
CSMED: Bridging the Dataset Gap in Automated Citation Screening for Systematic Literature Reviews

Wojciech Kusa^{1*}, Oscar E. Mendoza², Matthias Samwald³, Petr Knoth⁴, Allan Hanbury¹
¹TU Wien ²University Milano-Bicocca ³Medical University of Vienna ⁴The Open University

Abstract

Systematic literature reviews (SLRs) play an essential role in summarising, synthesising and validating scientific evidence. In recent years, there has been a growing interest in using machine learning techniques to automate the identification of relevant studies for SLRs. However, the lack of standardised evaluation datasets makes comparing the performance of such automated literature screening systems difficult. In this paper, we analyse the citation screening evaluation datasets, revealing that many of the available datasets are either too small, suffer from data leakage or have limited applicability to systems treating automated literature screening as a classification task, as opposed to, for example, a retrieval or question-answering task. To address these challenges, we introduce CSMED, a meta-dataset consolidating nine publicly released collections, providing unified access to 325 SLRs from the fields of medicine and computer science. CSMED serves as a comprehensive resource for training and evaluating the performance of automated citation screening models. Additionally, we introduce CSMED-FT, a new dataset designed explicitly for evaluating the full text publication screening task. To demonstrate the utility of CSMED, we conduct experiments and establish baselines on new datasets.

1 Introduction

Systematic literature reviews (SLRs, or meta-reviews) are a critical tool in scientific research, used for synthesising and summarising evidence from multiple studies. The SLR process involves several stages, including *citation screening (CS, or selection of primary studies)* which is, in itself, a time-consuming step [58, 10]. CS involves identifying studies relevant to the SLR based on a set of, often complex, inclusion and exclusion criteria (e.g., the study must be examining the efficacy of Drug X on Condition Y).

In recent years, there has been an increasing interest in automating the SLR process [80, 60, 56, 22, 2], with works often focusing on improving the CS step by (a) using machine learning (ML) [44, 45], (b) natural language processing (NLP) [27, 36, 79], and (c) information retrieval (IR) [70, 90] techniques. Automated CS systems have the potential to significantly reduce the time and resources required for this critical step, thereby speeding up the SLRs production [75].

The development of standards provides invaluable resources for evaluating and comparing different models. Benchmarks, such as BEIR [74], GLUE [85] or BLURB [20] have shown improvements in reproducibility and progress tracking of machine learning models in various domains. Unfortunately, in the context of SLR automation, the absence of standard benchmarks and evaluation methodologies still hampers progress and inhibits the development of reliable and effective solutions.

With the fast-evolving landscape of machine learning, identifying state-of-the-art performance has become especially challenging and inefficient in the context of CS. The notorious proliferation of small custom CS datasets and single-usage evaluation approaches further exacerbates this issue.

*Corresponding author: wojciech.kusa@tuwien.ac.at

We show that current CS datasets exhibit several shortcomings that hinder their applicability for comprehensive and standardised evaluations. These datasets are poorly documented, with most lacking datasheets, clear licenses and terms of use. In addition, the limited applicability of older datasets arises from their small size and lack of crucial metadata, restricting their use to classification tasks. Finally, data leakage and dataset overlap is another issue, with some SLRs present in multiple collections.

To address these limitations, we present **CSMED** (**C**itation **S**creening **M**eta-**D**ataset), a comprehensive collection of CS datasets that can be used to benchmark and evaluate automated screening systems. Our collection builds upon nine existing datasets, and a new dataset for evaluating the full text classification task, counting 325 SLRs from the fields of medicine and computer science. Thanks to the data harmonisation, our new collection can mitigate the issues of lack of canonical splits, limited applicability, and dataset overlap. Our contributions are as follows:

1. We create CSMED, a meta-collection of nine datasets comprising 325 SLRs. CSMED is built upon BIGBIO [16] and can be used to evaluate and benchmark automated CS systems. We also provide a comprehensive summary of existing citation screening datasets.
2. We extend CSMED with additional metadata after analysing issues on the existing collections and previous evaluation frameworks. Our extended dataset can be used to evaluate CS as question answering or textual pairs classification tasks.
3. Using new metadata, we introduce CSMED-FT, a new dataset for the task of full text screening. To the best of our knowledge, this is the first dataset designed explicitly on screening long documents in SLR. This dataset can be used for the evaluation of the inference capabilities based on a very long context (4,000+ words).

The remainder of the paper is structured as follows. In Section 2, we define the task of citation screening for systematic literature reviews. Section 3 provides an overview of related work in SLR automation and available benchmarks. Section 4 describes the CSMED meta-dataset in detail, including its creation, analysis and extension. In Section 5, we introduce the full text screening dataset together with baseline results on this dataset, and in Section 6, we discuss the implications of our work and potential extensions.

2 Task formulation

We start by introducing the task of citation screening for SLRs and presenting the notation used for its formulation. An SLR is characterised by various attributes, including the title, abstract, research question \mathcal{RQ} , and eligibility criteria \mathcal{C} . We refer to all these attributes as the SLR protocol. Eligibility criteria comprise a set of rules and conditions that a document must meet for inclusion in the SLR. Given a large pool of documents denoted as \mathcal{D} , the main goal of automated citation screening is to assist researchers in identifying relevant publications for inclusion in an SLR. Each document $d \in \mathcal{D}$ has attributes such as its title, abstract, main content, authors, and publication year. The task of CS for SLRs can be formally defined as follows:

Definition 2.1 (CS). Given a set of documents \mathcal{D} and a set of eligibility criteria \mathcal{C} , the task of CS for SLR is to determine for each document $d \in \mathcal{D}$ whether it satisfies the criteria \mathcal{C} . This decision can be represented as a binary label $y_d \in \{0, 1\}$, where $y_d = 1$ if document d satisfies the criteria \mathcal{C} , and $y_d = 0$ otherwise.

It is important to note that the manual CS is conducted in two steps, as shown in Figure 1: title and abstract screening and full text screening. In the first step, the relevance of each document is evaluated based on its title and abstract, while in the second step, a more thorough assessment is performed by examining the full text of the document.

Document retrieval The initial step involves document retrieval, which aims to generate a set of potentially relevant documents $\mathcal{D}' \subseteq \mathcal{D}$ given \mathcal{RQ} . This step commonly involves querying bibliographic databases with specific keywords and Boolean expressions. We can formulate this step as a retrieval function r , such that $r(\mathcal{RQ}, \mathcal{C}) = \mathcal{D}'$. However, the retrieved set \mathcal{D}' may contain a large number of false positives (irrelevant documents).

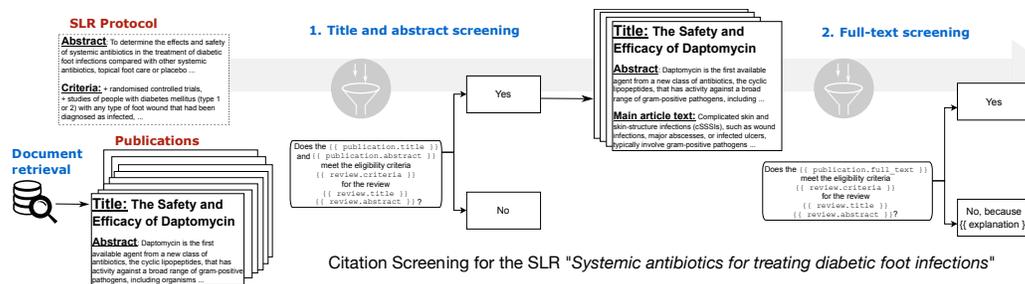


Figure 1: Illustration of the citation screening process, separated into two tasks (1) title and abstract screening and (2) full text screening. Tasks are represented as a specific example of question-answering when a single question asks for a fulfillment of all eligibility criteria \mathcal{C} at once.

Binary classification for relevance prediction Following document retrieval, the primary task is to assess the relevance of each document in the set \mathcal{D}' concerning the eligibility criteria \mathcal{C} . This is conducted in two stages, differing in which attributes of documents are considered (titles and abstracts vs. full texts). We treat this as a binary classification problem, where each document $d \in \mathcal{D}'$ is assigned a binary label $y_d \in \{0, 1\}$ to indicate its relevance ($y_d = 1$) or irrelevance ($y_d = 0$) to the SLR per the criteria \mathcal{C} .

Question answering for relevance An alternative formulation of the citation screening task is to frame it as a question-answering problem. In this approach, we transform the eligibility criteria \mathcal{C} into a set of questions $\mathcal{Q} = \{q_1, \dots, q_{|\mathcal{C}|}\}$, where each question q_k corresponds to a specific criterion in \mathcal{C} . For each document $d \in \mathcal{D}'$, we obtain a set of predicted answers $\hat{A}^d = \{\hat{a}_k^d | \text{meets}(q_k, \hat{a}_k^d)\}$, where $\text{meets}(q_k, \hat{a}_k^d)$ denotes that the document d should meet the criterion expressed by the question q_k . The final relevance label \hat{y}_d of a document d can be determined by aggregating the predicted answers \hat{A}^d using a logical combination function, such as the logical AND operation.

This question-answering formulation offers a more fine-grained assessment of a document's relevance concerning various aspects of the eligibility criteria \mathcal{C} . Other similar formulations of the CS task include document ranking or natural language inference (NLI).

3 Related work

We first motivate the work by providing context on the importance of SLRs and then focus on reviewing citation screening automation methods. Finally, we outline limitations of existing CS datasets.

3.1 Systematic literature reviews

SLRs are particularly important in the medical domain [29]. The Cochrane Collaboration,² the largest organisation responsible for creating SLRs in medicine, has created the foundations of Evidence-Based Medicine [24]. There are more than 220,000 records published between 2000 and 2022 tagged as SLRs in PubMed³ meaning that, on average, there were 10,000 SLRs published per year.

As SLRs focus on reproducibility and finding all relevant evidence about a given topic, the traditional framework involves tasks mainly done manually. It includes steps like defining the search strategy (designing complex Boolean queries) or the screening of every document by at least two reviewers, resulting in an average production time of more than one year [77].

Previous research focused on evaluating automation capabilities for several steps of the traditional framework, such as citation screening (CS) [11, 27, 75], search query (re-)formulation [70, 72], data extraction [59], SLR summarisation [86] or generation of reviews based on the title [94].

²<https://www.cochrane.org>

³<https://pubmed.ncbi.nlm.nih.gov/>

3.2 Citation screening automation

As described in Section 2, CS can be seen as a binary classification problem. However, due to a large number of retrieved studies, the significant class imbalance, and the need to identify *nearly all* relevant documents, this task is inherently complex. Screening automation is a general term for various approaches aimed at reducing workload during the CS stage [60]. These approaches can be classified as either screening reduction, which involves using classification or ranking algorithms to automatically exclude non-relevant publications or screening prioritisation, which focuses on ranking relevant records earlier in the screening process [56]. Automated screening systems leverage techniques from NLP, ML, IR, and statistics, all with the common objective of reducing manual screening time. The disparity in strategies from different fields hinders direct comparison and benchmarking. Next, we discuss some points of disagreement.

NLP approaches typically focus on the level of individual SLRs, treating each review as an independent dataset; whereas IR approaches would consider a set of reviews as a collection, the topics of the reviews analogous to queries, and report aggregated evaluation. Moreover, different publications across various venues adopt diverse evaluation measures, making even more complex the assessment of similar, if not identical, tasks. Evaluation of automatic approaches traditionally relies on binary relevance ratings, very often obtained from the title and abstract screening [60, 32]. When the screening problem is treated as a ranking task, such as screening prioritisation or stopping prediction; the performance is measured in terms of rank-based metrics and metrics at a fixed cut-off, such as $nDCG@n$, $Precision@n$, and *last relevant found* [71, 28]. On the other hand, when the screening problem is treated as a classification task, the performance in this case is measured based on the confusion matrix and the notions of Precision and Recall are commonly used [41, 60]. One challenge arising from these two distinct approaches is the difficulty in going beyond simple effectiveness measures and comparing the real-world savings for users. Further details on datasets and evaluation approaches can be found in a comprehensive review in the Appendix A.

3.3 Limitations of existing datasets

Through our review (see Appendix A), we identified twelve CS datasets reported in former research papers, of which ten have been publicly released. During this analysis, we identified several shortcomings; some are also prevalent in other machine learning problems. Below, we summarise our findings, highlighting the key issues.

Poor documentation One major concern with previous datasets is the lack of sufficient documentation. None of the datasets we examined implement a datasheet [18], which is an essential tool for ensuring transparency and reproducibility. Additionally, seven datasets do not provide clear licenses or terms of use. An inconsistency was also found for one of the datasets [71] in terms of the number of the available content: the paper states 93 SLRs, but we found a list of 176 reviews on the corresponding GitHub repository.

Limited applicability Previous datasets are often small and lack crucial metadata like SLR research question or eligibility criteria, limiting their use to only evaluation of classification tasks. Older datasets typically provide only the title of the review, which limits their applicability for the comprehensive evaluation of neural language understanding models. The most widely used dataset to date [11] was released in 2006. As ML and NLP techniques continue to advance rapidly, it is crucial to have up-to-date datasets that reflect the complexities and nuances of the current research landscape.

Lack of canonical splits Another significant challenge of previous datasets is the absence of canonical train-test splits. Depending on the field of research, practices may vary. As discussed before, in the ML and NLP domains, the prevailing practice is to use inter-review splits, where each review is treated as an individual dataset, and a set of citations is selected for training and testing. Conversely, IR publications often report intra-review splits, treating each review as a “topic” or query, and averaging the results across multiple queries.

In this sense, only the three TAR⁴ datasets contain pre-defined canonical splits, yet only at the intra-review level. For three other datasets [11, 82, 27], previous works have demonstrated significant

⁴TAR stands for Technology-Assisted Reviews and was a shared task organised at CLEF between 2017 and 2019 by Kanoulas et al. [30, 31, 32].

variability in model evaluation based on the selection of cross-validation splits, particularly for the smallest datasets that contain a limited number of relevant documents [79, 38]. The lack of standardised splits, especially in collections with fewer SLRs, makes it challenging to compare different approaches and hinders the fair evaluation of models' performance.

Dataset overlap We also evaluated overlapping throughout the previous datasets and discovered that at least 11 SLRs were present in multiple collections [71, 30, 31, 32]. Additionally, the TAR 2019 dataset contains three SLRs that are present both in its training and test splits, accounting for approximately 6% of the test partition [32]. While this overlap is not a significant concern when evaluating unsupervised methods like BM25 [69], it poses a potential threat to conducting fair comparisons with large language models (LLMs). Machine learning models, and especially LLMs, have the capability to memorise their training data, making it critical to address dataset overlap to ensure unbiased evaluations [25] (see Appendix E for a detailed analysis of the overlapping in previous datasets).

Lack of common evaluation Another notable deficiency among the previous datasets is the absence of a common set of evaluation measures. Only the three TAR datasets provide scripts for evaluating submissions. For example, the most widely used dataset by Cohen et al. [11] was evaluated using several disparate evaluation measures such as *WSS* [11], *AUC* or *Precision@r%*. However, recent research has exposed limitations and problems with both *WSS* and *AUC* as metrics for this task [40].

Availability in biomedical benchmarks Recent efforts have focused on creating larger collections of more diverse datasets to evaluate the performance of biomedical NLP models. These efforts include benchmarks like BLUE [64], HunFlair [92], BLURB [20], and BigBio [16], which provide datasets and tasks for evaluating biomedical language understanding and reasoning. Additionally, there are biomedical datasets geared towards prompt-based learning and evaluation of few and zero-shot classification, such as Super-NaturalInstructions [91] and BoX [62]. Out of all benchmarks mentioned above, only BoX contains one CS dataset covering five SLRs, however, this dataset is private. Coverage for other SLR tasks is also limited.

To summarise, previous datasets exhibit certain drawbacks that limit their suitability for comprehensive and standardised evaluation. While the TAR 2017-19 collections stand out as the only ones containing canonical splits and a set of evaluation measures, some of their topics overlap with another previous dataset [71], and we also identified data leakage in the newest TAR 2019 dataset. Consequently, we believe that developing a new collection is necessary to address these issues and establish a robust foundation for evaluation of CS and SLR automation.

4 The CSMED meta-dataset

The recent advancements and paradigm shifts in NLP and ML; with the extensive use of pre-trained models and transfer learning [46, 15], and the more recent prompt-based learning [49, 9]; have significantly transformed the field and enhanced the predictive capabilities of models across various tasks. Inspired by the success of benchmark collections in the field of biomedical NLP, we conducted a thorough review of available datasets and benchmarks to identify the most representative datasets for the task of citation screening, finding that this task is heavily underrepresented. The available datasets still primarily cater to training supervised algorithms, lacking the scale and granularity necessary to evaluate state-of-the-art models. To address these limitations and provide a more comprehensive resource for training and evaluating data-centric methods in SLR automation, we create CSMED, consolidating nine previously released collections of SLRs. We further extend a subset of SLRs in CSMED with additional metadata coming from the review protocol.

Our data analysis methodology involved creating visualisations and summary tables based on curated datasets. We analyse dataset statistics like available data splits, licensing information, dataset and reviews size as well as dataset overlap. This allows us to provide both a detailed view of individual reviews and an overview of datasets containing multiple reviews (see Appendix B for further details on visualisations).

4.1 Dataset creation

Currently, nine out of ten public CS datasets we identified have been included in CSMED, with the remaining one to be included. We provide a summary of the datasets in Table 1, and further details can be found in Appendix A. In total, CSMED consists of 325 SLRs, making it the largest publicly available collection in this domain and the only one providing access to the datasets via a harmonised API.

Table 1: A list of source citation screening datasets included in the CSMED. The first four datasets contain non-Cochrane SLRs, whereas the other five are based on Cochrane reviews. ‘Avg. ratio of included’ column presents ratio of included publication from the title and abstract screening stage, ‘Avg. size’ refers to averaged across SLRs document count in the dataset. The ‘Additional data’ column describes if the review contains metadata other than coming from the citation list: (A): Search queries, (B): Review protocol containing review title, abstract and search strategy, (C): Review updates consisting of changes to included papers. ‘DTA’ stands for diagnostic test accuracy reviews. ‡ – indicates a discrepancy in the number of reviews in the paper versus the GitHub repository. † – indicates the total count of reviews from all nine datasets before duplicates were removed.

Source	# reviews	Domain	Avg. size	Avg. ratio of included	Additional data	Cochrane reviews
[11]	15	Drug	1,249	7.7%	—	—
[82]	3	Clinical	3,456	7.9%	—	—
[27]	5	Mixed	19,271	4.6%	—	—
[22]	7	Comp. Science	340	11.7%	B	—
[71]	93/176‡	Clinical	1,159	1.2%	A	✓
[30]	50	DTA	5,339	4.4%	B	✓
[31]	30	DTA	7,283	4.7%	B	✓
[32]	49	Mixed	2,659	8.9%	B	✓
[2]	25	Clinical	4,402	0.4%	C	✓
Total	360†		3,471	4.4%		

To ensure interoperability and facilitate the ease of use, we designed data loaders for the datasets in accordance with the BigBio text classification schema [16]. This choice offers several advantages. BigBio has the largest coverage of biomedical datasets and supports access to the datasets via API. Moreover, it is compatible with popular libraries such as Hugging Face’s datasets [47] and the EleutherAI Language Model Evaluation Harness [17], thereby reducing the experimental costs.

Taking advantage of the lists of publications that most of the sources of datasets share as PubMed IDs, we extend the BigBio data loaders to enable the download of PubMed articles. Our harmonised data loaders selectively load the PubMed articles that are a part of each dataset. The single exception is the dataset by Hannousse and Yahiouche [22], which is the only publicly available collection of non-medical SLRs. For this dataset, we extract the content using the SemanticScholar API.⁵ As a result, CSMED serves also as the first resource that gathers SLRs from diverse domains.

4.2 Extending metadata

Our goal at this stage is not to create yet another gold standard dataset for SLRs, but rather improve the quality of current data and provide insights into promising avenues for future research. We begin by presenting the possibilities of extending the subset of Cochrane SLRs to experiment with screening beyond supervised classification.

We then categorise CSMED datasets into two groups: (1) datasets containing Cochrane medical SLRs and (2) datasets comprising other SLRs. This distinction is made because from following the Cochrane protocol, more extensive information on the review is provided. We use the additional data available on reviews websites to extend CSMED. Among the new information, the one we find most valuable is the eligibility criteria, which no longer limits the data to the evaluation of supervised binary classification but opens its application to question-answering or language inference tasks.

⁵<https://www.semanticscholar.org/product/api>

We carefully examine the subset of SLRs produced by Cochrane, aiming to identify potential enhancements and extensions that would help mitigating the existing limitations of previous datasets. Every Cochrane SLR first registers and publishes the protocol containing the review title, abstract, search strategy and the eligibility criteria. This information is all that human experts need to produce the final review, i.e., they first find the relevant studies and then conduct the meta-analysis of their results. As described in Section 2, the screening process can be also modelled as a question-answering, where every publication is compared against the eligibility criteria in order to make the decision about the inclusion,⁶ similar to the clinical decision support task of matching clinical trials to patients [67, 68].

To expand CSMED, we searched the Cochrane Library⁷ for all SLRs from the meta-dataset based on the Cochrane review ID and take their latest open-access version. We extract available information about the review: review title and abstract, eligibility criteria, search strategy and references. Cochrane reports a list of included and excluded publications at the full text screening stage (this can be treated as approximately all included publications during the title and abstract screening stage). Moreover, all excluded publications have a reason for exclusion selected by a reviewer. As the original relevance judgements were limited to publications from the PubMed database, we assign PubMed IDs to these publications. We also define appropriate BigBio data loaders so the task can be seen as question-answering or textual pairs classification task.

Table 2: Details of the CSMED expanded meta-dataset. Column ‘#docs’ refers to the total number of documents included in all SLRS within the dataset, ‘#included’ mentions number of included documents on the title and abstract screening stage and ‘Avg. %included’ the percentage of included publications averaged from all reviews.

Dataset name	#reviews	#docs	#included	Avg. #docs	Avg. % included	Avg. #words in document
CSMED-TRAIN-BASIC	30	128,438	7,958	4,281	9.6%	229
CSMED-TRAIN-COCHRANE	195	372,422	7,589	1,910	21.9%	180
CSMED-DEV-COCHRANE	100	229,376	4,365	2,294	20.8%	201
CSMED-ALL	325	730,236	19,912	2,247	20.5%	195

Details of the new expanded CSMED are provided in Table 2. We were not able to find suitable data for all SLRs, hence the expanded CSMED is smaller than the original meta-dataset. In total, the new expanded dataset consists of 295 unique Cochrane SLRs and 30 non-Cochrane SLRs. The entire set of basic SLRs is designated for training. From the Cochrane reviews, we randomly selected 195 to the training split and the remaining 100 to the development split. We abstain from designating a test split because CSMED aggregates existing datasets. Given the constraints of these datasets, creating a new, unbiased test collection is recommended.

4.3 Baseline experiment

We evaluate five models in a zero-shot setting: two statistical models BM25 and TF-IDF, and three Transformer-based models. Predictions are run on the CSMED-DEV-COCHRANE split, and we test four different input representations using the following sections from the SLR protocol: (1) title, (2) abstract, (3) search strategy and (4) eligibility criteria section. We evaluate models using $TNR@95%$ [40, 41], $nP@95%$ [41], $Last Rel$ [30, 31, 32], $nDCG@10$, MAP , Recall at rank k with k in $\{10, 50, 100\}$ ($R@k$). Detailed experimental setup and expanded results can be found in Appendix F.

Table 3 presents the summary results of best-performing statistical and neural models, while Table 9 in Appendix F contains complete results. The MiniLM neural model consistently outperforms the BM25 variant across nearly all query representations and evaluation measures. For the BM25 model, using the systematic review abstract text as a query representation results in the highest performance in all metrics compared to using the SLR title and criteria sections. Conversely, for

⁶In the current approach, we consider only binary relevance (included versus excluded). However, in practice, more categories can be defined by reviewers (e.g. a study can be assigned as a background publication or meta-analysis).

⁷<https://www.cochranelibrary.com/cdsr/reviews>

Table 3: Summary of results on CSMED-DEV-COCHRANE dataset. **Bold** values indicate best score.

Model	Representation	TNR@95%	nP@95%	Last Rel	NDCG@10	MAP	R@10	R@50	R@100
BM25	Title	0.469	0.142	72.2	0.438	0.388	0.349	0.623	0.704
	Abstract	0.474	0.170	63.6	0.503	0.453	0.379	0.657	0.757
	Search strategy	0.379	0.093	72.1	0.336	0.311	0.268	0.507	0.625
	Criteria	0.430	0.145	67.0	0.452	0.417	0.345	0.629	0.725
MPNet	Title	0.467	0.230	66.6	0.476	0.429	0.376	0.684	0.774
	Abstract	0.516	0.265	63.8	0.556	0.482	0.420	0.692	0.777
	Search strategy	0.429	0.181	68.6	0.400	0.372	0.328	0.614	0.699
	Criteria	0.545	0.216	58.5	0.514	0.488	0.393	0.691	0.784

the MiniLM model, using eligibility criteria leads to superior scores in TNR@95% and three other evaluation measures. This indicates that larger models are beginning to effectively leverage the criteria information. Notably, the top-performing model in the zero-shot setting, MiniLM using SLR eligibility criteria, achieves $TNR@95%$ equal to almost 0.55. This means that, on average, this model can remove more than half of the true negatives when achieving a recall of 95%. Exploring more advanced language models might reveal the potential for further harnessing the criteria information.

5 CSMED-FT: full text classification dataset

In this section, we introduce CSMED-FT full text screening dataset and present baseline experiments.

5.1 Dataset creation

LLM advancements have enabled processing long text snippets [7, 95, 55, 21]. Commercial tools now support inputs of up to 32k [61] or even up to 100k tokens [4]. We propose CSMED-FT, the full text screening dataset to enable research associated with the comprehensive understanding of very long documents, and evaluate such capabilities. We first gather full text versions of publications from CSMED SLRs, and then create the appropriate setting with canonical splits.

We use SemanticScholar and CORE [35, 34] APIs to find URLs to open-access full text documents. This process successfully finds URLs to, on average, 27% of all included and excluded publications from SLRs. After downloading full text PDFs, we use GROBID [1] to parse the content of these documents into an XML format.

We establish canonical splits considering the timestamps, such that the newest reviews belong to the test set. Specifically, we select 31 Cochrane reviews published in the last year (between 01/06/2022 and 31/05/2023) to create a test set, another 60 reviews (mentioned in Nussbaumer-Streit et al. [57]) for the development set, and 176 reviews (listed by Scells et al. [71]) as the training set. Filtering out reviews with no associated available full text publications results in 148/36/29 reviews in train/dev/test splits.

Details of CSMED-FT are presented in Table 4. We also release a subset of 50 randomly selected documents from the test set as CSMED-FT-TEST-SMALL. At the moment of writing this publication, creating a prompt for LLMs with an input of few thousands tokens is feasible albeit costly,⁸ See Appendix D for further details on the creation of CSMED-FT.

Table 4: Details of the CSMED-FT dataset. Column ‘#docs’ refers to the total number of documents included in the dataset and ‘#included’ mentions number of included documents on the full text step. CSMED-FT-TEST-SMALL is a subset of CSMED-FT-TEST.

Dataset name	#reviews	#docs.	#included	% included	Avg. #words in document	Avg. #words in review
CSMED-FT-TRAIN	148	2,053	904	44.0%	4,535	1,493
CSMED-FT-DEV	36	644	202	31.4%	4,419	1,402
CSMED-FT-TEST	29	636	278	43.7%	4,957	2,318
CSMED-FT-TEST-SMALL	16	50	22	44.0%	5,042	2,354

⁸According to the OpenAI model pricing on 22/09/2023, screening 500 full text documents with the GPT-4-32k model would cost more than 400 USD.

CSMED-FT could be a proxy for a very long document natural language inference (NLI) task. Popular NLI datasets (SciTail [33], McTest [66] or DocNLI [93]) contain both hypotheses and premises of an average length considerably shorter than 1,000 words; whereas in CSMED-FT, the premise (review protocol) has an average length of more than 1,000 words, and the hypothesis (publication) contains more than 4,000 words.

5.2 Experiment

We present how CSMED-FT can be used to evaluate LLMs capabilities in making eligibility decisions on very long documents. We run experiments both on fine-tuning of Transformer models and zero-shot prompting of GPT models.

Model selection As the combined input size of systematic review and publication can be big (9,246 mean number of tokens on a training split measured with a GPT-4 tokeniser), we only select models that allow inputs of at least 4k tokens context. We fine-tune the open-domain Longformer and BigBird, and domain-specific models pre-trained on clinical data: Clinical-BigBird and ClinicalLongformer. For zero-shot evaluation, we select GPT-3.5-turbo-0301, GPT-4-8k and GPT-3.5-turbo-16k accessed via OpenAI API. GPT-4-8k and GPT-3.5-turbo-16k are the only models capable of handling more than 4k-input tokens, with context window size of 8k and 16k tokens respectively.

Preprocessing and evaluation For all models, we concatenate the screening protocol with each publication; we truncate the review description text to half of the available context window (2,000 tokens for 4k models, 4,000 tokens for 8k model and 8,000 tokens for 16k model) and complete the input with a publication.

For GPT models, if a whole publication text does not fit the context window, we run multiple predictions with a sliding window and aggregate the results. In the case of GPT-3.5-turbo-16k model, only for 4 out of 50 documents the model was unable to process the full text of combined review and publication inside one prompt.

We fine-tune the Transformer models on CSMED-FT-TRAIN for four epochs and run evaluation on CSMED-FT-DEV. Due to the budget limitation, for the GPT-4-8k model, we run the evaluation only on CSMED-FT-TEST-SMALL (see Appendix G for further details on experimental settings). Finally, we evaluate the models using macro-averaged Precision, Recall and F1-score measures.

Results Results of the full text experiment are summarised in Table 5. On CSMED-FT-TEST-SMALL, GPT-4-8k strongly outperforms other models. However, this difference is not statistically significant. The GPT-3.5-turbo-16k achieves the highest Precision; this improvement can be attributed to the model's expanded context window and the limitations other GPT-based models have with our simple aggregation rules. However, this might also be caused by overfitting towards the positive class, as this model includes almost twice as many publications as other models. On CSMED-FT-TEST set, Clinical-BigBird, significantly outperforms zero-shot GPT-3.5 model and pre-trained models based on the LongFormer architecture.

Interestingly, both BigBird-based models outperform their counterparts using the Longformer architecture. The typical overall tendency to domain-pre-trained models achieving higher scores over their open-domain counterparts is also preserved. We believe that fine-tuning the Transformer models first on larger NLI/QA corpora could help improve the results.

6 Discussion

In this paper, we have addressed the challenge of standardised evaluation in CS automation. By revisiting existing screening datasets, we evaluated their suitability as benchmarks in the context of modern ML methods. Our analysis revealed limitations such as small size and data leakage issues.

To overcome these challenges, we introduced CSMED, a meta-dataset consolidating nine publicly released collections, providing programmatic access to 325 SLRs. CSMED serves as a comprehensive resource for training and evaluating automated citation screening models and can be used for tasks that involve textual pairs classification, question answering and NLI. Additionally, we included a

Table 5: Results of the full text screening experiment averaged over documents. The statistical significance was assessed with a McNemar’s t-test ($p < 0.05$) with Bonferroni correction for multiple testing. *Clinical-BigBird* on the CSMED-FT-TEST split showed statistically significant improvements compared to the *stratified random* baseline, *Longformer*, *Clinical-Longformer*, and *GPT-3.5-turbo-16k*, indicated by †. Stratified baseline is averaged from 100 different random seeds. ‘% incl.’ describes the percentage of documents predicted as relevant by models.

	CSMED-FT-TEST-SMALL				CSMED-FT-TEST			
	% incl.	Precision	Recall	F1-score	% incl.	Precision	Recall	F1-score
ORACLE	44%	—	—	—	43.7%	—	—	—
stratified random	50%	0.497	0.498	0.495	—	0.499	0.499	0.498
‘include all’	100%	0.220	0.500	0.306	100%	0.219	0.500	0.304
Longformer [7]	40%	0.467	0.468	0.466	40.4%	0.398	0.400	0.398
BigBird-roberta-base [95]	42%	0.572	0.571	0.572	45.1%	0.575	0.575	0.575
Clinical-Longformer [48]	36%	0.547	0.544	0.542	35.1%	0.436	0.441	0.435
Clinical-BigBird [48]	36%	0.590	0.584	0.583	32.8%	0.623 †	0.611 †	0.609 †
GPT-3.5-turbo-0301	54%	0.585	0.586	0.580	—	—	—	—
GPT-4-8k-0314	58%	0.674	0.672	0.660	—	—	—	—
GPT-3.5-turbo-16k	80%	0.712	0.638	0.576	75.9%	0.538	0.528	0.475

new dataset within CSMED for evaluating full text publication classification and conducted initial experiments showing that there is a room for improvement in understanding long contexts.

The focus of CSMED on providing unified access over a number of diverse citation screening datasets has many benefits. First, the evaluation code can be re-used, making sure that the evaluation is handled properly. Secondly, integration with the BigBio framework enables quick prototyping of prompts. We also improve the documentation for existing datasets and provide a comprehensive data card for CSMED-FT. Our extended version of CSMED is also deduplicated. Finally, it is a step towards providing a multi-domain SLR dataset and bridging the gap between IR and NLP research in the domain of screening automation, enabling direct comparisons of the methods.

Limitations and future work While we attempted to extract the data protocols as accurately as possible, extraction of data was not possible for all previous reviews. This was primarily due to the changing standards in Cochrane reviews throughout the years. In future work, ideally, direct access to Cochrane metadata would be needed to make sure that all information is covered. Even though the PubMed publications most probably will not change, what can change is the API and scripts necessary to download the data. There exists also the possibility that one of the sources will introduce a restriction on using their data for training and evaluation of machine learning models. We tried to further mitigate this potential issue by selecting open-access SLRs produced by Cochrane. Finally, we acknowledge that using machine assistance for citation selection can raise concerns about research quality, emphasising the vital role of human oversight throughout the process.

Future work will focus on further improving data quality, connecting the output reviews from screening tools like *CRUISE-Screening* [39], adding datasets covering other domains and different SLR tasks and designing a dataset for a prospective evaluation of review automation which could ensure no data leakage [12]. For the prospective dataset, predictions could be made as soon as the protocol is published, and the gold standard data becomes available when the review is eventually published, albeit with the drawback of a potentially long waiting time for review publication.

7 Conclusion

Our paper introduces CSMED, a meta-dataset that addresses the lack of standardisation in SLR automation. By consolidating datasets and providing a unified access point, CSMED facilitates the development and evaluation of automated citation screening and full text classification models. We believe it has the potential to advance the field and lead to more robust automated SLR systems. We envision CSMED as a living, evolving collection, and we invite researchers to contribute to expanding it with SLR datasets from other domains.

Acknowledgments and Disclosure of Funding

This work was supported by the EU Horizon 2020 ITN/ETN on Domain Specific Systems for Information Extraction and Retrieval – DoSSIER (H2020-EU.1.3.1., ID: 860721).

References

- [1] Grobid. <https://github.com/kermitt2/grobid>, 2008–2022.
- [2] Amal Alharbi and Mark Stevenson. A dataset of systematic review updates. *SIGIR 2019 - Proceedings of the 42nd International ACM SIGIR Conference on Research and Development in Information Retrieval*, pages 1257–1260, 7 2019. doi: 10.1145/3331184.3331358.
- [3] Amal Alharbi and Mark Stevenson. Refining Boolean queries to identify relevant studies for systematic review updates. *Journal of the American Medical Informatics Association*, 27(11):1658–1666, 10 2020. ISSN 1527-974X. doi: 10.1093/jamia/ocaa148. URL <https://doi.org/10.1093/jamia/ocaa148>.
- [4] Anthropic. claude-v1.3-100k, Blog post 'Introducing 100K Context Windows'. <https://www.anthropic.com/index/100k-context-windows>, 2023.
- [5] Maisie Badami, Boualem Benatallah, and Marcos Baez. Adaptive search query generation and refinement in systematic literature review. *Information Systems*, page 102231, 2023. ISSN 0306-4379. doi: <https://doi.org/10.1016/j.is.2023.102231>. URL <https://www.sciencedirect.com/science/article/pii/S0306437923000674>.
- [6] Elias Bassani. retriv: A Python Search Engine for the Common Man, May 2023. URL <https://github.com/AmenRa/retriv>.
- [7] Iz Beltagy, Matthew E Peters, and Arman Cohan. Longformer: The long-document transformer. *arXiv preprint arXiv:2004.05150*, 2020.
- [8] Kathrin Blagec, Jakob Kraiger, Wolfgang Frühwirt, and Matthias Samwald. Benchmark datasets driving artificial intelligence development fail to capture the needs of medical professionals. 1 2022. URL <https://arxiv.org/abs/2201.07040v1>.
- [9] Tom Brown, Benjamin Mann, Nick Ryder, Melanie Subbiah, Jared D Kaplan, Prafulla Dhariwal, Arvind Neelakantan, Pranav Shyam, Girish Sastry, Amanda Askell, et al. Language models are few-shot learners. *Advances in neural information processing systems*, 33:1877–1901, 2020.
- [10] Justin Clark, Paul Glasziou, Chris Del Mar, Alexandra Bannach-Brown, Paulina Stehlik, and Anna Mae Scott. A full systematic review was completed in 2 weeks using automation tools: a case study. *Journal of clinical epidemiology*, 121:81–90, 2020.
- [11] A. M. Cohen, W. R. Hersh, K. Peterson, and Po Yin Yen. Reducing workload in systematic review preparation using automated citation classification. *Journal of the American Medical Informatics Association*, 13(2):206–219, 3 2006. ISSN 10675027. doi: 10.1197/jamia.M1929. URL [/pmc/articles/PMC1447545/](https://www.ncbi.nlm.nih.gov/pmc/articles/PMC1447545/). URL <https://www.ncbi.nlm.nih.gov/pmc/articles/PMC1447545/?report=abstract><https://www.ncbi.nlm.nih.gov/pmc/articles/PMC1447545/>.
- [12] Aaron M Cohen, Kyle Ambert, and Marian McDonagh. A prospective evaluation of an automated classification system to support evidence-based medicine and systematic review. In *AMIA annual symposium proceedings*, volume 2010, page 121. American Medical Informatics Association, 2010.
- [13] Siddhartha R Dalal, Paul G Shekelle, Susanne Hempel, Sydne J Newberry, Aneesa Motala, and Kanaka D Shetty. A pilot study using machine learning and domain knowledge to facilitate comparative effectiveness review updating. *Medical Decision Making*, 33(3):343–355, 2013.
- [14] Jay DeYoung, Iz Beltagy, Madeleine van Zuylen, Bailey Kuehl, and Lucy Lu Wang. Ms²: Multi-document summarization of medical studies. In *EMNLP*, 2021.
- [15] Li Dong, Nan Yang, Wenhui Wang, Furu Wei, Xiaodong Liu, Yu Wang, Jianfeng Gao, Ming Zhou, and Hsiao-Wuen Hon. Unified language model pre-training for natural language understanding and generation. *Advances in neural information processing systems*, 32, 2019.

- [16] Jason Fries, Leon Weber, Natasha Seelam, Gabriel Altay, Debajyoti Datta, Samuele Garda, Sunny Kang, Rosaline Su, Wojciech Kusa, Samuel Cahyawijaya, Fabio Barth, Simon Ott, Matthias Samwald, Stephen Bach, Stella Biderman, Mario Sanger, Bo Wang, Alison Callahan, Daniel Le3n Perifi3n, Th3o Gigant, Patrick Haller, Jenny Chim, Jose Posada, John Giorgi, Karthik Rangasai Sivaraman, Marc P3mies, Marianna Nezhurina, Robert Martin, Michael Cullan, Moritz Freidank, Nathan Dahlberg, Shubhanshu Mishra, Shamik Bose, Nicholas Broad, Yanis Labrak, Shlok Deshmukh, Sid Kiblawi, Ayush Singh, Minh Chien Vu, Trishala Neeraj, Jonas Golde, Albert Villanova del Moral, and Benjamin Beilharz. Bigbio: A framework for data-centric biomedical natural language processing. In S. Koyejo, S. Mohamed, A. Agarwal, D. Belgrave, K. Cho, and A. Oh, editors, *Advances in Neural Information Processing Systems*, volume 35, pages 25792–25806. Curran Associates, Inc., 2022. URL https://proceedings.neurips.cc/paper_files/paper/2022/file/a583d2197eafc4afdd41f5b8765555c5-Paper-Datasets_and_Benchmarks.pdf.
- [17] Leo Gao, Jonathan Tow, Stella Biderman, Sid Black, Anthony DiPofi, Charles Foster, Laurence Golding, Jeffrey Hsu, Kyle McDonell, Niklas Muennighoff, Jason Phang, Laria Reynolds, Eric Tang, Anish Thite, Ben Wang, Kevin Wang, and Andy Zou. A framework for few-shot language model evaluation, September 2021. URL <https://doi.org/10.5281/zenodo.5371628>.
- [18] Timit Gebru, Jamie Morgenstern, Briana Vecchione, Jennifer Wortman Vaughan, Hanna M. Wallach, Hal Daum3 III, and Kate Crawford. Datasheets for datasets. *Commun. ACM*, 64(12):86–92, 2021. doi: 10.1145/3458723. URL <https://doi.org/10.1145/3458723>.
- [19] Maura R Grossman, Gordon V Cormack, and Adam Roegiest. TREC 2016 Total Recall Track Overview. In *TREC*, 2016.
- [20] Yu Gu, Robert Tinn, Hao Cheng, Michael Lucas, Naoto Usuyama, Xiaodong Liu, Tristan Naumann, Jianfeng Gao, and Hoifung Poon. Domain-specific language model pretraining for biomedical natural language processing. *ACM Trans. Comput. Heal.*, 3(1):2:1–2:23, 2022. doi: 10.1145/3458754. URL <https://doi.org/10.1145/3458754>.
- [21] Mandy Guo, Joshua Ainslie, David Uthus, Santiago Ontanon, Jianmo Ni, Yun-Hsuan Sung, and Yinfei Yang. Long5: Efficient text-to-text transformer for long sequences. *arXiv preprint arXiv:2112.07916*, 2021.
- [22] Abdelhakim Hannousse and Salima Yahiouche. A semi-automatic document screening system for computer science systematic reviews. In *Mediterranean Conference on Pattern Recognition and Artificial Intelligence*, pages 201–215. Springer, 2022.
- [23] Kazuma Hashimoto, Georgios Kontonatsios, Makoto Miwa, and Sophia Ananiadou. Topic detection using paragraph vectors to support active learning in systematic reviews. *Journal of Biomedical Informatics*, 62: 59–65, 8 2016.
- [24] Julian PT Higgins, James Thomas, Jacqueline Chandler, Miranda Cumpston, Tianjing Li, Matthew J Page, and Vivian A Welch. *Cochrane handbook for systematic reviews of interventions*. John Wiley & Sons, 2019.
- [25] hitz-zentroa. LM Contamination Index. GitHub repository, 2023. URL <https://github.com/hitz-zentroa/lm-contamination>.
- [26] Jingwen Hou, Xiaochen Wang, Jean-Jacques Dubois, R. Byron Rice, Amanda Haddock, and Yue Wang. Extreme systematic reviews: A large literature screening dataset to support environmental policymaking. In *Proceedings of the 31st ACM International Conference on Information & Knowledge Management*, CIKM ’22, page 4029–4033, New York, NY, USA, 2022. Association for Computing Machinery. ISBN 9781450392365. doi: 10.1145/3511808.3557600. URL <https://doi.org/10.1145/3511808.3557600>.
- [27] Brian E. Howard, Jason Phillips, Kyle Miller, Arpit Tandon, Deepak Mav, Mihir R. Shah, Stephanie Holmgren, Katherine E. Pelch, Vickie Walker, Andrew A. Rooney, Malcolm Macleod, Ruchir R. Shah, and Kristina Thayer. SWIFT-Review: A text-mining workbench for systematic review. *Systematic Reviews*, 5(1):1–16, 5 2016. ISSN 20464053. doi: 10.1186/s13643-016-0263-z. URL <https://link.springer.com/articles/10.1186/s13643-016-0263-zhttps://link.springer.com/article/10.1186/s13643-016-0263-z>.
- [28] Brian E. Howard, Jason Phillips, Arpit Tandon, Adyasha Maharana, Rebecca Elmore, Deepak Mav, Alex Sedykh, Kristina Thayer, B. Alex Merrick, Vickie Walker, Andrew Rooney, and Ruchir R. Shah. SWIFT-Active Screener: Accelerated document screening through active learning and integrated recall estimation. *Environment International*, 138:105623, 5 2020. ISSN 0160-4120. doi: 10.1016/J.ENVINT.2020.105623.

- [29] Akers Jo, Aguiar-Ibáñez Raquel, Burch Jane, Chambers Duncan, Eastwood Alison, Fayter Debra, Hempel Susanne, Light Kate, Rice Stephen, Rithalia Amber, Stewart Lesley, Stock Christian, Wilson Paul, and Woolacott Nerys. *Systematic Reviews: CRD's guidance for undertaking reviews in health care*. CRD, University of York, York, 1 2009. ISBN 978-1-900991-19-3. URL www.york.ac.uk/inst/crd.
- [30] Evangelos Kanoulas, Dan Li, Leif Azzopardi, and Rene Spijker. CLEF 2017 technologically assisted reviews in empirical medicine overview. *CEUR Workshop Proceedings*, 1866:1–29, 9 2017. ISSN 1613-0073. URL <https://pureportal.strath.ac.uk/en/publications/clef-2017-technologically-assisted-reviews-in-empirical-medicine->.
- [31] Evangelos Kanoulas, Dan Li, Leif Azzopardi, and Rene Spijker. CLEF 2018 technologically assisted reviews in empirical medicine overview. *CEUR Workshop Proceedings*, 2125, 7 2018. ISSN 1613-0073. URL <https://pureportal.strath.ac.uk/en/publications/clef-2018-technologically-assisted-reviews-in-empirical-medicine->.
- [32] Evangelos Kanoulas, Dan Li, Leif Azzopardi, and Rene Spijker. CLEF 2019 Technology Assisted Reviews in Empirical Medicine Overview. In *CLEF*, 2019.
- [33] Tushar Khot, Ashish Sabharwal, and Peter Clark. Scitail: A textual entailment dataset from science question answering. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 32, 2018.
- [34] Petr Knoth and Zdenek Zdrahal. Core: three access levels to underpin open access. *D-Lib Magazine*, 18 (11/12), 2012. URL <http://oro.open.ac.uk/35755/>.
- [35] Petr Knoth, Drahomira Herrmannova, Matteo Cancellieri, Lucas Anastasiou, Nancy Pontika, Samuel Pearce, Bikash Gyawali, and David Pride. Core: A global aggregation service for open access papers. *Nature Scientific Data*, 10(1):366, June 2023.
- [36] Georgios Kontonatsios, Austin J. Brockmeier, Piotr Przybyła, John McNaught, Tingting Mu, John Y. Goulermas, and Sophia Ananiadou. A semi-supervised approach using label propagation to support citation screening. *Journal of Biomedical Informatics*, 72:67–76, 8 2017.
- [37] Georgios Kontonatsios, Sally Spencer, Peter Matthew, and Ioannis Korkontzelos. Using a neural network-based feature extraction method to facilitate citation screening for systematic reviews. *Expert Systems with Applications: X*, 6:100030, 7 2020. ISSN 25901885. doi: 10.1016/j.eswx.2020.100030.
- [38] Wojciech Kusa, Allan Hanbury, and Petr Knoth. Automation of citation screening for systematic literature reviews using neural networks: A replicability study. In Matthias Hagen, Suzan Verberne, Craig Macdonald, Christin Seifert, Krisztian Balog, Kjetil Nørkvåg, and Vinay Setty, editors, *Advances in Information Retrieval*, pages 584–598, Cham, 2022. Springer International Publishing. ISBN 978-3-030-99736-6. URL <https://arxiv.org/abs/2201.07534v1>.
- [39] Wojciech Kusa, Petr Knoth, and Allan Hanbury. CRUISE-Screening: Living Literature Reviews Toolbox. In *Proceedings of the 32nd ACM International Conference on Information and Knowledge Management, CIKM '23*, page 5071–5075, New York, NY, USA, 2023. Association for Computing Machinery. ISBN 9798400701245. doi: 10.1145/3583780.3614736. URL <https://doi.org/10.1145/3583780.3614736>.
- [40] Wojciech Kusa, Aldo Lipani, Petr Knoth, and Allan Hanbury. An Analysis of Work Saved over Sampling in the Evaluation of Automated Citation Screening in Systematic Literature Reviews. *Intelligent Systems with Applications*, 18:200193, 2023. ISSN 2667-3053. doi: <https://doi.org/10.1016/j.iswa.2023.200193>. URL <https://www.sciencedirect.com/science/article/pii/S2667305323000182>.
- [41] Wojciech Kusa, Aldo Lipani, Petr Knoth, and Allan Hanbury. VoMBaT: A Tool for Visualising Evaluation Measure Behaviour in High-Recall Search Tasks. In *Proceedings of the 46th International ACM SIGIR Conference on Research and Development in Information Retrieval (SIGIR '23)*, page 5, Taipei, Taiwan, July 23–27 2023. ACM. URL <https://doi.org/10.1145/3539618.3591802>.
- [42] Wojciech Kusa, Oscar E. Mendoza, Petr Knoth, Gabriella Pasi, and Allan Hanbury. Effective Matching of Patients to Clinical Trials using Entity Extraction and Neural Re-ranking. *Journal of Biomedical Informatics*, JBI, 2023.
- [43] Wojciech Kusa, Guido Zuccon, Petr Knoth, and Allan Hanbury. Outcome-based evaluation of systematic review automation. In *Proceedings of the 2023 ACM SIGIR International Conference on Theory of Information Retrieval, ICTIR '23*, page 125–133, New York, NY, USA, 2023. Association for Computing Machinery. ISBN 9798400700736. doi: 10.1145/3578337.3605135. URL <https://doi.org/10.1145/3578337.3605135>.

- [44] Eric W Lee and Joyce C Ho. Pgb: A pubmed graph benchmark for heterogeneous network representation learning. *arXiv preprint arXiv:2305.02691*, 2023.
- [45] Eric W Lee and Joyce C Ho. Sr-comber: Heterogeneous network embedding using community multi-view enhanced graph convolutional network for automating systematic reviews. In *European Conference on Information Retrieval*, pages 553–568. Springer, 2023.
- [46] Mike Lewis, Yinhan Liu, Naman Goyal, Marjan Ghazvininejad, Abdelrahman Mohamed, Omer Levy, Veselin Stoyanov, and Luke Zettlemoyer. BART: Denoising sequence-to-sequence pre-training for natural language generation, translation, and comprehension. In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 7871–7880, Online, July 2020. Association for Computational Linguistics. doi: 10.18653/v1/2020.acl-main.703. URL <https://aclanthology.org/2020.acl-main.703>.
- [47] Quentin Lhoest, Albert Villanova del Moral, Yacine Jernite, Abhishek Thakur, Patrick von Platen, Suraj Patil, Julien Chaumond, Mariama Drame, Julien Plu, Lewis Tunstall, et al. Datasets: A community library for natural language processing. *arXiv preprint arXiv:2109.02846*, 2021.
- [48] Yikuan Li, Ramsey M Wehbe, Faraz S Ahmad, Hanyin Wang, and Yuan Luo. Clinical-longformer and clinical-bigbird: Transformers for long clinical sequences. *arXiv preprint arXiv:2201.11838*, 2022.
- [49] Pengfei Liu, Weizhe Yuan, Jinlan Fu, Zhengbao Jiang, Hiroaki Hayashi, and Graham Neubig. Pre-train, Prompt, and Predict: A Systematic Survey of Prompting Methods in Natural Language Processing. 7 2021. URL <https://arxiv.org/abs/2107.13586v1>.
- [50] Ilya Loshchilov and Frank Hutter. Decoupled weight decay regularization. In *International Conference on Learning Representations*, 2018.
- [51] Iain J Marshall, Joël Kuiper, and Byron C Wallace. Robotreviewer: evaluation of a system for automatically assessing bias in clinical trials. *Journal of the American Medical Informatics Association*, 23(1):193–201, 2016.
- [52] Stan Matwin, Alexandre Kouznetsov, Diana Inkpen, Oana Frunza, and Peter O’Blenis. A new algorithm for reducing the workload of experts in performing systematic reviews. *Journal of the American Medical Informatics Association*, 17(4):446–453, 7 2010. doi: 10.1136/JAMIA.2010.004325.
- [53] Leland McInnes, John Healy, Nathaniel Saul, and Lukas Grossberger. Umap: Uniform manifold approximation and projection. *The Journal of Open Source Software*, 3(29):861, 2018.
- [54] Makoto Miwa, James Thomas, Alison O’Mara-Eves, and Sophia Ananiadou. Reducing systematic review workload through certainty-based screening. *Journal of Biomedical Informatics*, 51:242–253, 10 2014. ISSN 1532-0464. doi: 10.1016/J.JBI.2014.06.005.
- [55] Amirkeivan Mohtashami and Martin Jaggi. Landmark attention: Random-access infinite context length for transformers. *arXiv preprint arXiv:2305.16300*, 2023.
- [56] Christopher Norman. *Systematic review automation methods*. PhD thesis, Université Paris-Saclay ; Universiteit van Amsterdam, 2 2020. URL <https://tel.archives-ouvertes.fr/tel-03060620>.
- [57] Barbara Nussbaumer-Streit, Irma Klerings, Gernot Wagner, Thomas L. Heise, Andreea I. Dobrescu, Susan Armijo-Olivo, Jan M. Stratil, Emma Persad, Stefan K. Lhachimi, Megan G. Van Noord, Tarquin Mittermayr, Hajo Zeeb, Lars Hemkens, and Gerald Gartlehner. Abbreviated literature searches were viable alternatives to comprehensive searches: a meta-epidemiological study. *Journal of Clinical Epidemiology*, 102:1–11, 2018. ISSN 0895-4356. doi: <https://doi.org/10.1016/j.jclinepi.2018.05.022>. URL <https://www.sciencedirect.com/science/article/pii/S0895435618300179>.
- [58] Barbara Nussbaumer-Streit, Moriah Ellen, Irma Klerings, Raluca Sfetcu, Nicoletta Riva, Mersiha Mahmić-Kaknjo, Georgios Poulentzas, P Martinez, Eduard Baladia, Liliya Eugenevna Ziganshina, et al. Resource use during systematic review production varies widely: a scoping review. *Journal of Clinical Epidemiology*, 139:287–296, 2021.
- [59] Benjamin Nye, Junyi Jessie Li, Roma Patel, Yinfei Yang, Iain J Marshall, Ani Nenkova, and Byron C Wallace. A corpus with multi-level annotations of patients, interventions and outcomes to support language processing for medical literature. In *Proceedings of the conference. Association for Computational Linguistics. Meeting*, volume 2018, page 197. NIH Public Access, 2018.
- [60] Alison O’Mara-Eves, James Thomas, John McNaught, Makoto Miwa, and Sophia Ananiadou. Using text mining for study identification in systematic reviews: A systematic review of current approaches. *Systematic Reviews*, 4(1):5, 1 2015. ISSN 20464053. doi: 10.1186/2046-4053-4-5.

- [61] OpenAI. Gpt-4 technical report, 2023.
- [62] Mihir Parmar, Swaroop Mishra, Mirali Purohit, Man Luo, Murad Mohammad, and Chitta Baral. In-BoXBART: Get instructions into biomedical multi-task learning. In *Findings of the Association for Computational Linguistics: NAACL 2022*, pages 112–128, Seattle, United States, July 2022. Association for Computational Linguistics. doi: 10.18653/v1/2022.findings-naacl.10. URL <https://aclanthology.org/2022.findings-naacl.10>.
- [63] Mihir Prafullsinh Parmar. Automation of title and abstract screening for clinical systematic reviews. Master's thesis, Arizona State University, 2021. URL https://keep.lib.asu.edu/_flysystem/fedora/c7/Parmar_asu_0010N_21179.pdf.
- [64] Yifan Peng, Shankai Yan, and Zhiyong Lu. Transfer learning in biomedical natural language processing: An evaluation of BERT and ELMo on ten benchmarking datasets. In *Proceedings of the 18th BioNLP Workshop and Shared Task*, pages 58–65, Florence, Italy, August 2019. Association for Computational Linguistics. doi: 10.18653/v1/W19-5006. URL <https://aclanthology.org/W19-5006>.
- [65] Nils Reimers and Iryna Gurevych. Sentence-BERT: Sentence Embeddings using Siamese BERT-Networks. *EMNLP-IJCNLP 2019 - 2019 Conference on Empirical Methods in Natural Language Processing and 9th International Joint Conference on Natural Language Processing, Proceedings of the Conference*, pages 3982–3992, 8 2019. URL <https://arxiv.org/abs/1908.10084v1>.
- [66] Matthew Richardson, Christopher J.C. Burges, and Erin Renshaw. MCTest: A challenge dataset for the open-domain machine comprehension of text. In *Proceedings of the 2013 Conference on Empirical Methods in Natural Language Processing*, pages 193–203, Seattle, Washington, USA, October 2013. Association for Computational Linguistics. URL <https://aclanthology.org/D13-1020>.
- [67] Kirk Roberts, Dina Demner-Fushman, Ellen M Voorhees, William R Hersh, Steven Bedrick, Alexander J Lazar, and Shubham Pant. Overview of the trec 2017 precision medicine track. In *TREC*, 2017.
- [68] Kirk Roberts, Dina Demner-Fushman, Ellen M Voorhees, Steven Bedrick, and Willian R Hersh. Overview of the trec 2021 clinical trials track. In *Proceedings of the Thirtieth Text REtrieval Conference (TREC 2021)*, 2021.
- [69] Stephen E Robertson, Steve Walker, Susan Jones, Micheline M Hancock-Beaulieu, Mike Gatford, et al. Okapi at trec-3. *Nist Special Publication Sp*, 109:109, 1995.
- [70] Harrisen Scells and Guido Zuccon. Generating better queries for systematic reviews. In *The 41st international ACM SIGIR conference on research & development in information retrieval*, pages 475–484, 2018.
- [71] Harrisen Scells, Guido Zuccon, Bevan Koopman, Anthony Deacon, Leif Azzopardi, and Shlomo Geva. A test collection for evaluating retrieval of studies for inclusion in systematic reviews. *SIGIR 2017 - Proceedings of the 40th International ACM SIGIR Conference on Research and Development in Information Retrieval*, pages 1237–1240, 8 2017. doi: 10.1145/3077136.3080707.
- [72] Harrisen Scells, Guido Zuccon, and Bevan Koopman. A comparison of automatic boolean query formulation for systematic reviews. *Information Retrieval Journal*, 24(1):3–28, 2021.
- [73] Gaurav Singh, James Thomas, and John Shawe-Taylor. Improving active learning in systematic reviews. *arXiv preprint arXiv:1801.09496*, 2018.
- [74] Nandan Thakur, Nils Reimers, Andreas Rücklé, Abhishek Srivastava, and Iryna Gurevych. BEIR: A Heterogenous Benchmark for Zero-shot Evaluation of Information Retrieval Models. 4 2021. URL <https://arxiv.org/abs/2104.08663v4>.
- [75] Guy Tsafnat, Adam Dunn, Paul Glasziou, and Enrico Coiera. The automation of systematic reviews, 2013.
- [76] Guy Tsafnat, Paul Glasziou, Miew K. Choong, Adam Dunn, Filippo Galgani, and Enrico Coiera. Systematic review automation technologies, 4 2014. ISSN 20464053.
- [77] Guy Tsafnat, Paul Glasziou, George Karystianis, and Enrico Coiera. Automated screening of research studies for systematic reviews using study characteristics. *Systematic Reviews 2018 7:1*, 7(1):1–9, 4 2018. ISSN 2046-4053. doi: 10.1186/S13643-018-0724-7. URL <https://link.springer.com/articles/10.1186/s13643-018-0724-7><https://link.springer.com/article/10.1186/s13643-018-0724-7>.
- [78] Raymon van Dinter, Cagatay Catal, and Bedir Tekinerdogan. A decision support system for automating document retrieval and citation screening. *Expert Systems with Applications*, 182, 11 2021.

- [79] Raymon van Dinter, Cagatay Catal, and Bedir Tekinerdogan. A Multi-Channel Convolutional Neural Network approach to automate the citation screening process. *Applied Soft Computing*, 112:107765, 11 2021. ISSN 1568-4946. doi: 10.1016/J.ASOC.2021.107765.
- [80] Raymon van Dinter, Bedir Tekinerdogan, and Cagatay Catal. Automation of systematic literature reviews: A systematic literature review. *Information and Software Technology*, 136:106589, 8 2021. ISSN 09505849. doi: 10.1016/j.infsof.2021.106589. URL <https://linkinghub.elsevier.com/retrieve/pii/S0950584921000690>.
- [81] Byron C. Wallace, Kevin Small, Carla E. Brodley, and Thomas A. Trikalinos. Active learning for biomedical citation screening. *Proceedings of the ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, pages 173–181, 2010. doi: 10.1145/1835804.1835829.
- [82] Byron C Wallace, Thomas A Trikalinos, Joseph Lau, Carla Brodley, and Christopher H Schmid. Semi-automated screening of biomedical citations for systematic reviews. *BMC bioinformatics*, 11(1):1–11, 2010.
- [83] Byron C Wallace, Kevin Small, Carla E Brodley, and Thomas A Trikalinos. Who should label what? instance allocation in multiple expert active learning. In *Proceedings of the 2011 SIAM international conference on data mining*, pages 176–187. SIAM, 2011.
- [84] Byron C. Wallace, Sayantani Saha, Frank Soboczinski, and Iain James Marshall. Generating (factual?) narrative summaries of rcts: Experiments with neural multi-document summarization. *AMIA Annual Symposium*, abs/2008.11293, 2020.
- [85] Alex Wang, Amanpreet Singh, Julian Michael, Felix Hill, Omer Levy, and Samuel R Bowman. GLUE: A multi-task benchmark and analysis platform for natural language understanding. *arXiv preprint arXiv:1804.07461*, 2018.
- [86] Lucy Lu Wang, Jay DeYoung, and Byron Wallace. Overview of MSLR2022: A shared task on multi-document summarization for literature reviews. In *Proceedings of the Third Workshop on Scholarly Document Processing*, pages 175–180, Gyeongju, Republic of Korea, October 2022. Association for Computational Linguistics. URL <https://aclanthology.org/2022.sdp-1.20>.
- [87] Shuai Wang, Harris Scells, Ahmed Mourad, and Guido Zuccon. Seed-driven Document Ranking for Systematic Reviews: A Reproducibility Study. 12 2021. URL <https://arxiv.org/abs/2112.04090v1>.
- [88] Shuai Wang, Harris Scells, Justin Clark, Bevan Koopman, and Guido Zuccon. From little things big things grow: A collection with seed studies for medical systematic review literature search. *arXiv preprint arXiv:2204.03096*, 2022.
- [89] Shuai Wang, Hang Li, and Guido Zuccon. Mesh suggester: A library and system for mesh term suggestion for systematic review boolean query construction. In *Proceedings of the Sixteenth ACM International Conference on Web Search and Data Mining, WSDM '23*, page 1176–1179, New York, NY, USA, 2023. Association for Computing Machinery. ISBN 9781450394079. doi: 10.1145/3539597.3573025. URL <https://doi.org/10.1145/3539597.3573025>.
- [90] Shuai Wang, Harris Scells, Bevan Koopman, and Guido Zuccon. Can ChatGPT write a good boolean query for systematic review literature search? *arXiv preprint arXiv:2302.03495*, 2023.
- [91] Yizhong Wang, Swaroop Mishra, Pegah Alipoormolabashi, Yeganeh Kordi, Amirreza Mirzaei, Atharva Naik, Arjun Ashok, Arut Selvan Dhanasekaran, Anjana Arunkumar, David Stap, Eshaan Pathak, Giannis Karamanolakis, Haizhi Lai, Ishan Purohit, Ishani Mondal, Jacob Anderson, Kirby Kuznia, Krima Doshi, Kuntal Kumar Pal, Maitreya Patel, Mehrad Moradshahi, Mihir Parmar, Mirali Purohit, Neeraj Varshney, Phani Rohitha Kaza, Pulkit Verma, Ravsehaj Singh Puri, Rushang Karia, Savan Doshi, Shailaja Keyur Sampat, Siddhartha Mishra, Sujay Reddy A, Sumanta Patro, Tanay Dixit, and Xudong Shen. Super-NaturalInstructions: Generalization via declarative instructions on 1600+ NLP tasks. In *Proceedings of the 2022 Conference on Empirical Methods in Natural Language Processing*, pages 5085–5109, Abu Dhabi, United Arab Emirates, December 2022. Association for Computational Linguistics. URL <https://aclanthology.org/2022.emnlp-main.340>.
- [92] Leon Weber, Mario Sanger, Jannes Munchmeyer, Maryam Habibi, Ulf Leser, and Alan Akbik. Hunflair: an easy-to-use tool for state-of-the-art biomedical named entity recognition. *Bioinformatics*, 37(17): 2792–2794, 2021.

- [93] Wenpeng Yin, Dragomir Radev, and Caiming Xiong. DocNLI: A large-scale dataset for document-level natural language inference. In *Findings of the Association for Computational Linguistics: ACL-IJCNLP 2021*, pages 4913–4922, Online, August 2021. Association for Computational Linguistics. doi: 10.18653/v1/2021.findings-acl.435. URL <https://aclanthology.org/2021.findings-acl.435>.
- [94] Hye Sun Yun, Iain J Marshall, Thomas Trikalinos, and Byron C Wallace. Appraising the potential uses and harms of LLMs for medical systematic reviews. *arXiv preprint arXiv:2305.11828*, 2023.
- [95] Manzil Zaheer, Guru Guruganesh, Kumar Avinava Dubey, Joshua Ainslie, Chris Alberti, Santiago Ontanon, Philip Pham, Anirudh Ravula, Qifan Wang, Li Yang, et al. Big bird: Transformers for longer sequences. *Advances in Neural Information Processing Systems*, 33:17283–17297, 2020.
- [96] Ningyu Zhang, Mosha Chen, Zhen Bi, Xiaozhuan Liang, Lei Li, Xin Shang, Kangping Yin, Chuanqi Tan, Jian Xu, Fei Huang, Luo Si, Yuan Ni, Guotong Xie, Zhifang Sui, Baobao Chang, Hui Zong, Zheng Yuan, Linfeng Li, Jun Yan, Hongying Zan, Kunli Zhang, Buzhou Tang, and Qingcai Chen. CBLUE: A Chinese biomedical language understanding evaluation benchmark. In *Proceedings of the 60th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 7888–7915, Dublin, Ireland, May 2022. Association for Computational Linguistics. doi: 10.18653/v1/2022.acl-long.544. URL <https://aclanthology.org/2022.acl-long.544>.