

SciDef: Datasets and Tools for Automated Definition Extraction from Scientific Literature with LLMs

Filip Kučera
National Institute of Informatics (NII)
Tokyo, Japan
University of Würzburg
Würzburg, Germany
f.kucera@media-bias-research.org

Christoph Mandl
National Institute of Informatics (NII)
Tokyo, Japan
University of Passau
Passau, Germany
c.mandl@media-bias-research.org

Isao Echizen
National Institute of Informatics (NII)
Tokyo, Japan
iechizen@nii.ac.jp

Radu Timofte
University of Würzburg (JMU)
Würzburg, Germany
radu.timofte@uni-wuerzburg.de

Timo Spinde
National Institute of Informatics (NII)
Tokyo, Japan
t.spinde@media-bias-research.org

Abstract

Scientific concepts are often defined inconsistently across papers, making it difficult to compare findings, reuse terminology, and build reliable downstream resources. We present SciDef, a resource suite for scientific definition extraction. The suite contains DefExtra, a benchmark of 268 human-validated author-stated definitions from 75 academic papers; DefSim, 60 human-labeled definition-pair similarity judgments; and an open LLM-based pipeline for PDF pre-processing, chunking, definition extraction, prompt optimization, and evaluation. We validate the resources by benchmarking 16 language models across prompting strategies and chunking schemes. The strongest set-level configuration achieves a score of 0.397, while the highest-coverage configuration matches at least one prediction to 86.4% of gold definitions but over-generates candidate definitions. We further show that an NLI-based matching metric agrees strongly with human DefSim judgments. These results position SciDef as a reusable benchmark and tooling layer for definition-centric literature analysis, while highlighting relevance-aware filtering as the key bottleneck for fully automatic definition extraction. Code & datasets are available at <https://github.com/Media-Bias-Group/SciDef>.

CCS Concepts

• **Information systems** → *Data mining*; **Information extraction**;
• **Computing methodologies** → *Natural language processing*;
Language resources.

Keywords

definition extraction, large language models, scientific literature mining, evaluation metrics, natural language inference, semantic similarity

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than the author(s) must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.

CIKM '26, Rome, Italy

© 2026 Copyright held by the owner/author(s). Publication rights licensed to ACM.

ACM Reference Format:

Filip Kučera, Christoph Mandl, Isao Echizen, Radu Timofte, and Timo Spinde. 2026. SciDef: Datasets and Tools for Automated Definition Extraction from Scientific Literature with LLMs. In *Proceedings of The 35th ACM International Conference on Information and Knowledge Management (CIKM '26)*. ACM, New York, NY, USA, 6 pages.

1 Introduction

Definitions are foundational to scientific work because they determine how concepts are introduced, interpreted, compared, and reused across studies [13, 22, 33]. Without shared conceptual structures, research can become fragmented and difficult to reproduce [8]. However, as publication volumes grow, manually finding and consolidating definitions for a given keyword becomes increasingly infeasible [10]. Motivated by this growing problem, we propose developing reusable resources for extracting, comparing, and evaluating scientific definitions from academic literature.

For definition extraction, large language models (LLMs) are promising tools, as they can recognize definitional statements beyond surface-level keyword patterns [35]. Yet, reproducible research on the extraction of scientific definitions using LLMs remains limited. Existing work lacks a public benchmark of definitions extracted from academic papers, a dedicated dataset for evaluating definition similarity, and reusable infrastructure for comparing extraction pipelines and prompting strategies [2, 13, 22, 37, 40, 47].

Therefore, we release a resource suite for extracting and evaluating scientific definitions. *DefExtra* is a benchmark of 268 human-extracted definitions from 75 curated academic papers, including context spans, metadata, and explicit/implicit definition labels. *DefSim* contains 60 definition pairs with human semantic-similarity labels for validating definition-matching metrics. *SciDef* is a reproducible LLM pipeline that supports extractor and metric comparison, as well as definition collection for literature reviews, taxonomy construction, ontology engineering, and domain mapping.

We evaluate the resource suite using media bias as a challenging case study, since the domain is widely studied but inconsistently defined across disciplines [36]. Using *SciDef*, we benchmark 16 language models under multiple prompting strategies, including one-step, multi-step, few-shot, and DSPy-optimized prompting [16]. We also compare similarity metrics and show that an NLI-based

method aligns most reliably with human judgments from *DefSim*. Our extraction experiments show that LLMs can recover many scientific definitions, with the highest-recall configuration extracting 86.4% of definitions from *DefExtra*. At the same time, models often over-generate candidate definitions, highlighting relevance-aware filtering as a key open challenge.

The key contributions of this resource paper are:

- **DefExtra**, a benchmark of 268 human-extracted definitions from 75 academic papers, including source metadata, definition type labels, out-of-domain indicators, and copyright-protecting metadata for context reconstruction.
- **DefSim**, a task-specific validation resource of 60 definition-pair similarity judgments for evaluating semantic matching metrics on extracted scientific definitions.
- **SciDef**, an open, documented, and versioned LLM-based pipeline for PDF preprocessing, chunking, definition extraction, prompt optimization, and evaluation.
- **Resource validation experiments** showing how the released datasets and pipeline support comparison across 16 LLMs, multiple prompting strategies, chunking schemes, and semantic matching metrics.

2 Related Work

Definitions enable consistent communication and reuse of scientific concepts [8], but terminological consensus is difficult in many fields [15, 23]. The domain of media bias illustrates this problem: prior work defines biases such as partisan reporting [27], linguistic bias [38], or bias by word choice [12], which describe similar concepts under different names, limiting comparability [36]. Systematic reviews and content analyses can consolidate such definitions, but they are labor-intensive and difficult to scale [18].

To find solutions to this problem, prior definition extraction work includes structured systems such as DefIE [7], sequence-labeling models [44], term-extraction and prompt-engineering approaches [2, 43, 47], and broader ontology-learning methods [25]. Existing resources include WCL for Wikipedia-based definition and hypernym extraction [28, 29], W00 for term/definition labels in ACL workshop papers [14], DEFT for free and semi-structured text from textbooks and SEC filings [34], and dictionaries such as Oxford and Urban Dictionary [31, 42]. TaxoMatic uses LLM prompting for taxonomy-oriented definition extraction [37].

Finally, to compare contents such as definitions, existing work explored various semantic similarity metrics. Transformer embeddings and LLM representations improve over lexical matching [1, 4, 11, 32], while NLI models capture entailment-style relations useful for asymmetric or partial matches [3, 5, 46]. Benchmarks such as STS3k [10], SICK [26], MSRP [9], and QQP [6] support general metric selection, but do not evaluate scientific definition pairs. *DefSim* provides such task-specific validation data.

3 Resource Construction and Evaluation

Our resources comprise two datasets, *DefExtra* and *DefSim*, and the LLM-based *SciDef* pipeline. *DefExtra* benchmarks scientific-definition extractors; *DefSim* validates definition-matching metrics; and *SciDef* provides a reproducible workflow for PDF preprocessing, chunking, LLM extraction, prompt optimization, and evaluation.

Source	Size	Input → Label(s)
DefExtra: extraction benchmark		
[37]	113 papers / 123 defs.	Paper ID → author-stated and annotator-synthesized definitions
This work	75 papers / 268 defs.	Paper ID → term, context, definition, exp./imp. type, metadata, OOD flag
DefSim: definition similarity benchmark		
This work	60 pairs	Definition pair + context → human label (1–5), NLI score/label (1–5)

Table 1: Overview of the two released datasets. The source column and color coding distinguish prior work from this work; OOD = out-of-domain.

Availability, licensing, ethics, and provenance. Following FAIR [45], *DefExtra*, *DefSim*, and *SciDef* are archived with persistent DOIs [19–21] and released through Zenodo, Hugging Face, and GitHub; *SciDef* is released under Apache-2.0, and *DefExtra/DefSim* under CC BY 4.0. To reduce copyright risk in *DefExtra*, we release paper identifiers, metadata, and hydration scripts instead of source PDFs. The datasets contain no private or sensitive information. Human annotation was conducted by authors and academic collaborators; under the authors’ institutional guidelines, this work did not require IRB review. Provenance, schemas, hydration instructions, full appendix, and run configurations are documented in the released repositories.

3.1 DefExtra: Definition Extraction Benchmark

DefExtra is a benchmark for extracting author-stated definitions from academic papers. It focuses primarily on media bias, a domain where definitions are diverse, contested, and spread across disciplinary boundaries [39]. As summarized in Table 1, *DefExtra* builds on the TaxoMatic annotation effort [37], which provided a media-bias-centered search space, Semantic Scholar retrieval, GROBID preprocessing, and first-pass annotations. In TaxoMatic, terms from an existing media-bias taxonomy [36] were expanded with GPT-3.5-turbo, open-access PDFs were parsed with GROBID [24], and six trained annotators¹ screened highly cited papers for relevance and extracted initial definitions from relevant full texts.

This work builds on these initial annotations and builds the *DefExtra* benchmark by implementing stricter validation, reannotation, and extension. Two authors of this work manually filtered the candidate papers and definitions to retain only author-stated definitions and exclude loose descriptions, annotator-synthesized definitions, and definitions that could not be verified in the source text. We validated definitions from retained papers and, where necessary, reannotated them from the original paper text; we kept only items on which both authors agreed. From the original 113 papers, we kept only 60 by this procedure. To broaden the benchmark, we further annotated 15 additional papers, including out-of-domain papers from areas such as energy, ecology, and medicine. These additions were annotated by 11 annotators selected under the same criteria as in TaxoMatic [37], then checked by two authors and retained only under unanimous agreement. After validation, extension, and duplicate removal, *DefExtra* contains 268 definitions from

¹Chosen among past or present collaborators with academic media-bias domain expertise.

75 curated papers, published between 1987 and 2025. All extraction and validation were conducted using Label Studio [41].

Each annotation contains the definition text, source metadata, a context span, and a definition-type label. The context span includes the surrounding sentences, supporting source localization and distinguishing grounded extractions from hallucinated candidates. We labeled the definitions as *explicit* when they can be quoted directly from the paper, and as *implicit* when a clear author-provided definition is lightly reworded while preserving its meaning, e.g., by removing irrelevant parenthetical notes. We marked out-of-domain papers in the released data (i.e., not media-bias related).

3.2 DefSim: Definition Similarity Resource

DefSim supports evaluation of extracted definitions by providing human judgment of definition-pair similarity. This is necessary because LLM outputs may not match human annotations verbatim: a prediction may paraphrase the ground truth, partially overlap with it, or express a narrower concept. We construct *DefSim* from 60 definition pairs drawn from *DefExtra* “gold” annotations and top-performing extractor predictions described in Section 3.3. The set contains 30 gold-prediction, 15 gold-gold, and 15 prediction-prediction pairs, covering exact, partial, and non-matching cases.

Each pair was labeled by three authors of this paper. Annotators saw both definitions and their local contexts, and rated the conceptual similarity of the defined terms on a 1–5 scale, where low scores indicate unrelated definitions, intermediate scores indicate partial overlap or scope differences, and high scores indicate semantic equivalence. We release the mean and majority human labels, along with pair metadata and our selected NLI metric score (Section 3.3).

3.3 SciDef: Extraction and Evaluation Pipeline

We release *SciDef*, an extensible extraction and evaluation pipeline for scientific definitions. Given a paper, it parses the document, chunks the text, prompts an LLM to extract candidate definitions, and optionally evaluates the resulting definition set against *DefExtra*. The pipeline is modular: users can replace the model backend, chunking strategy, prompt template, or evaluation metric.

We parse PDFs with GROBID [24] and evaluate four paper-chunking strategies: section-level, paragraph-level, sentence-level, and a three-sentence sliding window using regex-based splitting, the latter matching the context-span format in *DefExtra*. For extraction, we compare *OneStep* prompting, which directly extracts definitions from a chunk, and *MultiStep* prompting, which first determines whether a chunk contains a definition and extracts only if it does. We also evaluate a few-shot variant, *OneStep-FS* and *MultiStep-FS*, using training examples. All extractors output a definition text and an explicit/implicit type label. In addition, we implement DSPy-based versions of the one-step and multi-step extractors and optimize prompts with BootstrapFewShot, BootstrapFewShotWithRandomSearch, and MIPROv2 [30]. During DSPy optimization, models also predict the local context span as an auxiliary grounding signal (to mitigate the generation of definitions not present in the context), but context is not used during inference.

We split *DefExtra* into 60%/20%/20% train/development/test sets corresponding to 45/15/15 papers and 133/54/81 definitions, respectively. The training split is used for few-shot samples and DSPy

optimization, the development for prompt selection, and the test for final comparison across 16 LLMs and prompting strategies. We evaluate both open- and proprietary-LLMs.²

SciDef also includes the evaluation protocol for comparing extractor outputs with *DefExtra*. Since exact matching fails for paraphrased, partial, or broader definitions, we first compare candidate pairwise metrics on established semantic similarity and paraphrase benchmarks: STS3k-{all,adv,non}, STS-B, SICK, MSRP, and QQP [6, 9, 10, 26]. We compare embedding cosine similarity, bidirectional NLI, and LLM-as-a-Judge scoring (full LLM-as-a-judge model & prompt list available in our repository’s appendix), select an NLI-based metric from these general benchmarks, and then use *DefSim* as a task-specific validation resource to assess agreement with human judgments on definition pairs. *DefSim* is therefore not used to tune or select the metric but rather to validate it.

For each paper, we evaluate the predicted definition set P against the gold set G . Given the selected pairwise metric M_{def} (i.e., NLI-based metric), we threshold weak matches at $\tau = 0.25$ and combine semantic similarity with explicit/implicit type agreement:

$$S(g, p) = \begin{cases} 0, & M_{\text{def}}(g, p) < \tau, \\ \frac{M_{\text{def}}(g, p) + \mathbb{1}[\text{type}(g) = \text{type}(p)]}{2}, & \text{otherwise.} \end{cases} \quad (1)$$

We then compute a bidirectional best-match score:

$$\text{Score}(G, P) = \frac{1}{2} \left(\frac{1}{|G|} \sum_{g \in G} \max_{p \in P} S(g, p) + \frac{1}{|P|} \sum_{p \in P} \max_{g \in G} S(g, p) \right). \quad (2)$$

The first term measures coverage of human annotations, while the second penalizes over-generation. If P is empty, the final score is set to zero. We also report *GT coverage*, the percentage of gold definitions with at least one prediction above τ .³

4 Resource Validation Results

We report validation results for all the resources. Detailed per-model/strategy results are available in the repository appendix⁴.

Table 2 shows representative sample-label pairs from the two released datasets. For *DefExtra*, the input contains context from a paper, and the label is the extracted term, definition, and definition type. For *DefSim*, the input is a definition pair, and the label is the human similarity judgment and the similarity judged by our metric.

4.1 Metric Selection and DefSim Validation

Before running the extraction benchmark, we selected the pairwise similarity metric using established semantic similarity datasets. Across the seven benchmarks, NLI achieved the strongest performance, with a mean F1 of 0.824, a median F1 of 0.843, and the best score on six of seven benchmarks. Embeddings achieved a mean F1 of 0.682 and a median F1 of 0.719, performing best only on MSRP. LLM-as-a-Judge scoring achieved a mean F1 of 0.687 and a median F1 of 0.713, but did not perform best on any benchmark. We therefore use NLI as the pairwise metric for extraction scoring.

²The full model list, prompts, provider details, and configurations are available in the released repository.

³Full implementation details are available in our repository.

⁴<https://github.com/Media-Bias-Group/SciDef/blob/main/APPENDIX.md>

<i>DefExtra</i> : context → extraction label		<i>DefSim</i> : definition pair → similarity label	
Sample context	Label	Definition pair	Label
"Media bias is defined by researchers as slanted news coverage or internal bias, reflected in news articles."	<i>media bias</i> : "slanted news coverage or internal bias, reflected in news articles"; explicit	<i>publication bias</i> : "certain types of statistical results are more likely to be published ..."	5/5
"Publication bias arises if certain types of statistical results are more likely to be published than other results."	<i>publication bias</i> : "certain types of statistical results are more likely to be published than other results"; explicit	↔ <i>publication bias</i> : "arises if certain types of statistical results are more likely to be published ..."	
"... people often use terms that are highly offensive to certain groups but in a qualitatively different manner."	<i>offensive language</i> : "Terms that are highly offensive to certain groups but are used in a qualitatively different manner"; implicit	<i>task planning</i> : "a complex process that involves reasoning, decision-making, and coordination to ..."	3/5
		↔ <i>task planning</i> : "the process of generating a sequence of actions to achieve a goal"	
		<i>ad hominem bias</i> : "a journalist attacks another person instead of their argument ..."	1/5
		↔ <i>climatic sampling bias</i> : "bias that occurs when trials are performed under unrepresentative seasonal climate conditions"	

Table 2: Sample-label examples from the released datasets. Contexts are shortened for display. *DefExtra* labels contain the extracted term, definition, and explicit/implicit type; *DefSim* labels are human definition-similarity judgments on a 1–5 scale.

Table 3: Extraction validation on *DefExtra* (test set). Mean is averaged across configurations. GT, score, and prediction counts are reported for the best observed configuration in each strategy family.

Strategy	Mean	Best Configuration	GT	Score	Pred. avg.	Pred. med.
DSPy	0.216	Qwen3-30B-A3B-Thinking-2507-FP8, section	69.7%	0.397	12.20	5.00
MultiStep-FS	0.219	Magistral-Small-2509, paragraph	66.7%	0.382	11.07	6.00
MultiStep	0.221	Magistral-Small-2509, paragraph	57.6%	0.341	9.93	5.00
OneStep	0.158	gpt-5-nano, section	77.3%	0.274	27.53	18.00
OneStep-FS	0.147	deepseek-chat-v3-0324, section	86.4%	0.258	39.13	22.00

After selecting NLI, we use *DefSim* to validate whether the metric transfers to scientific definition similarity. For the released definition-pair labels ($N = 60$), annotators show very high agreement, with Krippendorff’s α [17] of $\alpha = 0.924$, and the NLI metric aligns closely with the mean human labels with Pearson correlation of $\rho = 0.937$. As a second validation (i.e., outside of *DefSim*), we evaluate the full paper-level scoring function in Equation (2) on a separate set of 58 full-paper extraction outputs, each scored by human annotators and our automatic metric. Annotator agreement remains substantial with $\alpha = 0.740$, and metric-human agreement is high with $\rho = 0.918$. Together, these results support *DefSim* as a task-specific validation resource for definition matching and support the NLI-based score used in the extraction benchmark.

4.2 Extraction Benchmark with SciDef

We use *SciDef* to validate *DefExtra* as an extraction benchmark and to validate the NLI-based metric. Table 3 reports the average test score by strategy, the best observed configuration in each strategy family, its ground-truth coverage, final score, and average/median number of predicted definitions per paper. The strongest configuration is a DSPy-optimized extractor using a *Qwen-family* model, reaching a test score of 0.397. The highest ground-truth coverage is achieved by a one-step few-shot extractor, which recovers 86.4% of gold definitions but obtains a lower score because it also extracts

many irrelevant candidates. The mismatch between coverage, score, and prediction counts demonstrates the main empirical value of the benchmark: it distinguishes between simply finding many plausible definitions and finding the relevant definitions.

The results show that LLMs can extract definitions from scientific papers, but extraction coverage alone is insufficient. High-recall extractors often overgenerate irrelevant or non-definitional statements, while the best-scoring extractors balance coverage with specificity. Eight of the top 10 extractors are DSPy-based, suggesting prompt optimization can improve this balance. All top-10 rankings and score-vs.-count plots are available in the repository appendix.

5 Conclusion

We released two resources for scientific definition extraction: *DefExtra*, a benchmark of human-extracted definitions from academic papers, and *DefSim*, a human-labeled definition-similarity resource. To support reuse, we also released *SciDef*, a reproducible LLM-based extraction and evaluation pipeline. Our validation experiments show that LLMs can recover many scientific definitions. They also show that DSPy-optimized prompting improves extraction quality and that an NLI-based metric provides a useful basis for semantic matching between predicted and human definitions.

Beyond benchmarking, the resources support practical literature-analysis workflows. Researchers can use *DefExtra* to compare new extractors, *DefSim* to validate definition-matching metrics, and *SciDef* to collect candidate definitions for downstream tasks such as literature reviews, taxonomy construction, ontology development, and domain mapping. The practical lesson is that current LLMs are useful as high-recall definition discovery tools but should not be treated as fully automatic definition selectors: models often retrieve many plausible candidates, making relevance-aware filtering and, in high-stakes settings, human validation a target for future work.

Acknowledgments

This work was partially supported by JSPS KAKENHI JP24H00732, by JST CREST Grants JPMJCR20D3 and JPMJCR2562, including the AIP challenge program, and by JST K Program Grant JPMJKP24C2, Japan. It was also funded by the German Federal Ministry of Education and Research (BMBF) through the DAAD (German Academic Exchange Service) and by the Humboldt Foundation.

GenAI Usage Disclosure

Generative AI systems were used in this work in four ways. First, LLMs are the subject and components of the SciDef experiments: we evaluate LLM-based extraction pipelines across models, prompting strategies, and chunking schemes. Second, in the prior TaxoMatic stage that DefExtra builds on, GPT-3.5-turbo was used for search-space expansion as described in Section 3.1. Third, generative AI tools were used for language editing and drafting assistance during manuscript preparation. Fourth, generative AI tools were used for documentation drafting. All AI-assisted text, code, documentation, analyses, and claims were reviewed, corrected, and approved by the authors. No LLM-generated definition was accepted as a DefExtra gold annotation without human validation. DefSim labels were assigned by human annotators, and annotators were not instructed to defer to model or metric outputs. The released code and documentation identify AI-assisted components where applicable.

References

- [1] Eneko Agirre, Daniel Cer, Mona Diab, and Aitor Gonzalez-Agirre. 2012. SemEval-2012 Task 6: A Pilot on Semantic Textual Similarity. In **SEM 2012: The First Joint Conference on Lexical and Computational Semantics – Volume 1: Proceedings of the main conference and the shared task, and Volume 2: Proceedings of the Sixth International Workshop on Semantic Evaluation (SemEval 2012)*. Eneko Agirre, Johan Bos, Mona Diab, Suresh Manandhar, Yuval Marton, and Deniz Yuret (Eds.). Association for Computational Linguistics, Montréal, Canada, 385–393. <https://aclanthology.org/S12-1051>
- [2] Shubhanker Banerjee, Bharathi Raja Chakravarthi, and John Philip McCrae. 2024. Large Language Models for Few-Shot Automatic Term Extraction. In *Natural Language Processing and Information Systems*. Amon Rapp, Luigi Di Caro, Farid Meziane, and Vijayan Sugumaran (Eds.). Springer Nature Switzerland, Cham, 137–150.
- [3] Samuel R. Bowman, Gabor Angeli, Christopher Potts, and Christopher D. Manning. 2015. A large annotated corpus for learning natural language inference. In *Proceedings of the 2015 Conference on Empirical Methods in Natural Language Processing*, Lluís Màrquez, Chris Callison-Burch, and Jian Su (Eds.). Association for Computational Linguistics, Lisbon, Portugal, 632–642. doi:10.18653/v1/D15-1075
- [4] Tom B. Brown, Benjamin Mann, Nick Ryder, Melanie Subbiah, Jared Kaplan, Prafulla Dhariwal, Arvind Neelakantan, Pranav Shyam, Girish Sastry, Amanda Askell, Sandhini Agarwal, Ariel Herbert-Voss, Gretchen Krueger, Tom Henighan, Rewon Child, Aditya Ramesh, Daniel M. Ziegler, Jeffrey Wu, Clemens Winter, Christopher Hesse, Mark Chen, Eric Sigler, Mateusz Litwin, Scott Gray, Benjamin Chess, Jack Clark, Christopher Berner, Sam McCandlish, Alec Radford, Ilya Sutskever, and Dario Amodei. 2020. Language Models are Few-Shot Learners. In *Advances in Neural Information Processing Systems 33: Annual Conference on Neural Information Processing Systems 2020, NeurIPS 2020, December 6–12, 2020, virtual*. Hugo Larochelle, Marc'Aurelio Ranzato, Raia Hadsell, Maria-Florina Balcan, and Hsuan-Tien Lin (Eds.). <https://proceedings.neurips.cc/paper/2020/hash/1457c0d6bfc4967418bfb8ac142f64a-Abstract.html>
- [5] Xi Chen, Ali Zeynali, Chico Camargo, Fabian Flöck, Devin Gaffney, Przemysław Grabowicz, Scott Hale, David Jurgens, and Mattia Samory. 2022. SemEval-2022 Task 8: Multilingual news article similarity. In *Proceedings of the 16th International Workshop on Semantic Evaluation (SemEval-2022)*. Guy Emerson, Natalie Schluter, Gabriel Stanovsky, Ritesh Kumar, Alexis Palmer, Nathan Schneider, Sidharth Singh, and Shyam Ratan (Eds.). Association for Computational Linguistics, Seattle, United States, 1094–1106. doi:10.18653/v1/2022.semeval-1.155
- [6] DataCanary, hilfalkaff, Jiang Lili, Risdal Meg, Dandekar Nikhil, and tmtung. 2017. Quora Question Pairs. <https://kaggle.com/competitions/quora-question-pairs>
- [7] Claudio Delli Bovi, Luca Telesca, and Roberto Navigli. 2015. Large-Scale Information Extraction from Textual Definitions through Deep Syntactic and Semantic Analysis. *Transactions of the Association for Computational Linguistics* 3 (2015), 529–543. doi:10.1162/tacl_a_00156
- [8] Patrick Diaba-Nuhoho and Michael Amponsah-Offeh. 2021. Reproducibility and research integrity: the role of scientists and institutions. *BMC Research Notes* 14 (2021), 451. doi:10.1186/s13104-021-05875-3
- [9] William B. Dolan and Chris Brockett. 2005. Automatically Constructing a Corpus of Sentential Paraphrases. In *Proceedings of the Third International Workshop on Paraphrasing (IWP2005)*. <https://aclanthology.org/105-5002>
- [10] James Fodor, Simon De Deyne, and Shinsuke Suzuki. 2025. Compositionality and Sentence Meaning: Comparing Semantic Parsing and Transformers on a Challenging Sentence Similarity Dataset. *Computational Linguistics* 51, 1 (March 2025), 139–190. doi:10.1162/coli_a_00536
- [11] Tianyu Gao, Xingcheng Yao, and Danqi Chen. 2021. SimCSE: Simple Contrastive Learning of Sentence Embeddings. In *Proceedings of the 2021 Conference on Empirical Methods in Natural Language Processing*, Marie-Francine Moens, Xuanjing Huang, Lucia Specia, and Scott Wen-tau Yih (Eds.). Association for Computational Linguistics, Online and Punta Cana, Dominican Republic, 6894–6910. doi:10.18653/v1/2021.emnlp-main.552
- [12] Felix Hamborg, Karsten Donnay, and Bela Gipp. 2019. Automated identification of media bias in news articles: an interdisciplinary literature review. *International Journal on Digital Libraries* 20, 4 (2019), 391–415. doi:10.1007/s00799-018-0261-y
- [13] Minhao Jiang, Xiangchen Song, Jieyu Zhang, and Jiawei Han. 2022. TaxoEnrich: Self-Supervised Taxonomy Completion via Structure-Semantic Representations. In *WWW '22: The ACM Web Conference 2022, Virtual Event, Lyon, France, April 25–29, 2022*, Frédérique Laforest, Raphaël Troncy, Elena Simperl, Deepak Agarwal, Aristides Gionis, Ivan Herman, and Lionel Médini (Eds.). ACM, 925–934. doi:10.1145/3485447.3511935
- [14] Yiping Jin, Min-Yen Kan, Jun-Ping Ng, and Xiangnan He. 2013. Mining Scientific Terms and their Definitions: A Study of the ACL Anthology. In *Proceedings of the 2013 Conference on Empirical Methods in Natural Language Processing*, David Yarowsky, Timothy Baldwin, Anna Korhonen, Karen Livescu, and Steven Bethard (Eds.). Association for Computational Linguistics, Seattle, Washington, USA, 780–790. <https://aclanthology.org/D13-1073>
- [15] Brett J. Kagan, Michael Mahlis, Anjali Bhat, Josh Bongard, Victor M. Cole, Phillip Corlett, Christopher Gyngell, Thomas Hartung, Bianca Jupp, Michael Levin, Tamra Lysaght, Nicholas Opie, Adeel Razi, Lena Smirnova, Ian Tennant, Peter Theustrup Wade, and Ge Wang. 2024. Toward a nomenclature consensus for diverse intelligent systems: Call for collaboration. *The Innovation* 5, 5 (09 Sep 2024). doi:10.1016/j.xinn.2024.100658
- [16] Omar Khattab, Arnab Singhvi, Paridhi Maheshwari, Zhiyuan Zhang, Keshav Santhanam, Sri Vardhamanan, Saiful Haq, Ashutosh Sharma, Thomas T. Joshi, Hanna Moazam, Heather Miller, Matei Zaharia, and Christopher Potts. 2023. DSPy: Compiling Declarative Language Model Calls into Self-Improving Pipelines. <https://arxiv.org/abs/2310.03714>
- [17] Klaus Krippendorff. 2011. Computing Krippendorff’s alpha-reliability. (2011).
- [18] Klaus Krippendorff. 2019. *Content Analysis: An Introduction to Its Methodology* (4 ed.). SAGE Publications, Inc., Thousand Oaks, CA, USA. doi:10.4135/9781071878781
- [19] Filip Kučera, Christoph Mandl, Isao Echizen, Radu Timofte, and Timo Spinde. 2026. *DefExtra*. doi:10.5281/zenodo.20302624
- [20] Filip Kučera, Christoph Mandl, Isao Echizen, Radu Timofte, and Timo Spinde. 2026. *DefSim*. doi:10.5281/zenodo.20302704
- [21] Filip Kučera, Christoph Mandl, Isao Echizen, Radu Timofte, and Timo Spinde. 2026. *SciDef*. doi:10.5281/zenodo.18501198
- [22] Dongha Lee, Jiaming Shen, Seongku Kang, Susik Yoon, Jiawei Han, and Hwanjo Yu. 2022. TaxoCom: Topic Taxonomy Completion with Hierarchical Discovery of Novel Topic Clusters. In *WWW '22: The ACM Web Conference 2022, Virtual Event, Lyon, France, April 25–29, 2022*, Frédérique Laforest, Raphaël Troncy, Elena Simperl, Deepak Agarwal, Aristides Gionis, Ivan Herman, and Lionel Médini (Eds.). ACM, 2819–2829. doi:10.1145/3485447.3512002
- [23] Heinrich R. Liesefeld, Dominique Lamy, Nicholas Gaspelin, Joy J. Geng, Dirk Kerzel, Jeffrey D. Schall, Harriet A. Allen, Brian A. Anderson, Sage Boettcher, Niko A. Busch, Nancy B. Carlisle, Hans Colonius, Dejan Draschkow, Howard Eggett, Andrew B. Leber, Hermann J. Müller, Jan Philipp Röer, Anna Schubö, Heleen A. Slagter, Jan Theeuwes, and Jeremy Wolfe. 2024. Terms of debate: Consensus definitions to guide the scientific discourse on visual distraction. *Attention, Perception, & Psychophysics* 86, 5 (01 Jul 2024), 1445–1472. doi:10.3758/s13414-023-02820-3
- [24] Patrice Lopez. 2008–2025. GROBID. <https://github.com/kermitt2/grobid.swh:1:dir:dab86b296e3c3216e2241968f0d63b6e8209d3c>
- [25] Alexander Maedche and Steffen Staab. 2001. Ontology Learning for the Semantic Web. *IEEE Intelligent Systems* 16, 2 (2001), 72–79. doi:10.1109/5254.920602
- [26] Marco Marelli, Stefano Menini, Marco Baroni, Luisa Bentivogli, Raffaella Bernardi, and Roberto Zamparelli. 2014. A SICK cure for the evaluation of compositional distributional semantic models. In *Proceedings of the Ninth International Conference on Language Resources and Evaluation (LREC'14)*, Nicoletta Calzolari, Khalid Choukri, Thierry Declerck, Hrafn Loftsson, Bente Maegaard, Joseph Mariani, Asuncion Moreno, Jan Odijk, and Stelios Piperidis (Eds.). European Language Resources Association (ELRA), Reykjavik, Iceland, 216–223. http://www.lrec-conf.org/proceedings/lrec2014/pdf/363_Paper.pdf
- [27] Jeffrey Milyo and Tim Groseclose. 2005. A Measure of Media Bias. *The Quarterly Journal of Economics* 120 (2005), 1191–1237. doi:10.1162/003355505775097542
- [28] Roberto Navigli and Paola Velardi. 2010. Learning Word-Class Lattices for Definition and Hypernym Extraction. In *Proceedings of the 48th Annual Meeting of the Association for Computational Linguistics*, Jan Hajič, Sandra Carberry, Stephen Clark, and Joakim Nivre (Eds.). Association for Computational Linguistics, Uppsala, Sweden, 1318–1327. <https://aclanthology.org/P10-1134/>
- [29] Roberto Navigli, Paola Velardi, and Juana Maria Ruiz-Martinez. 2010. An Annotated Dataset for Extracting Definitions and Hypernyms from the Web.

- 581 In *Proceedings of the Seventh International Conference on Language Resources*
582 *and Evaluation (LREC'10)*, Nicoletta Calzolari, Khalid Choukri, Bente Mae-
583 gaard, Joseph Mariani, Jan Odiijk, Stelios Piperidis, Mike Rosner, and Daniel
584 Tapias (Eds.). European Language Resources Association (ELRA), Valletta, Malta.
585 http://www.lrec-conf.org/proceedings/lrec2010/pdf/20_Paper.pdf
- [30] Krista Opsahl-Ong, Michael J Ryan, Josh Purcell, David Broman, Christopher
586 Potts, Matei Zaharia, and Omar Khattab. 2024. Optimizing Instructions and
587 Demonstrations for Multi-Stage Language Model Programs. In *Proceedings of*
588 *the 2024 Conference on Empirical Methods in Natural Language Processing*, Yaser
589 Al-Onaizan, Mohit Bansal, and Yun-Nung Chen (Eds.). Association for Computa-
590 tional Linguistics, Miami, Florida, USA, 9340–9366. doi:10.18653/v1/2024.emnlp-
591 main.525
- [31] Oxford University Press. 2025. Oxford English Dictionary Online. <https://www.oed.com/>. Accessed: 2025-09-30.
- [32] Nils Reimers and Iryna Gurevych. 2019. Sentence-BERT: Sentence Embeddings
592 using Siamese BERT-Networks. In *Proceedings of the 2019 Conference on Empirical*
593 *Methods in Natural Language Processing and the 9th International Joint Confer-*
594 *ence on Natural Language Processing (EMNLP-IJCNLP)*, Kentaro Inui, Jing Jiang,
595 Vincent Ng, and Xiaojun Wan (Eds.). Association for Computational Linguistics,
596 Hong Kong, China, 3982–3992. doi:10.18653/v1/D19-1410
- [33] Cezar Sas and Andrea Capiluppi. 2024. Automatic Bottom-Up Taxonomy Con-
597 struction: A Software Application Domain Study. <https://arxiv.org/abs/2409.15881>
- [34] Sasha Spala, Nicholas A. Miller, Yiming Yang, Franck Dernoncourt, and Carl
598 Dockhorn. 2019. DEFT: A corpus for definition extraction in free- and semi-
599 structured text. In *Proceedings of the 13th Linguistic Annotation Workshop*. An-
600 nemarie Friedrich, Deniz Zeyrek, and Jet Hoek (Eds.). Association for Computa-
601 tional Linguistics, Florence, Italy, 124–131. doi:10.18653/v1/W19-4015
- [35] Timo Spinde. 2025. *Automated Detection of Media Bias: From the Conceptualization*
602 *of Media Bias to its Computational Classification*. Springer Vieweg Wiesbaden.
603 doi:10.1007/978-3-658-47798-1
- [36] Timo Spinde, Smi Hinterreiter, Fabian Haak, Terry Ruas, Helge Giese, Norman
604 Meuschke, and Bela Gipp. 2023. The Media Bias Taxonomy: A Systematic
605 Literature Review on the Forms and Automated Detection of Media Bias. <https://arxiv.org/abs/2312.16148>
- [37] Timo Spinde, Luyang Lin, Smi Hinterreiter, and Isao Echizen. 2025. Leveraging
606 Large Language Models for Automated Definition Extraction with TaxoMatic –
607 A Case Study on Media Bias. In *Proceedings of the International AAAI Conference*
608 *on Web and Social Media (ICWSM'25)*, Vol. 19. AAAI, Copenhagen, Denmark.
609 <https://media-bias-research.org/wp-content/uploads/2025/04/spinde2025.pdf>
- [38] Timo Spinde, Manuel Plank, Jan-David Krieger, Terry Ruas, Bela Gipp, and Akiko
610 Aizawa. 2021. Neural Media Bias Detection Using Distant Supervision With
611 BABE - Bias Annotations By Experts. In *Findings of the Association for Computa-*
612 *tional Linguistics: EMNLP 2021*, Marie-Francine Moens, Xuanjing Huang, Lucia
613 Specia, and Scott Wen-tau Yih (Eds.). Association for Computational Linguistics,
614 Punta Cana, Dominican Republic, 1166–1177. doi:10.18653/v1/2021.findings-
615 emnlp.101
- [39] Timo Spinde, Lada Rudnitskaia, Jelena Mitrović, Felix Hamborg, Michael Gran-
616 itzer, Bela Gipp, and Karsten Donnay. 2021. Automated identification of bias
617 inducing words in news articles using linguistic and context-oriented features.
618 *Information Processing & Management* 58, 3 (2021), 102505. doi:10.1016/j.ipm.
619 2021.102505
- [40] Yutian Sun and Hai Zhuge. 2023. Discovering Patterns of Definitions and Methods
620 from Scientific Documents. <https://arxiv.org/abs/2307.01216>
- [41] Maxim Tkachenko, Mikhail Malyuk, Andrey Holmanyuk, and Nikolai Li-
621 ubimov. 2020-2025. Label Studio: Data labeling software. [https://](https://github.com/HumanSignal/label-studio)
622 github.com/HumanSignal/label-studio Open source software available from
623 <https://github.com/HumanSignal/label-studio>.
- [42] Urban Dictionary. 2025. Urban Dictionary. <https://www.urbandictionary.com/>.
624 Accessed: 2025-09-30.
- [43] Shubham Vatsal and Harsh Dubey. 2024. A Survey of Prompt Engineering
625 Methods in Large Language Models for Different NLP Tasks. [https://arxiv.org/](https://arxiv.org/abs/2407.12994)
626 [abs/2407.12994](https://arxiv.org/abs/2407.12994)
- [44] Amir Veysseh, Franck Dernoncourt, Dejing Dou, and Thien Nguyen. 2020. A
627 Joint Model for Definition Extraction with Syntactic Connection and Semantic
628 Consistency. *Proceedings of the AAAI Conference on Artificial Intelligence* 34, 05
629 (Apr. 2020), 9098–9105. doi:10.1609/aaai.v34i05.6444
- [45] Mark D. Wilkinson, Michel Dumontier, IJsbrand Jan Aalbersberg, Gabrielle Apple-
630 ton, Myles Axton, Arie Baak, Niklas Blomberg, Jan-Willem Boiten, Luiz Bonino
631 da Silva Santos, Philip E. Bourne, Jildau Bouwman, Anthony J. Brookes, Tim
632 Clark, Mercè Crosas, Ingrid Dillo, Olivier Dumon, Scott Edmunds, Chris T. Evelo,
633 Richard Finkers, Alejandra Gonzalez-Beltran, Alasdair J.G. Gray, Paul Groth,
634 Carole Goble, Jeffrey S. Grethe, Jaap Heringa, Peter A.C. 't Hoen, Rob Hooft,
635 Tobias Kuhn, Ruben Kok, Joost Kok, Scott J. Lusher, Maryann E. Martone, Al-
636 bert Mons, Abel L. Packer, Bengt Persson, Philippe Rocca-Serra, Marco Roos,
637 Rene van Schaik, Susanna-Assunta Sansone, Erik Schultes, Thierry Sengstag,
638 Ted Slater, George Strawn, Morris A. Swertz, Mark Thompson, Johan van der
639 Lei, Erik van Mulligen, Jan Velterop, Andra Waagmeester, Peter Wittenburg,
640 Katherine Wolstencroft, Jun Zhao, and Barend Mons. 2016. The FAIR Guiding
641 Principles for scientific data management and stewardship. *Scientific Data* 3, 1
642 (15 Mar 2016), 160018. doi:10.1038/sdata.2016.18
- [46] Adina Williams, Nikita Nangia, and Samuel Bowman. 2018. A Broad-Coverage
643 Challenge Corpus for Sentence Understanding through Inference. In *Proceedings*
644 *of the 2018 Conference of the North American Chapter of the Association for*
645 *Computational Linguistics: Human Language Technologies, Volume 1 (Long Papers)*,
646 Marilyn Walker, Heng Ji, and Amanda Stent (Eds.). Association for Computational
647 Linguistics, New Orleans, Louisiana, 1112–1122. doi:10.18653/v1/N18-1101
- [47] Kang Xu, Yifan Feng, Qiandi Li, Zhenjiang Dong, and Jianxiang Wei. 2025. Survey
648 on Terminology Extraction from Texts. *Journal of Big Data* 12, 1 (Feb. 2025), 29.
649 doi:10.1186/s40537-025-01077-x