split into training and testing sets (85% and 15% respectively), ensuring that the authors' interactions were stratified to maintain the original distribution. The evaluation process involved comparing the recommendations generated by the model with the authors' actual interactions in the test set.

For each author in the test set, the articles they interacted with were identified. Then, the model generated a list of recommendations, excluding the articles already known to the author in the training set. To evaluate the recommendations, a set of articles with which the author did not interact was randomly selected to simulate the model's search space. This approach ensures a more robust evaluation of recommendation quality by including non-interacted items in the sample, mitigating potential biases that could arise from considering only interacted items.

The evaluation metrics included precision, recall, F1-score, NDCG, and MRR at different levels (top-5, top-10, and top-20). Ablation experiments were also conducted to deactivate the time-aware contextualization of the system and examine its results.

5. Experiments

This section presents our experimental setup and evaluation of the proposed recommendation system, focusing on its adaptability to evolving research trends. We also assess the impact of various techniques, such as *Doc2Vec*, BERT, and hybrid models, on recommendation accuracy.

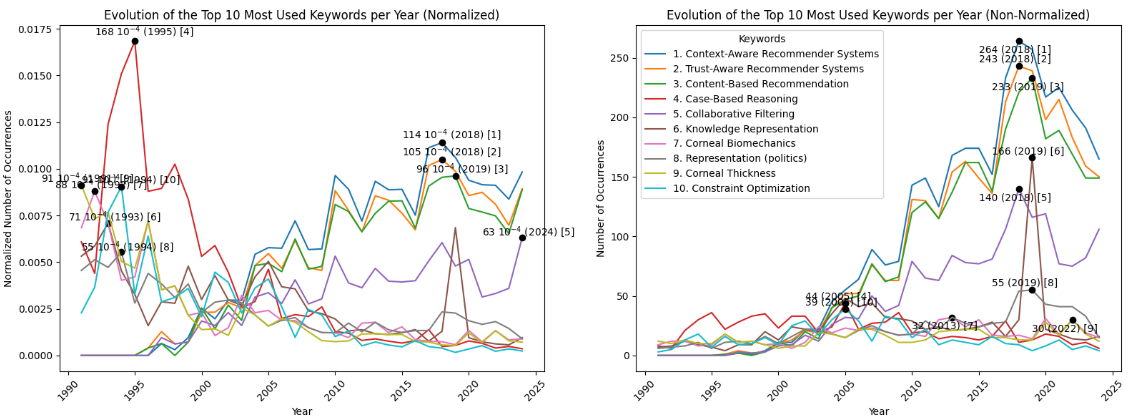

Figure 4. Evolution of the Most Used Keywords by Year

5.1. Evolution of Research Areas

We first explore the evolution of research areas in recommendation systems and related fields. The growth of scientific publications has driven the emergence of various themes, reflecting technological advancements and the adaptation of recommendation systems to evolving demands.

Publications were grouped by year, starting from 1990, covering 35 years. Keyword occurrences were counted annually, and the total number of publications per year was calculated. To analyze trends, keyword counts were normalized relative to the total publications each year. The top 10 most frequently used keywords overall were identified, and two graphs were generated:

- A normalized graph showing the evolution of keyword usage over time, adjusted for the total publications.
- A non-normalized graph showing raw keyword counts by year.

Figure 4 illustrates how the most commonly used keywords in scientific publications have evolved. The normalized graph highlights distinct trends: for example, a peak in the use of "Case-Based Reasoning" in 1995, followed by a period of more balanced topics around 2003–2004. From 2018–2019, there was a noticeable rise in the prominence of "Context-Aware Recommender Systems," "Trust-Aware Recommender Systems," and "Content-Based Recommendation."

These patterns indicate that frequently used keywords reflect the research interests and priorities of each era. Notably, many researchers focusing on modern recommendation systems previously explored topics like "Case-Based Reasoning," showcasing the evolution of research interests over time.

Table 2. Experiment Results

Metric	@K	Popularity	*Doc2Vec*	*Doc2Vec* + Time-aware		BERT	BERT + Time-aware	
	5	08.39%	10.15%	11.14%	Δ+09.75%	39.14%	43.89%	Δ+12.13%
Recall	10	08.39%	10.75%	11.63%	Δ+08.19%	40.30%	44.56%	Δ+10.56%
	20	08.39%	11.23%	11.96%	Δ+06.50%	43.19%	46.01%	Δ+06.52%
	5	06.17%	12.34%	13.40%	Δ+08.86%	41.76%	46.44%	Δ+11.20%
Precision	10	06.17%	11.32%	12.05%	Δ+06.45%	40.58%	44.89%	Δ+10.62%
	20	06.17%	09.87%	10.29%	Δ+04.25%	39.31%	42.65%	Δ+08.49%
	5	07.11%	11.14%	12.16%	Δ+09.22%	40.40%	45.12%	Δ+11.68%
F1-Score	10	07.11%	11.03%	11.84%	Δ+07.33%	40.43%	44.72%	Δ+10.59%
	20	07.11%	10.51%	11.06%	Δ+05.29%	41.15%	44.26%	Δ+07.55%
	5	05.13%	08.56%	09.81%	Δ+14.60%	33.72%	37.46%	Δ+11.09%
NDCG	10	05.56%	09.20%	10.36%	Δ+12.60%	34.60%	38.40%	Δ+10.98%
	20	05.60%	10.79%	11.91%	Δ+10.37%	36.44%	39.72%	Δ+09.01%
	5	10.63%	12.75%	13.80%	Δ+08.23%	43.39%	47.55%	Δ+09.87%
MRR	10	11.01%	12.99%	14.02%	Δ+07.92%	45.10%	48.76%	Δ+08.11%
	20	11.96%	13.56%	14.35%	Δ+05.59%	46.00%	49.35%	Δ+07.28%

5.2. Recommendation System Evaluation

The performance of the proposed recommendation system was evaluated through a series of experiments comparing different model configurations using standard evaluation metrics. Table 2 summarizes results for five metrics - Recall, Precision, F1-Score, NDCG, and MRR - evaluated at three cutoff levels (@K=5, 10, and 20). The configurations include a Popularity baseline, two variants of *Doc2Vec* (with and without time-awareness), and two variants of BERT (also with and without time-awareness).

The results demonstrate that incorporating time-awareness consistently enhances the performance of both *Doc2Vec* and BERT models across all metrics and cutoff levels. For example, *Doc2Vec* + Time-aware shows substantial improvements over the standard *Doc2Vec*, achieving up to a **9.75%** increase in Recall at @K=5. Similarly, BERT + Time-aware achieves even greater improvements, with a **12.13%** increase in Recall at the same cutoff level. These findings suggest that time-aware models are better suited for generating relevant recommendations, particularly in dynamic or time-sensitive contexts.

The trends for Precision and F1-Score align with those observed for Recall. Time-aware configurations consistently outperform their non-time-aware counterparts. For instance, BERT + Time-aware achieves a Precision of **46.44%** at @K=5, representing an **11.20%** improvement over standard BERT. The F1-Score, which balances Precision and Recall, reaches its highest value with BERT + Time-aware, further emphasizing the importance of incorporating temporal dynamics in recommendation models.

When examining ranking-sensitive metrics such as NDCG and MRR, the results reinforce the findings from Recall, Precision, and F1-Score. *Doc2Vec* + Time-aware improves NDCG by up to **14.60%** over the standard *Doc2Vec* at @K=5, indicating better alignment between top-ranked recommendations and user preferences. BERT + Time-aware achieves the highest NDCG and MRR scores, with MRR reaching **47.55%** at @K=5, a **9.87%** improvement over its non-time-aware counterpart.

In summary, the experimental results highlight the significant impact of time-awareness on recommendation quality across all configurations and metrics. While *Doc2Vec* + Time-aware delivers notable improvements over the basic *Doc2Vec*, BERT + Time-aware achieves the best overall performance. These findings underscore the importance of incorporating temporal dynamics into recommendation systems, particularly in domains where user preferences evolve or are influenced by time-sensitive factors.

6. Conclusions

This study demonstrated the effectiveness of a time-aware recommendation system in the academic domain, emphasizing the importance of incorporating the temporal dimension to align recommendations with the evolving trajectories of researchers. Significant improvements were observed in metrics such as Precision, Recall, and NDCG compared to traditional models that assume static user interests. These findings underscore the critical role of temporal modeling in capturing the dynamic nature of user priorities, particularly in contexts where interests change over time.

The analysis of publication data revealed distinct patterns in the evolution of researchers' interests, driven by factors such as new collaborations, advancements in emerging technologies, and adaptation to evolving scientific trends. This temporal perspective proved essential in capturing the complexity of academic trajectories, highlighting how researchers transition across different areas over time.

Overall, this study highlights that time-aware recommendation systems can not only enhance the relevance and accuracy of suggestions but also support the scientific community in efficiently discovering knowledge in a contextually meaningful manner. It is anticipated that these contributions will inspire novel approaches in recommendation system design and motivate future research on applying temporality to other domains and use cases.

References

Bai, X., Wang, M., Lee, I., Yang, Z., Kong, X., and Xia, F. (2019). Scientific paper recommendation: A survey. *IEEE Access*, 7:9324–9339.

Campos, P. G., Díez, F., and Cantador, I. (2013). Time-aware recommender systems: a comprehensive survey and analysis of existing evaluation protocols. *User Modeling and User-Adapted Interaction*, 24(1–2):67–119.

Chen, Y.-C., Hui, L., and Thaipisutikul, T. (2020). A collaborative filtering recommendation system with dynamic time decay. *The Journal of Supercomputing*, 77(1):244–262.

Devlin, J., Chang, M.-W., Lee, K., and Toutanova, K. (2019). BERT: Pre-training of deep bidirectional transformers for language understanding. In Burstein, J., Doran, C., and Solorio, T., editors, *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long and Short Papers)*, pages 4171–4186, Minneapolis, Minnesota. Association for Computational Linguistics.

Grover, A. and Leskovec, J. (2016). node2vec: Scalable feature learning for networks.

Gómez-Losada, A. and Duch-Brown, N. (2019). *Time Series Forecasting by Recommendation: An Empirical Analysis on Amazon Marketplace*, page 45–54. Springer International Publishing.

Ikoma, T. and Matsubara, S. (2023). Paper recommendation using citation contexts in scholarly documents. In Huang, C.-R., Harada, Y., Kim, J.-B., Chen, S., Hsu, Y.-Y., Chersoni, E., A, P., Zeng, W. H., Peng, B., Li, Y., and Li, J., editors, *Proceedings of the 37th Pacific Asia Conference on Language, Information and Computation*, pages 710–716, Hong Kong, China. Association for Computational Linguistics.

Jiang, C., Ma, X., Zeng, J., Zhang, Y., Yang, T., and Deng, Q. (2023). Taprec: time-aware paper recommendation via the modeling of researchers' dynamic preferences. *Scientometrics*, 128(6):3453–3471.

Kuznetsov, S. and Kordík, P. (2023). *Overcoming the Cold-Start Problem in Recommendation Systems with Ontologies and Knowledge Graphs*, page 591–603. Springer Nature Switzerland.

Le, Q. V. and Mikolov, T. (2014). Distributed representations of sentences and documents.

Ly, K., Kashnitsky, Y., Chamezopoulos, S., and Krzhizhanovskaya, V. (2024). Article classification with graph neural networks and multigraphs. In Calzolari, N., Kan, M.-Y., Hoste, V., Lenci, A., Sakti, S., and Xue, N., editors, *Proceedings of the 2024 Joint International Conference on Computational Linguistics, Language Resources and Evaluation (LREC-COLING 2024)*, pages 1539–1547, Torino, Italia. ELRA and ICCL.

Mena-Chalco, J. P., Digiampietri, L. A., Lopes, F. M., and Cesar, R. M. (2014). Brazilian bibliometric coauthorship networks. *Journal of the Association for Information Science and Technology*, 65(7):1424–1445.

Mikolov, T., Chen, K., Corrado, G., and Dean, J. (2013). Efficient estimation of word representations in vector space.

Perera, D. and Zimmermann, R. (2020). Towards comprehensive recommender systems: Time-aware unified recommendations based on listwise ranking of implicit cross-network data. *Proceedings of the AAAI Conference on Artificial Intelligence*, 34(01):189–197.