# Cutting Through the Clutter: The Potential of LLMs for Efficient Filtration in Systematic Literature Reviews



<sup>1</sup>University of Konstanz, Germany



Figure 1: Schematic overview of leveraging LLM-based agents for structured literature filtration in a systematic review (SLR). Keywordbased search in online libraries generates a large set of candidate papers that are classified by multiple LLMs based on title and abstract using a customized prompt. A consensus voting scheme determines inclusion or rejection, providing justifications that users can review and refine.

## Abstract

Systematic literature reviews (SLRs) are essential but labor-intensive due to high publication volumes and inefficient keywordbased filtering. To streamline this process, we evaluate Large Language Models (LLMs) for enhancing efficiency and accuracy in corpus filtration while minimizing manual effort. Our open-source tool LLMSurver presents a visual interface to utilize LLMs for literature filtration, evaluate the results, and refine queries in an interactive way. We assess the real-world performance of our approach in filtering over 8.3k articles during a recent survey construction, comparing results with human efforts. The findings show that recent LLM models can reduce filtering time from weeks to minutes. A consensus scheme ensures recall rates >98.8%, surpassing typical human error thresholds and improving selection accuracy. This work advances literature review methodologies and highlights the potential of responsible human-AI collaboration in academic research.

## **CCS Concepts**

• Human-centered computing  $\rightarrow$  Interactive systems and tools; • Computing methodologies  $\rightarrow$  Artificial intelligence; • Applied computing  $\rightarrow$  Publishing;

#### 1. Introduction

Literature reviews, and in particular systematic literature reviews (SLRs), have been described as the *gold standard* for conducting literature research in academia [ESA01, DMBM14]. They provide a transparent and reproducible approach to systematically synthesize and categorize research findings, providing a comprehensive overview of a research topic [Nig09]. Such reviews date back to the 18th century [vTC21] and help identify research gaps, future directions, and ensuring consistency and reliability in academic

© 2025 The Authors.

research [Lam19]. The creation of SLRs, however, is typically a highly manual and labor-intensive process [SJD21]. Egger et al. describe the standard process through a set of eight stages (1. research question, 2. criteria definition, 3. locating, 4. selection, 5. assessment, 6. data extraction, 7. presentation, 8. interpretation) [ESA01]. With *PRISMA* [LAT\*09], a well-established, standardized method exists, describing the process of retrieving a paper corpus (steps 3–6) by keyword-based search, duplicate removal, manual screening based on title and abstract, and the full-text manuscript review. One of the most time-consuming tasks in this pipeline is the man-

Proceedings published by Eurographics - The European Association for Computer Graphics.

This is an open access article under the terms of the Creative Commons Attribution License, which permits use, distribution and reproduction in any medium, provided the original work is properly cited.

ual title and abstract screening, particularly in domains or research fields where classical keyword-based filtering may lead to ambiguous results. According to Wallace et al. [WTL\*10], an experienced peer-reviewer can manually screen about two papers per minute based on title and abstract. At this rate, a corpus of about 8,000 potentially relevant publications for a large SLR requires approximately 66 person-hours (about one and a half full work weeks) of uninterrupted work time. Effects like fatigue, loss of accuracy, inefficiencies, dual verification, and other work commitments typically increase the required time frame significantly, resulting in survey latency times closer to a few months for the initial screening alone. Given the repetitive but still demanding nature of the tasks and the ever-faster progress in academia, it stands to reason if-and how-this process can be improved upon. Using automation for such a (relatively) well-defined classification task is not a new idea, with the first use of automation being reported in the mid-2000s [vTC21]. However, the recent advancements of Large Language Models (LLMs) prove promising for the tasks of initial literature filtration during the creation of SLRs primarily due to two reasons: (1) their capability to understand nuanced semantic ambiguities, potentially reaching a feasible accuracy (recall, precision) threshold, and (2) their unparalleled speed and cost-efficiency w.r.t. to human labor. While existing LLM chatbots have the ability to search external databases through function calls, limited research has been conducted on a schematic pipeline of the whole process from repository acquisition to final paper selection and its evaluation, ensuring completeness, reliability, and accountability, which is the focus of this research. In this work, we present a visualinteractive approach leveraging LLMs for literature filtration during SLR creation, allowing users to iteratively refine prompts, evaluate the results, and interactively create a consensus scheme leading to the desired classification result. Thereby, we make the following contributions:

- A conceptual schema for the structured literature filtration process leveraging LLM-based agents with consensus voting
- A visual-interactive open-source application, LLMSurver, implementing our framework and making it accessible to others
- A comprehensive **evaluation** for a large SLR (8.3k papers) with an extensive **discussion** on the pitfalls, potentials, and future prospects of leveraging AI agents for literature filtration

## 2. Related Work

With the recent successes of machine learning and in particular LLMs, an increasing number of publications show how language models can leverage some parts of the tasks [WTL\*10, SJD21] involved in the scientific publishing process [LWM\*23]. These tasks include, for instance, generating paper reviews (e.g., to improve the own work) [TSL\*24], reformulating paragraphs for clarity [GC23], identifying and avoiding biased arguments [HT23], or finding gaps in previous research for a given domain [LWM\*23].

Besides these general publishing tasks, LLMs have recently been shown to also support various aspects of conducting literature reviews [SR\*24]. While a large body of research on how to conduct literature reviews exists [Nig09, DMBM14, ESA01, Lam19], many of the developed methods are labor-intensive and repetitive. Therefore, it has been investigated how agent-based systems can help with the formulation, filtering, and search of a research domain [WH23, SR\*24, HT23], using LLMs for specific keyword generation and retrieval through RAG [ALCP24], or more generally, how machine learning [WTL\*10, vTC21, SJD21], but also LLMs [ACR23, Sus23, RMBK23, BSOM24, HT24, PBH\*24] can support the overall process. Further, the summarization step may be supported using LLMs [LCL\*24]. One particular aspect that has received less attention is the accurate filtering and classification of a (relatively) large body of potentially relevant research concerning a particular research question to speed up the paper pre-selection process. This is particularly relevant for topics or domains where keyword-based filtering is difficult to use, for example, due to semantic ambiguities or duplicated word use. Haryanto [Har24] explores the usability of LLMs for performing this specific task, focusing on the vote of individual LLMs. Also, fairly recently, automatic tooling approaches for SLR-generation using LLMs have been proposed [SHJ\*24, GLAACG24, Jaf24, SHR\*25]. Gehrmann et al. [GQB24] introduced the only LLM-based automated preselection approach, showing that negative prompting can boost accuracy. Building on this, we propose a similar pipeline for classifying large paper corpora, designed as a visual-interactive process that incorporates and evaluates voting schemes from multiple LLM agents and lets users iteratively refine prompts and LLMs. We also compare results with a manual SLR selection on the same dataset, offering insights into reliability and accuracy.

## 3. Methodology

To evaluate the applicability of LLMs for pre-filtering the paper corpus for an SLR, we followed a structured methodology, starting with a topic definition, using an early, preliminary version of our recent literature survey "Visual Network Analysis in Immersive Environments: A Survey" [JFR\*25]. This topic leads to a sufficiently large corpus of potential papers since all papers dealing with some immersive technology, such as Virtual Reality, Augmented Reality, and others, focusing on the widespread data type of graphs are of relevance. For the initial paper selection, we followed the PRISMA pipeline [LAT<sup>\*</sup>09], starting with a structural keyword-based search in paper titles and abstracts of potential paper candidates in major computer science repositories. In our case, we included papers from the ACM Digital Library, IEEE Xplore, and Eurographics. After unifying the format of the paper metadata, removing duplicates, and excluding all non-paper publications, we retrieved an initial corpus of 8,323 papers in the preliminary version (the published version used a later iteration with more results). Papers were manually screened in multiple iterations. This process was highly time-

You are a professor in computer science conducting a literature review. Please decide and classify if the following paper belongs to a specific research direction or not. For this, you are provided with the title and the abstract, which should give you sufficient information for an informed and accurate decision. The research direction is the topic of "ITILE". Therefore include papers that deal with ASPECT\_1, ASPECT\_2, ... Examples of AS-PECT\_1 are: term 1, term 2, ... You MUST discard papers that EXCLUSION\_EXCEPTION\_1, ... You MUST include papers that INCLUSION\_EXCEPTION\_1, ... Below is the title and abstract. You must only answer with INCLUDE or DISCARD and a 2-sentence reason of why.

Figure 2: Prompt template for the individual agents.



Figure 3: Number of papers (gray background) that were incorrectly (inc.) voted to be included (left) or excluded (right) by the agents, grouped by the number of incorrect agents involved in a decision. The individual bars show how many times a particular agent was involved in the wrong decision. It can be seen on the far right that only a single paper is unanimously misjudged by all agents (and therefore lost forever), demonstrating that N-Consensus voting is beneficial when prioritizing recall.

Table 1: Evaluation results of the LLM agents and two consensus schemes (all models and best models only) for our reference survey, with the validated human classification as ground truth.

	Metric	Llama-3 (8B) <sup>1</sup>	Llama-3 (70B) <sup>2</sup>	Gemini 1.5 Flash <sup>3</sup>	Claude 3.5 Sonnet <sup>4</sup>	GPT-40 <sup>5</sup>	Consensus (All) <sup>6</sup>	Consensus (Best) <sup>7</sup>
Counts	<b>TP</b> (↑)	86	85	67	76	80	87	87
	$FP(\downarrow)$	774	194	91	95	50	862	167
	TN $(\uparrow)$	7461	8041	8144	8140	8185	7373	8068
	$FN(\downarrow)$	2	3	21	12	8	1	1
Evaluation	Acc. $(\uparrow)$	90.68	97.63	98.65	98.71	99.30	89.63	97.98
	Prec. $(\uparrow)$	10.00	30.47	42.41	44.44	61.54	9.17	34.25
	Rec. $(\uparrow)$	97.73	96.59	76.14	86.36	90.91	98.86	98.86
	$F_1 \ (\uparrow)$	18.14	46.32	54.47	58.69	73.39	16.78	50.88

meta-llama-3-8b-instruct.Q8\_0 2 meta-llama-3-70b-instruct.Q4\_K\_M gemini-1.5-flash-001 <sup>4</sup> claude-3-5-sonneces... <u>9024-05-13</u> <sup>6</sup> Consensus between all (five) models. <sup>4</sup> claude-3-5-sonnet@20240620 <sup>5</sup> gpt-40-2024-05-13 7 Consensus between

models with  $F_1 > 50\%$  (i.e. without Llama3 variants).

consuming, involving multiple researchers for multiple weeks, but led to the ground truth categorization: for the initial corpus, we identified 88 papers that needed to be included and 8235 to discard. Based on the ground-truth categorization, we investigate the potential for LLMs to facilitate the laborious process of filtering a paper corpus, considering five open and commercial state-of-the-art foundation models (see Table 1 for details).

We initially tested the LLMs with different prompt styles, asking the models to classify each paper individually, and quickly found a basic prompt schema that works well. In this schema, we tell the LLM its context and role, the overall task, before concluding with an output format and the paper title and abstract. For the final prompt (see Figure 2), we added further exclusion and inclusion criteria, leading to the following results.

#### 4. Evaluation

The results of the individual LLM classifications are summarized in Table 1. In general, the LLMs performed well with an accuracy above 90 % across all models. However, there are still notable differences: While the open-source models-especially Llama3 8Bwere more conservative, including more papers in general (high FP rate), trying not to exclude any relevant papers (low FN rate), the commercial models discarded more papers (higher TN rate), but with the downside of having more papers erroneously excluded (higher FN rate). Interestingly, the falsely classified papers were

© 2025 The Authors Proceedings published by Eurographics - The European Association for Computer Graphics mostly different across the LLMs: Regarding the erroneous inclusions (FP), for most papers, only one LLM-often Llama3 8B-was responsible for the wrong classification (see Figure 3 left). The number of papers to exclude that were falsely included by multiple LLMs is drastically lower. This is also the case for relevant, incorrectly discarded papers (FN), where mostly individual LLMs (mostly Gemini 1.5 Flash) generated errors, but false exclusions by multiple LLMs were way lower (see Figure 3 right). Therefore, we also analyzed the performance of a consensus voting of all LLMs-Consensus (All)-and a selection of the best-performing LLMs with an F1 score above 50 %-Consensus (Best)-consisting of Gemini 1.5 Flash, Claude 3.5 Sonnet, and GPT-40. For consensus voting, a paper is only discarded if all of the involved LLMs agree to discard it-and included if at least one LLM includes it. The results of both consensus approaches (see Table 1, right column) are highly encouraging, showing great results, especially for the TP and the FN rates. By consensus voting, only one paper would be discarded that should be part of the survey (based on the human ground-truth data). A manual inspection revealed that this paper was also an edge case for the involved researchers, who might have excluded the paper based on the abstract and title but ultimately included it after investigating its content. While both consensus approaches lead to the same TP and FN rates, which are of most relevance for our use case, the Consensus (Best) approach comes with a lower FP rate (see Figure 4), reducing the manual filtering by 695 papers, and only requires three instead of five LLMs, reducing time and cost.

#### 5. Interactive Human-AI Collaboration

The results of our experiment show that LLMs can support the initial filtering process. However, relying solely on (individual) LLMs without human intervention is of high risk. Therefore, we propose a new pipeline (see Figure 1) to form the paper corpus of survey papers, incorporating a tight collaboration between the researcher (human) and LLMs (AI). The first steps of our suggested pipeline remain the same as for classical paper retrieval (see Section 3): Online repositories are searched for papers of relevance based on keywords, leading to a pre-processed initial paper corpus. Then, multiple LLMs classify each paper independently. The survey authors are an essential part of the process, as they iteratively create and adapt prompts (as in Figure 2) and investigate sampled LLM output through a visual-interactive interface until the results are of sufficient quality. Investigating the LLMs' justification of their de-



Figure 4: Pairwise comparison of the *incorrect* decisions by the agents. A paper's validated ground truth classification is shown in the top row (left part: included, right part: discarded), followed by the individual agents and then the two consensus methods. Incorrect exclusions (FN) can be seen as red discarded lines (left part), while incorrect inclusion (FP) can be seen as green included lines (larger right part).

cisions is often highly useful for evaluating the prompt and for refining it. When the preliminary results are sufficient, all papers are classified by each LLM (or the most promising ones), and the results are combined through a consensus voting. Again, the consensus results can be iteratively adapted and evaluated through a visual-interactive process, highlighting similarities and differences across the LLMs (similar to Figure 3). As our evaluation demonstrated, consensus voting is highly effective in reducing the number of papers, while the rate of erroneously removed papers remains very low. Human classification can similarly result in false exclusions, although these papers can typically be recovered in a subsequent snowballing step [Woh14]. Therefore, a small number of removals during the LLM step may be considered acceptable.

# 6. Application

4 of 6

We developed the interactive open-source application LLMSurver (https://github.com/dbvis-ukon/LLMSurver) featuring a user interface (UI) implementing our proposed pipeline. This tool demonstrates the practical application and provides support for researchers conducting their own literature surveys. The application is fully containerized and follows a frontend (single-page React), backend Python-based FastAPI) database (SQLite) architecture. The UI is structured as a dashboard, with visually distinct components reflecting the pipeline steps (see Figure 5). The main table <sup>(2)</sup> displays paper details from the corpus, populated by uploading Bibtex files or providing DOI numbers. A prompt editor component <sup>9</sup> allows users to craft and refine classification prompts for selected LLMs in component 0. Users can register new LLMs (local or remote) by entering necessary details such as API keys or hostnames. Classifications can be applied to subsets for testing or the entire corpus, with intermediate results saved. Classification results are visualized in the main table and can be exported as a CSV file. Users can view individual LLM outputs, particularly useful for ambiguous results (indicated by an orange error icon). The consensus component • enables run selection, statistics visualization, and LLM selection for consensus-building, with results reflected in the main table. To aid decision-making, component <sup>1</sup> presents two charts: one shows the classification distribution across LLMs, while the other visualizes agreement levels, highlighting outliers (e.g., LLama3 8B in our evaluation) that may reduce consensus quality. The opensource tool is adaptable for other use cases, supporting custom consensus methods or additional decision-making visualizations.

#### 7. Discussion

We have demonstrated that incorporating AI techniques, particularly LLM-based agents, into a structured analysis pipeline can effectively support the initial stages of a systematic literature review with surprisingly high quality. A key advantage is the speed and cost-efficiency of filtration. In our case study with 8,323 papers, GPT-40 processed 4,432,169 input tokens-approximately 532 tokens per paper (including prompts)-and 443,735 output tokens, or around 53 per paper. This entire process was completed in under 10 minutes for just \$28.81 (as of July 2024), demonstrating the scalability of LLMs. Their ability to operate continuously or scale up through additional GPU resources makes them ideal for large-scale literature reviews. Compared to manual filtering, which requires a minimum of 69 hours of concentrated human effort, this represents a significant reduction in both time and cost, making systematic reviews more accessible and efficient. When cost or confidentiality is a concern, smaller, open-source models-capable of running locally on a standard laptop-still achieve impressive recall rates of 97.73%, though with lower precision. Even so, these models allow researchers to explore research fields quickly, reducing the manual search space by nearly 90% (from 8,323 papers to 860) in just a few hours. Notably, the recall difference between top-performing models and Llama3 8B was minimal, with only one additional paper lost as a false positive. The model's bias towards inclusion requires further investigation. To enhance precision, using a consensus approach reduced false positives from 774 to 167-a 98% re-



Figure 5: The user interface of our application implementing the proposed pipeline, consisting of a paper table <sup>O</sup>, a prompt definition area <sup>O</sup>, a panel for LLM selection classification runs <sup>O</sup>, the consensus scheme with statistics <sup>O</sup>, and visual plots <sup>O</sup>.

duction in the validation space, missing only one out of 88. These results are comparable to **human error thresholds**, which vary depending on factors like task difficulty, familiarity, stress, and repetition [SS11]. Established frameworks such as HEART [Hum88], TESEO [BC80], and THERP [Kir88] suggest error rates ranging from 0.5% to 9%, placing our filtration results within or even below these ranges. Additionally, during our manual review process, 34 papers were reclassified after initial human filtration, further highlighting the strengths of our LLM-based approach.

A significant advantage of using LLMs is the ability to generate consistent and descriptive classification explanations through prompting, a task that would require considerable additional effort from human reviewers. Automating the initial filtration phase also leads to better resource allocation, allowing researchers to focus on higher-level analysis and interpretation, thereby improving overall productivity while reducing fatigue from repetitive tasks. This efficiency enables researchers to explore a more diverse range of research fields by lowering the entry cost of initial surveys. Additionally, it can help identify gaps in existing literature by semi-automatically gathering relevant publications for broader overviews of specific topics. The multilingual capabilities of LLMs further enhance the accessibility of non-English academic literature, facilitating the inclusion of relevant publications from specialized venues or fields with older literature not available in English. Finally, automation inherently improves data management. Using a pipeline architecture helps structure large datasets, making the literature easier to navigate compared to manual processes.

#### 7.1. Limitations and Future Work

The use of automation and generative models presents several challenges and limitations. Our study is based on a single large corpus and prompt, which may not generalize to other research areas. Also, our tool has not yet been evaluated in a controlled user study. A validity risk, in particular when avoiding snowballing [Woh14], is a careful selection of the initial set of bibliographical entries and source databases. Other factors, such as prompt design, corpus characteristics, or writing style, could also influence performance. Nevertheless, given the strong text comprehension abilities of LLMs, their potential for literature filtration remains promising, warranting further investigation into their capabilities and limitations. LLMs face well-known challenges, including hallucinations, biases from training data, and accuracy concerns. To ensure completeness, we limited the role of generative AI to classification within a structured schema, avoiding direct involvement in the search process and minimizing the risk of generating false references [HQS<sup>\*</sup>23]. Despite high accuracy rates, these models can still produce misleading outputs, with performance influenced by model quality, prompt formulation, and contextual understanding. Our study did not focus on prompt engineering [WFH\*23], which could potentially improve outcomes. While our approach primarily employed zero-shot learning with contextual examples, exploring few-shot learning could further enhance accuracy, albeit with increased token usage. Inherent biases, originating from training data or the Reinforcement Learning from Human Feedback (RLHF) process [BJN\*22], can also lead to skewed or incomplete results. Addressing these biases through interactive feedback loops and visual analytics [FHJ\*22] is essential for ensuring research ac-

© 2025 The Authors. Proceedings published by Eurographics - The European Association for Computer Graphics. curacy. Although state-of-the-art commercial models demonstrate the highest performance, they present **access limitations** due to cost, availability, and rate restrictions, potentially disadvantaging smaller research groups or independent researchers [BHA\*22]. Developing reliable **evaluation metrics** for LLM-generated literature surveys is an important area for future research. A potential risk is the **over-reliance on automation**, which could undermine researchers' critical thinking and analytical skills [BHA\*22]. Balancing automation with human oversight [FHJ\*22] remains essential.

Future research should explore the development of interactive literature review platforms where LLMs assist researchers in a **collaborative environment**, integrating user feedback mechanisms into the review process. While our approach facilitates keywordbased paper search, the potential of new LLMs with access to search engines for **retrieving the paper corpus** (prompt-based) should be investigated. Extending the use of LLMs to support **semi-automatic paper coding**–especially when full-text papers are available–could help evaluate the interpretative capabilities of language models more effectively. Additionally, applying this approach to conduct SLRs across **multiple disciplines** could enable broader, cross-disciplinary analyses, facilitating research efforts that were previously infeasible due to scale or complexity.

#### 8. Conclusion

This work evaluates the potential of LLMs to enhance filtration in academic literature reviews. We propose a semi-automated filtration schema for systematic reviews, leveraging recent foundation models-Llama3 (8B and 70B), Gemini 1.5 Flash, Claude 3.5 Sonnet, and GPT-4o-as classification agents to filter large corpora of publications relevant to specific research questions. Our method addresses the limitations of traditional keyword-based filtering, which often struggles with semantic ambiguities and inconsistent terminology, requiring time-consuming manual checks. With our opensource tool LLMSurver, users can iteratively test different prompts and LLMs while interactively evaluating the results. We assess LLM performance during the construction of a recent literature survey [JFR\*25], comparing results against human filtering on a dataset of 8,323 articles. The findings show that LLMs can drastically accelerate the review process, shrinking search space by an order of magnitude and reducing weeks of effort to minutes, while maintaining recall (> 98%), even below typical human error rates. This efficiency not only enhances SLR but also holds promise for broader academic applications. Overall, this study highlights the effective use of LLMs to streamline academic research.

#### Acknowledgements

The authors gratefully acknowledge financial support by the Federal Ministry for Economic Affairs and Climate Action (BMWK, grant No. 03EI1048D) and the Deutsche Forschungsgemeinschaft (DFG) – Project-ID 251654672 – TRR 161.

#### References

[ACR23] ANTU S. A., CHEN H., RICHARDS C. K.: Using llm to improve efficiency in literature review for undergraduate research. 2

6 of 6

- [ALCP24] AGARWAL S., LARADJI I. H., CHARLIN L., PAL C.: Litllm: A toolkit for scientific literature review, 2024. doi:10.48550/ ARXIV.2402.01788.2
- [BC80] BELLO G., COLOMBARI V.: The human factors in risk analyses of process plants: The control room operator model 'teseo'. *Reliability engineering 1*, 1 (1980), 3–14. 5
- [BHA\*22] BOMMASANI R., HUDSON D. A., ADELI E., ALTMAN R., ARORA S., VON ARX S., BERNSTEIN M. S., BOHG J., BOSSELUT A., BRUNSKILL E., ET AL.: On the opportunities and risks of foundation models, 2022. doi:10.48550/arXiv.2108.07258.5
- [BJN\*22] BAI Y., JONES A., NDOUSSE K., ASKELL A., ET AL.: Training a helpful and harmless assistant with reinforcement learning from human feedback, 2022. doi:10.48550/arXiv.2204.05862.5
- [BSOM24] BOLANOS F., SALATINO A., OSBORNE F., MOTTA E.: Artificial intelligence for literature reviews: Opportunities and challenges. *Artificial Intelligence Review* 57, 10 (2024), 259. doi:10.1007/ s10462-024-10902-3.2
- [DMBM14] DAVIS J., MENGERSEN K., BENNETT S., MAZEROLLE L.: Viewing systematic reviews and meta-analysis in social research through different lenses. *SpringerPlus 3* (2014), 1–9. doi:10.1186/ 2193-1801-3-511. 1, 2
- [ESA01] EGGER M., SMITH G. D., ALTMAN D.: Systematic reviews in health care: meta-analysis in context. 2001. doi:10.1002/ 9780470693926.1,2
- [FHJ\*22] FISCHER M. T., HIRSBRUNNER S. D., JENTNER W., MILLER M., KEIM D. A., HELM P.: Promoting ethical awareness in communication analysis: Investigating potentials and limits of visual analytics for intelligence applications. In *Proc. FAcct* '22 (2022), ACM, pp. 877–889. doi:10.1145/3531146.3533151.5
- [GC23] GILAT R., COLE B. J.: How will artificial intelligence affect scientific writing, reviewing and editing? the future is here... Arthroscopy: The Journal of Arthroscopic & Related Surgery 39, 5 (2023), 1119–1120. doi:10.1016/j.arthro.2023.01.014.2
- [GLAACG24] GANA B., LEIVA-ARAOS A., ALLENDE-CID H., GAR-CÍA J.: Leveraging llms for efficient topic reviews. *Applied Sciences 14*, 17 (2024), 7675. doi:10.3390/app14177675. 2
- [GQB24] GEHRMANN J., QUAKULINSKI L., BEYAN O.: Large language models for literature reviews-an exemplary comparison of llmbased approaches with manual methods. In *Proc. FLLM* (2024), IEEE, pp. 385–391. doi:10.1109/FLLM63129.2024.10852447.2
- [Har24] HARYANTO C. Y.: Llassist: Simple tools for automating literature review using large language models, 2024. doi:10.48550/ arXiv.2407.13993.2
- [HQS\*23] HADI M. U., QURESHI R., SHAH A., IRFAN M., ET AL.: A survey on large language models: Applications, challenges, limitations, and practical usage. *Authorea Preprints* (2023). 5
- [HT23] HUANG J., TAN M.: The role of chatgpt in scientific communication: writing better scientific review articles. AJCR 13, 4 (2023). 2
- [HT24] HAWKINS J., TIVEY D.: Efficient systematic reviews: Literature filtering with transformers & transfer learning, 2024. doi:10.48550/arXiv.2405.20354.2
- [Hum88] HUMPHREYS P.: Human reliability assessors guide: an overview. Human factors and decision making: their influence on safety and reliability (1988). 5
- [Jaf24] JAFARI S. M. A.: Streamlining the selection phase of systematic literature reviews (slrs) using ai-enabled gpt-4 assistant api, 2024. doi: 10.48550/arXiv.2402.18582.2
- [JFR\*25] JOOS L., FISCHER M. T., RAUSCHER J., KEIM D. A., DWYER T., SCHREIBER F., KLEIN K.: Visual network analysis in immersive environments: A survey, 2025. doi:10.48550/arXiv. 2501.08500.2,5
- [Kir88] KIRWAN B.: A comparative evaluation of five human reliability assessment techniques. In *Human Factors and Decision Making: Their* influence on safety and reliability. 1988. 5

- [Lam19] LAME G.: Systematic literature reviews: An introduction. In Proceedings of the design society: Int. conf. on engineering design (2019), pp. 1633–1642. doi:10.1017/dsi.2019.169.1,2
- [LAT\*09] LIBERATI A., ALTMAN D. G., TETZLAFF J., MULROW C., GØTZSCHE P. C., ET AL.: The prisma statement for reporting systematic reviews and meta-analyses of studies that evaluate healthcare interventions. *BMJ 339* (2009). doi:10.1136/bmj.b2700.1,2
- [LCL\*24] LI Y., CHEN L., LIU A., YU K., WEN L.: Chatcite: Llm agent with human workflow guidance for comparative literature summary, 2024. doi:10.48550/arXiv.2403.02574.2
- [LWM\*23] LUND B. D., WANG T., MANNURU N. R., NIE B., SHIM-RAY S., WANG Z.: Chatgpt and a new academic reality: Artificial intelligence-written research papers and the ethics of the large language models in scholarly publishing. J. of the Ass. for Inf. Science and Tech 74, 5 (2023), 570–581. doi:10.1002/asi.24750.2
- [Nig09] NIGHTINGALE A.: A guide to systematic literature reviews. Surgery (Oxford) 27, 9 (2009), 381–384. Determining surgical efficacy. doi:10.1016/j.mpsur.2009.07.005. 1, 2
- [PBH\*24] PEINL R., BAERNTHALER J., HABERL A., CHOUGULEY S. R., THALMANN S.: Using llms to improve reproducibility of literature reviews. *Proc. of 2024 Pre-ICIS SIGDSA Symposium* (2024). 2
- [RMBK23] RATHI H., MALIK A., BEHERA D., KAMBOJ G.: P21 a comparative analysis of large language models (llm) utilised in systematic literature review. *Value in Health 26*, 12 (2023), S6. doi: 10.1016/j.jval.2023.09.030.2
- [SHJ\*24] SCHERBAKOV D., HUBIG N., JANSARI V., BAKUMENKO A., LENERT L. A.: The emergence of large language models (llm) as a tool in literature reviews: an llm automated systematic review, 2024. doi: 10.48550/arXiv.2409.04600.2
- [SHR\*25] SUSNJAK T., HWANG P., REYES N. H., BARCZAK A. L. C., MCINTOSH T. R., RANATHUNGA S.: Automating research synthesis with domain-specific large language model fine-tuning. ACM Trans. Knowl. Discov. Data (2025). doi:10.1145/3715964.2
- [SJD21] SILVA JÚNIOR E. M. D., DUTRA M. L.: A roadmap toward the automatic composition of systematic literature reviews. *IJSMC 1*, 2 (2021), 1–22. doi:10.47909/ijsmc.52.1,2
- [SR\*24] SAMI A. M., RASHEED Z., ET AL.: System for systematic literature review using multiple ai agents: Concept and an empirical evaluation, 2024. doi:10.48550/ARXIV.2403.08399.2
- [SS11] SMITH D. J., SIMPSON K. G.: Chapter 5 reliability modeling techniques. In Safety Critical Systems Handbook. 2011, pp. 89–106. doi:10.1016/B978-0-08-096781-3.10005-7.5
- [Sus23] SUSNJAK T.: Prisma-dfilm: An extension of prisma for systematic literature reviews using domain-specific finetuned large language models, 2023. doi:10.48550/arXiv.2306.14905.2
- [TSL\*24] TYSER K., SEGEV B., LONGHITANO G., ET AL.: Ai-driven review systems: Evaluating llms in scalable and bias-aware academic reviews, 2024. doi:10.48550/arXiv.2408.10365.2
- [VTC21] VAN DINTER R., TEKINERDOGAN B., CATAL C.: Automation of systematic literature reviews: A systematic literature review. *IST 136* (2021). doi:10.1016/j.infsof.2021.106589. 1, 2
- [WFH\*23] WHITE J., FU Q., HAYS S., SANDBORN M., OLEA C., ET AL.: A prompt pattern catalog to enhance prompt engineering with chatgpt, 2023. doi:10.48550/arXiv.2302.11382.5
- [WH23] WHITFIELD S., HOFMANN M. A.: Elicit: Ai literature review research assistant. *Public Services Quarterly 19*, 3 (2023), 201–207. doi:10.1080/15228959.2023.2224125.2
- [Woh14] WOHLIN C.: Guidelines for snowballing in systematic literature studies and a replication in software engineering. In *Proc. EASE* (2014), ACM. doi:10.1145/2601248.2601268.4,5
- [WTL\*10] WALLACE B. C., TRIKALINOS T. A., LAU J., BRODLEY C., SCHMID C. H.: Semi-automated screening of biomedical citations for systematic reviews. *BMC bioinformatics 11* (2010), 1–11. doi: 10.1186/1471-2105-11-55.2

© 2025 The Authors. Proceedings published by Eurographics - The European Association for Computer Graphics.