

# Freshness and Informativity Weighted Cognitive Extent and Its Correlation with Cumulative Citation Count

Zihe Wang<sup>1,†</sup>, Jian Wu<sup>2,\*,†</sup>

<sup>1</sup>Ohio State University, Columbus, OH, United States

<sup>2</sup>Old Dominion University, Norfolk, VA, United States

## Abstract

In this paper, we revisit cognitive extent, originally defined as the number of unique phrases in a quota. We introduce Freshness and Informativity Weighted Cognitive Extent (FICE), calculated based on two novel weighting factors, the lifetime ratio and informativity of scientific entities. We model the lifetime of each scientific entity as the time-dependent document frequency, which is fit by the composition of multiple Gaussian profiles. The lifetime ratio is then calculated as the cumulative document frequency at the publication time  $t_0$  divided by the cumulative document frequency over its entire lifetime. The informativity is calculated by normalizing the document frequency across all scientific entities recognized in a title. Using the ACL Anthology, we verified the trend formerly observed in several other domains that the number of unique scientific entities per quota increased gradually at a slower rate. We found that FICE exhibits a strong correlation with the average cumulative citation count within a quota. Our code is available at <https://github.com/ZiheHerzWang/Freshness-and-Informativity-Weighted-Cognitive-Extent>

## Keywords

cognitive extent, citation impact, entity recognition, document frequency

## 1. Introduction

Cognitive extent is an approach to quantify the extent of cognitive domains of scientific fields based on the concept of lexical diversity [1]. The metric was originally calculated by counting the number of unique concepts (phrases) appearing in the titles of statistically large unit quotas of scientific literature, which reflects the extent of the cognitive territory covered in that literature. Cognitive extent has been used as a representation of knowledge gained by scientists. It has been shown that cognitive extent calculated in multiple academic fields (Physics, Astronomy, and Biomedicine) grew at a slower rate.

However, this definition of cognitive extent has two limitations. First, it only accounts for the occurrence of a phrase as a dichotomous value within a quota and ignores *when* the phrase occurs. Specifically, a phrase is novel when it appears the first time in the title of Paper A, but when it appears again in the title of Paper B, it still maintains a level of freshness, and thus still reflects the scientist of Paper B's cognitive knowledge except that the knowledge is no longer new. Second, the definition treats all phrases with the same weight. However, certain phrases may be more informative than others. For example, at a particular time the phrase "entity recognition" occurred in many titles but the phrase "nuclearity rhetorical relation" occurred in only a small number of titles. From a reader's perspective, the latter phrase is more informative because the former phrase has been seen in many papers.

To overcome the limitations, we propose Freshness and Informativity Weighted Cognitive Extent (FICE), which is calculated as a weighted occurrence of disambiguated unique scientific entities extracted from paper titles in a quota. Different from the original cognitive extent, FICE accounts for contributions of *freshness* and *informativity* of scientific entities extracted from documents within a corpus. The freshness is based on the lifetime ratio and the informativity is based on time-dependent document

---

IMSC Workshop@JCDL 2024, December 20, 2024, Hong Kong, CN

\*Corresponding author.

†These authors contributed equally.

✉ wang.14629@osu.edu (Z. Wang); j1wu@odu.edu (J. Wu)

🌐 <https://www.cs.odu.edu/~jwu/> (J. Wu)

🆔 0009-0007-4293-3449 (Z. Wang); 0000-0003-0173-4463 (J. Wu)



© 2024 Copyright for this paper by its authors. Use permitted under Creative Commons License Attribution 4.0 International (CC BY 4.0).

frequency across scientific entities in a document. Here the document can be any form of scholarly text. Throughout this paper, we focus our study on paper titles.

We verified a previous finding using ACL Anthology papers that although the number of papers increased exponentially, the unique number of scientific entities per quota gradually increased with time at a slower rate. One property of FICE is its relationship with citation impact factors. We found that FICE exhibits a strong correlation with the logarithm of 5-year average cumulative citation counts for papers in the ACL Anthology.

## 2. Related Work

The definition of cognitive extent is closely relevant to lexical diversity, defined as the extent of vocabulary disparity within a given language sample [2]. Early metrics determined the lexical diversity only by single words [3, 4]. Berube et al. proposed the type-token ratio as a metric of lexical diversity [2]. Both methods were based on word-level tokens and ignored their connections, therefore they do not necessarily reflect the *knowledge*, which is better captured by phrases and entities.

Milojević was the first to propose using concepts (phrases) to quantify the extent of cognitive domains of scientific fields [1]. This method was later applied to study the properties of paper titles in various domains [5]. Recently, a method was proposed to use neural embedding of paper titles to represent the cognitive content of a cluster of papers [6]. Although neural embedding has been widely used to capture the semantics of text and compare the semantic similarities between two pieces of text, the embedding itself cannot be directly converted to a numeral that represents the cognitive extent of a single or a corpus of documents.

Bibliographic impact factors have been extensively studied and several citation impact factors have been proposed [7]. The total number of citations (raw citations) is usually criticized as a good indicator because of domain discrepancies, data completeness, and other random factors. The average number of citations per publication of a research unit is frequently used but also criticized because the average value can be biased due to the skewness of citation distribution within the research unit [8]. However, it was argued that short-term citations can be considered as currency on the research front and long-term citations can contribute to the codification of knowledge claims into bodies of knowledge [9]. Determining the exact boundary between short-term and long-term is non-trivial. Here we adopt  $C_5(y)$ , which is the 5-year average cumulative citation count of papers within a quota as an estimate of the quota’s average impact over the short- and long-term.

## 3. Methodology

### 3.1. Scientific Entity Recognition and Disambiguation

A scientific entity is defined as a noun phrase that delivers domain knowledge of interest [10, 11]. Scientific entities can be extracted using sequence tagging models, and constructing knowledge graphs, e.g., [12, 11]. Recently, large language models (LLMs) have shown superior performance on named entity recognition, e.g., [13, 14].

Two scientific entities may have similar semantics. To investigate the impact of entity disambiguation on the results, we conflate entities with similar semantic meanings. This was treated as a classification problem using a thresholded method on the similarity scores calculated using the Cross Encoder [15], a model that takes two entity names and outputs a similarity score. The threshold was calibrated based on the classification performance evaluated on a set of manually labeled entity pairs.

### 3.2. Lifetime Ratio

The lifetime of a scientific entity is defined as the period during which it appears in at least one document (in our case, a paper title). Here, we assume that all scientific entities have a finite lifetime, meaning there is a time point  $t_s$  when a scientific entity first appears and another time point  $t_e$  after which it no

longer presents in any documents. We borrowed the concept of document frequency from information retrieval to indicate whether a scientific entity  $e$  is still within its lifetime ( $df(e, t) > 0$ ) or its lifetime ends ( $df(e, t) \leq 0$ ). Here, we define the time-dependent document frequency  $df(e, t)$ , which is the number of documents that contain the scientific entity  $e$  and are published at time  $t$  (e.g., a certain year). The lifetime ratio of a scientific entity  $e$  at  $t_0$  is then defined as the number of accumulated documents up to  $t_0$  divided by the total number of documents that contain  $e$  over its entire lifetime  $[t_s, t_e]$ ,

$$r(e, t_0) = \frac{\sum_{t_s}^{t_0} df(e, t)}{\sum_{t_s}^{t_e} df(e, t)}. \quad (1)$$

The freshness of  $e$  is then calculated as  $1 - r(e, t_0)$ . Because any corpus can cover a limited period, the lifetime ratio can only be calculated based on the *observable period* covered by the corpus. We then model  $df(e, t)$  for each  $e$  as a composite of analytical profiles. By definition, the lifetime ratio provides an estimation of the relative freshness of a scientific entity. Specifically, a low lifetime ratio indicates a scientific entity is relatively new (so  $1 - r(e, t_0)$  is large) and a high lifetime ratio indicates a scientific entity is not new anymore. The model fitting provides a prediction of  $df(e, t)$  beyond the observable period based on data in the observable period.

### 3.3. Informativity Weight

In linguistics, informativity concerns the extent to which the contents of a text are already known or expected as compared to unknown or unexpected [16]. The informativity of paper titles has been commonly calculated by counting the number of “substantive” words. A diachronic analysis of informativity was conducted on chemical paper titles [17], in which non-substantive words were defined as words that convey little or no information, such as articles, prepositions, conjunctions, pronouns, and auxiliary verbs. This objective approach was then extensively used by scholars to study title informativity. We argue that whether a text (word, phrase, entity name) is informative or not depends not only on its semantics but also on the relative frequency it appears in existing papers. At a certain time point  $t$ , a scientific entity that appears in a large number of documents conveys relatively less new information than an entity that appears only in a few documents. We calculate informativity as the cumulative time-dependent document frequency  $DF$  normalized by its range of  $[DF_{\min}, DF_{\max}]$  across all scientific entities in a document,

$$\begin{aligned} w(e, t_0) &= 1 - \frac{DF - DF_{\min}}{DF_{\max} - DF_{\min}}, \quad DF(e, t_0) = \sum_{t_s}^{t_0} df(e, t) \\ DF_{\min} &= \min \{DF(e_i, t_0), e_i \in E\} \\ DF_{\max} &= \max \{DF(e_i, t_0), e_i \in E\} \end{aligned} \quad (2)$$

in which  $E$  is the set of scientific entities extracted from a document. The FICE of documents in a quota  $Q$  is calculated as

$$FICE = \sum_{d \in Q} \sum_{e \in d} w(e, t_d) (1 - r(e, t_d)), \quad (3)$$


in which  $t_d$  is the time when document  $d$  was published.

## 4. Data Processing


### 4.1. Data Collection

We downloaded the metadata of all papers in the ACL Anthology [18] in BibTeX format<sup>1</sup>. The publication year and title of each paper were extracted from the BibTeX file. The number of papers published each year from 1952 to 2020 is shown in Figure 1.

<sup>1</sup><https://aclanthology.org/anthology+abstracts.bib.gz>



**Figure 1:** The number of papers, scientific entities (undisambiguated), and disambiguated entities in the ACL Corpus.



**Figure 2:** The document frequency chart (blue) of an entity named *machine learning* and a fitting with 4 Gaussian profiles.

## 4.2. Scientific Entity Recognition

We compared three off-the-shelf models for scientific entity recognition. (1) **GPT-4** [19]. We construct a zero-shot prompting template to extract scientific entities using GPT-4. The temperature is set to zero to ensure consistent outputs. (2) **SciBERT** [20]. We used the named entity recognition implementation from the Hugging Face library developed by AllenAI<sup>2</sup>. (3) **SpaCy** [21]. Entity recognition is performed by invoking the entity recognition module from the Hugging Face library<sup>3</sup>.

To compare the performance of these models, we built a small benchmark dataset by manually annotating 200 titles randomly selected from all ACL Anthology papers, following the annotation guidelines in Wu et al. [22]. The F1-scores achieved by GPT-4, SciBERT, and SpaCy are 0.66, 0.05, and 0.07, respectively, indicating that GPT-4 outperforms the other two models, so we adopt GPT-4 for recognizing scientific entities of all paper titles. The numbers of papers and scientific entities recognized are shown in Figure 1. The average number of scientific entities per title is about 3.

## 4.3. Scientific Entity Disambiguation

Using the method described in Section 3.1, the threshold was calibrated based on the classification performance evaluated against 180 manually labeled entity pairs. Entity pairs extracted from titles within a quota were automatically labeled as “similar” or “not similar” based on the calculated similarity scores and a threshold of 0.5.


## 4.4. Document Frequency Curve Fitting

For each scientific entity  $e$ , we count the number of paper titles that contain  $e$  at Year  $t$  and plot the document frequency chart. An example is shown in Figure 2.

The document frequency chart for each scientific entity is fit using a composite of Gaussian profiles each having three parameters, peak, mean, and dispersion. Our fitting program employs a dynamic

<sup>2</sup>[https://huggingface.co/allenai/scibert\\_scivocab\\_cased](https://huggingface.co/allenai/scibert_scivocab_cased)

<sup>3</sup>[https://huggingface.co/spacy/en\\_core\\_web\\_sm](https://huggingface.co/spacy/en_core_web_sm)



**Figure 3:** The FICE calculated using disambiguated scientific entities (black dots) with  $|Q| = 125, 250, 500$ . The red and blue curves are the polynomial fittings of disambiguated and undisambiguated entities, respectively. For each year, we only plot data points that represent full quotas of papers.

tuning approach, where the number of peaks is inferred from the data using an algorithm based on a comparison of neighboring values<sup>4</sup>. The center of each Gaussian is initialized at the detected peak position; the amplitude is initialized randomly within a defined range, and the width of each Gaussian is initialized based on the year range divided by the number of peaks. The fitting process iteratively updates the parameters using gradient descent. We used ADAM as the optimizer [23] and the mean squared error (MSE) as the loss function. To prevent overfitting, a regularization term is added to the loss function, which penalizes excessively large or narrow peaks. The number of epochs and parameters was automatically adjusted by the fitting algorithm.

#### 4.5. Calculating FICE

After fitting the document frequency chart for a scientific entity  $e$ , the starting point for  $t_s$  is determined as the year when  $e$  first appeared and  $t_e$  is determined when the predicted document frequency is less than 1, which may be beyond the time span of the observable period. We calculated the lifetime ratio for each  $e$  in a title  $d$  using Eq. (1) and the informativeness weight using Eq. (2). The FICE for a given quota  $Q$  is calculated using Eq. (3).

#### 4.6. Cumulative Citation Count $C_5$

We obtain citations for each paper each year from the Semantic Scholarly Graph API [24]. Therefore, the average 5-year cumulative citation count in 2015 is  $C_5(2015) = \sum_{y=2015}^{2019} c(y)$ , in which  $c(y)$  is the citation received by a paper in Year  $y$ .

<sup>4</sup>[https://docs.scipy.org/doc/scipy/reference/generated/scipy.signal.find\\_peaks.html](https://docs.scipy.org/doc/scipy/reference/generated/scipy.signal.find_peaks.html)

**Table 1**

The linear fit slopes of entity-based cognitive extents.

Year Range	$ Q  = 125$	$ Q  = 250$	$ Q  = 500$
1980–2000	1.19	7.21	15.88
2000–2020	6.80	3.09	6.10

## 5. Results

### 5.1. Growth of Scientific Entity Diversity

To distinguish from the original cognitive extent, we define entity-based cognitive extent as the unique number of scientific entities extracted from paper titles within a quota. Similar to lexical diversity, entity-based cognitive extent can be seen as a measure of the scientific entity diversity. By plotting entity-based cognitive extent over time, we found that it gradually increases with time, which is consistent with the trend of the original cognitive extent based on paper titles in Astronomy, Physics, and Biomedical [1]. We compared the trends calculated using disambiguated and undisambiguated scientific entities (Figure 4) and found that undisambiguated entities linearly shifted the curve up but did not significantly change the growth rate.

To investigate how the growth rate of entity-based cognitive extent increases with time, we fit the disambiguated entity-based data points using a linear function from 1980 to 2000 and another linear function from 2000 to 2020, respectively. The slopes obtained for various quotas are tabulated in Table 1. The results indicate that the entity-based cognitive extent increases at a slower rate.


The entity-based cognitive extent vs. year relations for three quota sizes are illustrated in Figure 4. Similar to the original cognitive extent, calculating the entity-based cognitive extent within a quota is important to generate a consistent measure of the metric. Our quota is different from the ones used in Milojević in two aspects. First, instead of a fixed number of phrases, we use a fixed number of titles. This is because the informativity weight is normalized within scientific entities of a particular title before weights of multiple titles are aggregated. Second, the quota sizes used in our study are smaller than the quota used in the original cognitive extent, which is 3000 – 10000. For example,  $|Q| = 500$  converts to about 1500 scientific entities. Low quota data points may suffer a saturation, meaning that cognitive extent value increases with quota size for a given time point, which is seen in Figure 4. To fix the problem, a linear correction factor can be applied to “lift” data points in low quota to be aligned with unsaturated data points. The exact correction factor may be domain-dependent and will be investigated with a larger corpus in future studies.

### 5.2. Correlation Between FICE with $C_5$

We found that FICE exhibits a strong positive correlation with  $\log C_5$  as shown in Figure 4. The plots were made by first ranking papers by  $C_5(2015)$  in ascending order and binning the sequence by a given quota. For each data point, the  $x$ -coordinate is calculated as the logarithmic value of the average of  $C_5(2015)$  of papers in a quota, and the  $y$ -coordinate is the arithmetic average FICE of papers in the same quota. The error bars are calculated as the standard deviation assuming a Gaussian distribution. The Spearman correlation coefficients between FICE and  $\log(C_5)$  for three quota are shown in Table 2. To test the influence of disambiguation on the correlation. We calculated FICE without entity disambiguation and obtained similar Spearman correlation coefficients.

Note that Figure 4 reveals a *collective* instead of an individual correlation because FICE represents the weighted cognitive extent in a quota. This correlation does not apply to individual papers because of the small number of scientific entities in a title. Therefore, the correlation does not imply that one could increase the citation impact by simply using never-existing entity names. If the paper lacks true novelty and significant contributions, newly introduced entity names are unlikely to be adopted in subsequent research, rendering them “transient” with a short lifetime and a minimal contribution to FICE.





**Figure 4:** Average FICE calculated using undisambiguated entities per quota vs. the  $\log C_5$ . Paper titles are grouped into a bin size of 250.

We compare FICE with three simplified versions and demonstrate the contribution of the lifetime ratio and informativity weight in the correlation above. These simplified versions are (1) **Dichotomous Entity-based Cognitive Extent**, calculated by adding the number of disambiguated unique scientific entities in a quota. (2) **Weight Only**, calculated by simply summing up the normalized weights (Eq.(2)) of all scientific entities. (3) **Lifetime Ratio Only**, calculated by summing up the unweighted lifetime ratios  $1 - r(e, t)$  of all scientific entities. Table 2 indicates that FICE exhibits the strongest correlation against all baseline models. Both the lifetime ratio and the informativity weight contribute to this strong correlation. The quota size will influence the correlation coefficients. In particular, all simplified versions exhibit a strong correlation with  $|Q| = 500$ . The lifetime ratio consistently exhibits a strong correlation with  $\log C_5$ .

**Table 2**

Spearman correlation coefficients and  $p$ -values between FICE and three simplified versions: Dicho (Dichotomous), Weight (Weight only), and L. Ratio (Lifetime Ratio only).

Method	$ Q  = 125$	$ Q  = 250$	$ Q  = 500$
Dicho	-0.206(0.371)	-0.333(0.317)	-0.703(0.119)
Weight	0.228(0.320)	0.334(0.316)	0.715(0.110)
L. Ratio	0.603(0.004)	0.705(0.015)	0.744(0.090)
<b>FICE</b>	<b>0.766(&lt; 0.001)</b>	<b>0.748(0.008)</b>	<b>0.717(0.109)</b>

## 6. Conclusion

We proposed FICE, which extends the original cognitive extent. FICE is calculated based on the lifetime ratio and informativity of scientific entities extracted from paper titles within a quota. Using ACL Anthology, we found the number of unique scientific entities per quota increased with time, consistent with previous observations in other disciplines. FICE exhibits a strong positive correlation with the average 5-year cumulative citation count, which may be used for predicting collective citations for trending topics.

## 7. Acknowledgments

The paper is presented at the second Workshop on “Innovation Measurement for Scientific Communication (IMSC) in the Era of Big Data” at 2024 ACM/IEEE Joint Conference on Digital Libraries (JCDL).

## Declaration on Generative AI

The authors have not employed any Generative AI tools.

## References

- [1] S. Milojević, Quantifying the cognitive extent of science, *Journal of Informetrics* 9 (2015) 962–973.
- [2] N. Bérubé, M. Sainte-Marie, P. Mongeon, V. Larivière, Words by the tail: Assessing lexical diversity in scholarly titles using frequency-rank distribution tail fits, *PLOS ONE* 13 (2018) 1–31. URL: <https://doi.org/10.1371/journal.pone.0197775>. doi:10.1371/journal.pone.0197775.
- [3] S. Jarvis, M. Daller, Defining and measuring lexical diversity, *Vocabulary knowledge: Human ratings and automated measures*. Amsterdam, The Netherlands (2013).
- [4] D. Malvern, B. Richards, N. Chipere, P. Durán, *Lexical diversity and language development*, Springer, 2004.
- [5] S. Milojević, The length and semantic structure of article titles—evolving disciplinary practices and correlations with impact, *Frontiers in Research Metrics and Analytics* 2 (2017). URL: <https://www.frontiersin.org/journals/research-metrics-and-analytics/articles/10.3389/frma.2017.00002>. doi:10.3389/frma.2017.00002.
- [6] J. Eykens, R. Guns, T. C. Engels, F. Vandermoere, Cognitive and interdisciplinary mobility in the social sciences and humanities: Traces of increased boundary crossing, *Journal of Information Science* (2023) 01655515231171086. URL: <https://doi.org/10.1177/01655515231171086>. doi:10.1177/01655515231171086.
- [7] L. Waltman, A review of the literature on citation impact indicators, *Journal of Informetrics* 10 (2016) 365–391. URL: <https://www.sciencedirect.com/science/article/pii/S1751157715300900>. doi:<https://doi.org/10.1016/j.joi.2016.02.007>.
- [8] P. Albarrán, J. A. Crespo, I. Ortuño, J. Ruiz-Castillo, The skewness of science in 219 sub-fields and a number of aggregates, *Scientometrics* 88 (2011) 385–397.
- [9] L. Leydesdorff, L. Bornmann, J. A. Comins, S. Milojević, Citations: Indicators of quality? the impact fallacy, *Frontiers in Research Metrics and Analytics* 1 (2016). URL: <https://www.frontiersin.org/journals/research-metrics-and-analytics/articles/10.3389/frma.2016.00001>. doi:10.3389/frma.2016.00001.
- [10] J. Wu, M. R. Ul Hoque, G. W. Reiske, M. C. Weigle, B. T. Bradshaw, H. D. Gaff, J. Li, C. Kwan, A comparative study of sequence tagging methods for domain knowledge entity recognition in biomedical papers, in: *Proceedings of the ACM/IEEE Joint Conference on Digital Libraries in 2020*, 2020, pp. 397–400.
- [11] Z. Hong, R. Tchoua, K. Chard, I. Foster, Sciner: Extracting named entities from scientific literature, in: V. V. Krzhizhanovskaya, G. Závodszky, M. H. Lees, J. J. Dongarra, P. M. A. Sloot, S. Brissos, J. Teixeira (Eds.), *Computational Science – ICCS 2020*, Springer International Publishing, Cham, 2020, pp. 308–321.
- [12] J. Li, A. Sun, J. Han, C. Li, A survey on deep learning for named entity recognition, *IEEE transactions on knowledge and data engineering* 34 (2020) 50–70.
- [13] S. Wang, X. Sun, X. Li, R. Ouyang, F. Wu, T. Zhang, J. Li, G. Wang, Gpt-ner: Named entity recognition via large language models, *arXiv preprint arXiv:2304.10428* (2023).
- [14] H. Sousa, N. Guimarães, A. Jorge, R. Campos, Gpt struct me: Probing gpt models on narrative entity extraction, in: *2023 IEEE/WIC International Conference on Web Intelligence and Intelligent Agent Technology (WI-IAT)*, 2023, pp. 383–387. doi:10.1109/WI-IAT59888.2023.00063.
- [15] Y. Li, J. Li, Y. Suhara, A. Doan, W.-C. Tan, Deep entity matching with pre-trained language models, *arXiv preprint arXiv:2004.00584* (2020).
- [16] R.-A. De Beaugrande, W. U. Dressler, *Introduction to text linguistics*, volume 1, longman London, 1981.



- [17] J. J. Tocatlian, Are titles of chemical papers becoming more informative?, *Journal of the American Society for Information Science* 21 (1970) 345–350.
- [18] D. R. Radev, P. Muthukrishnan, V. Qazvinian, A. Abu-Jbara, The acl anthology network corpus, *Language Resources and Evaluation* 47 (2013) 919–944.
- [19] T. B. Brown, Language models are few-shot learners, *arXiv preprint arXiv:2005.14165* (2020).
- [20] I. Beltagy, K. Lo, A. Cohan, SciBERT: A pretrained language model for scientific text, in: *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing (EMNLP-IJCNLP)*, Association for Computational Linguistics, Hong Kong, China, 2019, pp. 3615–3620. URL: <https://www.aclweb.org/anthology/D19-1371>. doi:10.18653/v1/D19-1371.
- [21] Matthew Honnibal, spaCy’s NER Model, 2024.
- [22] J. Wu, M. R. Ul Hoque, G. W. Reiske, M. C. Weigle, B. T. Bradshaw, H. D. Gaff, J. Li, C. Kwan, A comparative study of sequence tagging methods for domain knowledge entity recognition in biomedical papers, in: *Proceedings of the ACM/IEEE Joint Conference on Digital Libraries in 2020, JCDL ’20*, Association for Computing Machinery, New York, NY, USA, 2020, p. 397–400. URL: <https://doi.org/10.1145/3383583.3398602>. doi:10.1145/3383583.3398602.
- [23] D. P. Kingma, J. Ba, Adam: A method for stochastic optimization, in: Y. Bengio, Y. LeCun (Eds.), *3rd International Conference on Learning Representations, ICLR 2015, San Diego, CA, USA, May 7–9, 2015, Conference Track Proceedings*, 2015. URL: <http://arxiv.org/abs/1412.6980>.
- [24] AllenAI, Semantic Scholar API, <https://www.semanticscholar.org/product/api>, 2024.