Proceedings of the 11th World Congress on New Technologies (NewTech'25)
Paris, France - August, 2025
Paper No. ICBB 161
DOI: 10.11159/icbb25.161

Paperparser: A Bioinformatic Tool That Synthesizes Scientific Literature through Advanced AI Techniques, Enhancing Scholarly Insights While Mimicking Human-Like Expression

Alberto Corradin[1], Francesco Ciccarese[1], Vittoria Raimondi[1], Loredana Urso[1], Micol Silic-Benussi[1]
[1]Veneto Institute of Oncology IOV – IRCCS
Via Gattamelata 64, Padova, Italy
alberto.corradin@iov.veneto.it

Abstract **-** In this output-obsessed era, as scientific publishing continues its exponential growth, we've unthinkingly created a paradox: the cumulative advancement of knowledge by humanity now threatens to drive us all into information overload. Into the breach steps, PaperParser, a bioinformatic conductor crafted to gap this mental chasm. At a high level, PaperParser leverages Scrapy's raw extraction strength to crawl PubMed articles, combines this rich knowledge stream with the abstractive summarization ability of the SciLitLLM-14B. It's in the tool's layered construction, however—combining RAG methodologies, an autonomous Read-Eval-Print Loop (REPL) agent and Large Language Models—that the system's complexity truly emerges.
By following users' queries fluently, PaperParser keeps dynamically grouping cohorts of relevant articles that come out from the huge number of articles of the PubMed repository. But the path from data to insight is anything but a straight one: initial corpora is refined through FAISS indexing for fast, relevance-weighted passage retrieval. Drafts are shaped with GPT-40, powered by RAG retrievers that lash generative capability to the evidentiary backbone of a bespoke vector emporium. Further growth of these drafts requires nothing more than an agile AI agent given the reins of GPT-4.1 inside of a Python REPL skeleton. This agile editorial loop yields semantic accuracy as well as emergent clarity. And then, the 'humanization' phase. Here, OpenAI's o3-mini-2025-01-31 is deftly tuned using a Prompt Engineering interface by invoking stylistic signal harvested from target author exemplars through LangChain. PaperParser does not speak as a pure automaton, but as an intermediary, reconciling machine processing speed with a thoughtful, human-touched expression. It allows academics to go beyond mindless reading, shifting their focus to innovative interpretation and critique. We hope to extend this to paywalled periodicals and to perform source relevance-weighting—a further step towards a more discerning and context-sensitive literature navigator. In conclusion: PaperParser represents a new form of scientific intelligence enhancement.

Keywords**:** Web scraping, public databases, Large Language Models, RAG, AI-agents, REPL, Prompt Engineering

1. Introduction

Scientific literature, which is continuously growing exponentially, has begun to be an engine of progress while serving as the main obstacle to it. Today's researchers are overwhelmed with the nonstop influx of new research every piece promising to shed light on some mystery or problem we face, but taken together with so much of the rest of it that we cannot move. The archetypal literature review now seems almost Sisyphean in the domain of this exponential pile. Manual synthesis is not just inefficient; it has already in many areas become practically infeasible.

Enter PaperParser: a game-changing tool that will revolutionize the way you discover and digest papers. In conjoining the arachnid-like strength of Scrapy's mechanics [1-3] with the cognitive finesse native to the state-of-the-art language model (SciLitLLM-14B [4]), PaperParser orchestrates a world in which the relevant scientific texts that it encounters do not merely collect, — they distill and regiment themselves. The suggestion is transformative: now freed from the cognitive mechanics of data assembly, scholars may instead focus on interfacing directly with synthesized knowledge.

At the heart of this evolution is the spawning impact of Large Language Models (LLMs) [5-7]. PaperParser leverages these developments by integrating techniques of Retrieval-Augmented Generation (RAG) [8], AI-enabled REPLs [9], and the subtle high-dimensional matching provided by FAISS-powered vector similarity search [10]. The system's combination of OpenAI's technologies [11] with clever Prompt Engineering [12] turns independent summaries into architected academic drafts. In other words, PaperParser is not a passive channel but rather an active co-driver: "interpretating", structuring, and taking literature review away from being a tiresome burden to become an academic asset.

2. Methods

At its heart, both conceptually and technically, PaperParser is powered by the Scrapy framework—an industry standard, open source platform for controlling web crawls and data harvesting, primarily written in the Python programming language. Scrapy's built-in support for asynchronous HTTP provides out tool with the power to handle requests concurrently. This makes it uniquely predisposed to systematically collecting large bibliographic corpuses found across the repositories of the rigor and complexity of PubMed. In fact, PaperParser relies on a custom spider that performs fine-grained extraction of metadata, full text entities, and abstracts.

When a user query is submitted, PaperParser launches a carefully choreographed assault against PubMed, pulling out articles. Harvested metadata includes canonical fields—title, author cohorts, journal sovereigns, year of imprint, and, importantly, abstracts. Where unmediated full-text access is provided for, the pipeline switches to XML parsing, leveraging R tools such as the `tidypmc` package [13] to achieve high-fidelity recovery of scientific prose. The 'rpy2' technology grants the necessary bridge through distinct programming languages [14]: it allows calling the R packages needed to interrogate the EuroPMC database [15]. Knowing that it's a never-ending arms race with anti-bot measures, PaperParser strategically adds some delays and also cycles user agent headers. This organic navigation simulation is non-trivial, sustaining access to primary sources: it avoids PaperParser's queries being interpreted as illegitimate (informatic) attacks to public servers.

Upon retrieval, the contents are subject to a summarization schema. TextRank, in the form of Sumy [16], performs extractive condensation, retrieving relevant sentences to construct a summarized view of the underlying debate. This to provide a preliminary overview of collected data to the user and a subsequent human check of Natural Language Processing (NLP) [17-19]. Abstractive summarization is applied invoking transformer architectures, i.e. Google T5, via Hugging Face's powerful `transformers` toolkit [20]. And, later, by the fine-tuned SciLitLLM-14B, whose output summaries are not mere reductions, but syntheses, that seek to maintain conceptual coherence and fine-grained context. Then, the informative content of abstracts and AI-generated summaries is organized through FAISS (Facebook AI Similarity Search), transformed into vectorized high dimensional representations with the HuggingFaceEmbeddings's model. FAISS isn't just for storage: it also provides semantically aware and fast retrieval via vector similarity search, which backs PaperParser's RAG retriever. The generation of the draft proceeds through multiple phases: an initial framework is generated thanks to GPT-4o-mini by retrieving information from the FAISS-derived semantic kernels. Every paragraph is extended to overcome limitations in the number of ouput tokens. A GPT-4.1-fueled AI-agent guides a Python REPL to identify paragraphs' titles and contents. These latter are analyzed and enhanced through a LangChain application [21] that leverages GPT-4o and RAG. The final manuscript also goes through one last dramatic reenactment by OpenAI's o3-mini-2025-01-31. State-of-the-art Prompt Engineering guides the rework of the draft by inserting the target author's previous written as a LangChain Runnable. This affects the narrative with a characteristic scientific rhythm, adjusting the register and pace of the text to the authorial guide prints [22].

3. Results

PaperParser acts as an enabler in scientific workflows by significantly reducing the temporal and cognitive costs incurred in the process of thorough literature retrieval and distillation. By automating extraction and summarization pipelines, the system helps researchers to take/generate panoramic though specific snapshots of massive scientific territories/questions in a fraction of the time they would normally need for such an undertaking. PaperParser can represent an efficient time-saving solution. Of central importance is the incorporation of SciLitLLM-14B, a fine-tuned LLM specifically developed to deal with scientific literature, whose summarization capability is detailed in Li et al. (2024). This was supported by a FAISS-powered high-dimensional vector similarity storage for data organization and retrieval.

Despite the progress, PaperParser is not immune to restriction. Currently, its "horizon" for literature retrieval is limited by what's available in "open access" repositories, leaving the vast collections of paywall journals dangling out of reach. Future versions will incorporate broader approaches to data harvesting—such as the slicing of paywalls through institutional access—and relevance vectors to include citation analytics, impact scores, and fine-grained user preference. These improvements will bring not only technical enhancements but also a philosophically richer approach to how machines can surface, sift, and synthesize the constantly expanding corpus of scientific knowledge.

The flow chart in Figure 1 summarizes the web scraping phase whereas Figure 2 describes the sequence of operations leading to text generation and refinement. Refer to [23-27] for details concerning the technologies reported in the figures. The core Scrapy spider was coded with Python scripts leveraging object-oriented programming. It can be run from Linux shell or from Jupyter notebooks (through system call). Connection to the Hugging Face repository shall be granted. In particular, user's identification is necessary to download Facebook's FAISS. The paywalled subscription to OpenAI is mandatory. Login/user name and password data are usually passed as environment variables. A cloud platform can support the launch of the text generation and refinement phase. For PaperParser's development, Google Colab provided the interface to Jupyter notebooks whereas A100 Nvidia GPU revealed necessary only to run SciLitLLM-14B. The drafting and humanizing steps were implemented with a second series of Python scripts, which were stored in Google Drive. All the scripts are publicly available in the author's Github web page: https://github.com/AlbertoCorradinPhD/PaperParser.

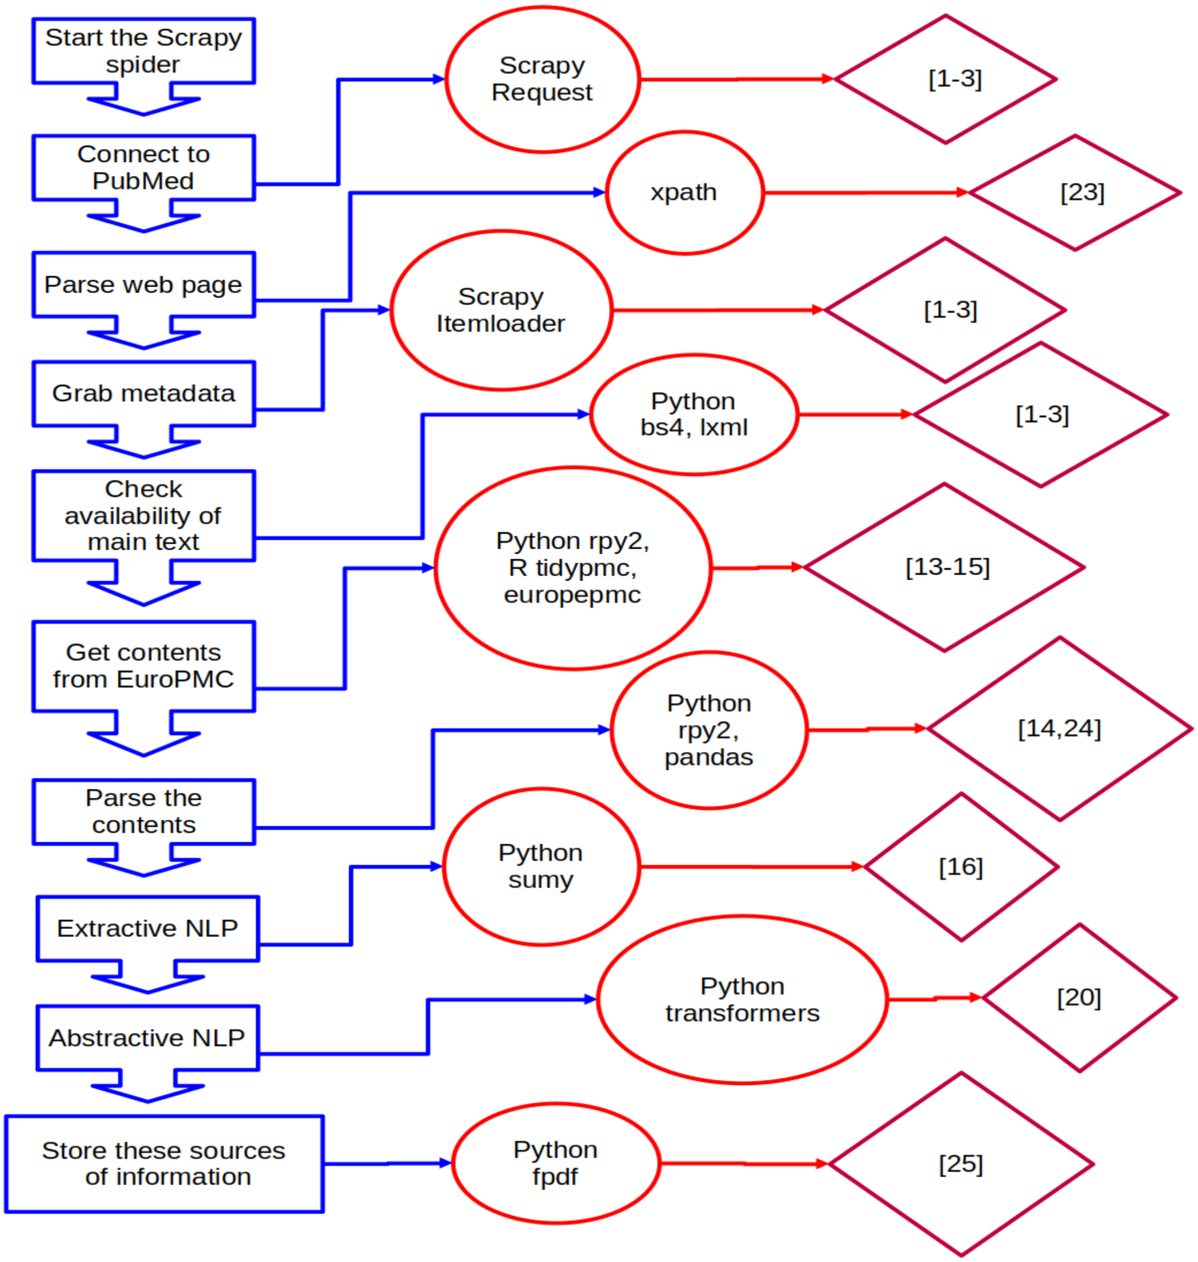

Fig. 1: Workflow of the Web Scraping phase. Python modules and R packages are listed inside red circles. Corresponding references are cited on the right side (inside violet diamonds)

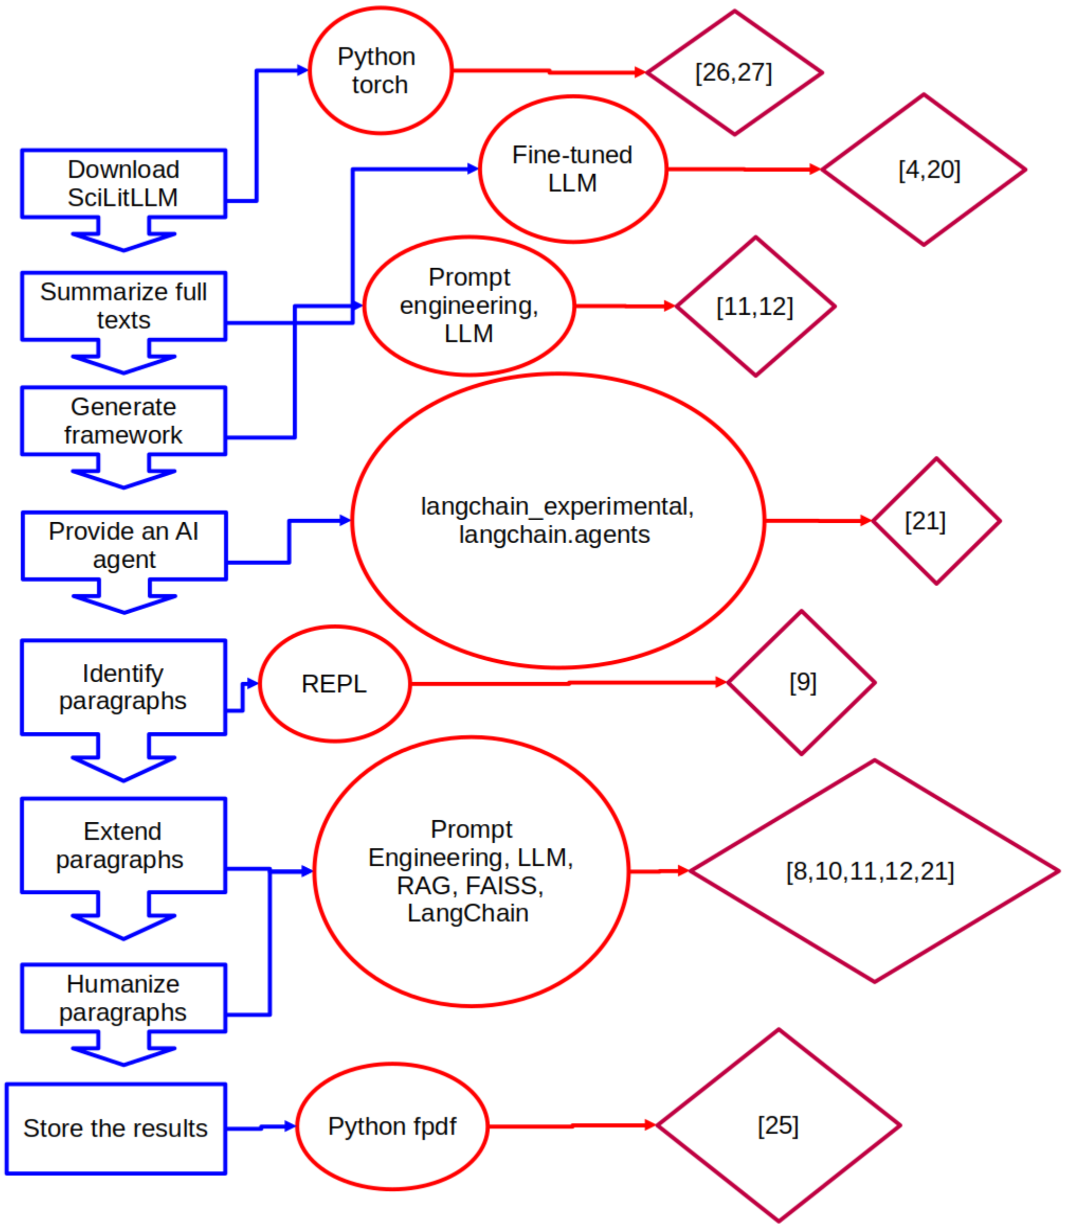

Fig. 2: Flow chart for text generation and refinement's operations. Python modules or other technologies are listed inside red circles. Corresponding references are cited on the right side (inside violet diamonds)

4. Discussion

PaperParser represents a potential game-changer to a crisis of biblical proportions when it comes to too many scientific information to manage. Meaning and ethics of LLMs and PaperParser validation concerns, reproducibility and ethics are obviously legitimate cards on the table, but we need to foreground how much potentials there are in LLMs such as ones used in the PaperParser. The future of medical and scientific communication may rely on artificial intelligence, which could radically reshape our scholarly, collaborative and discovery traditions. What LLMs are giving us is not merely a mechanical condensation of the text, but a partner in thinking, an extension of our cognitive capital, one that addresses the tension between exponential data growth and finite human cognitive bandwidth. The volume of published biomedical and life sciences literature is now so large that no team, let alone single person, can keep up with even a fraction of what is pertinent. In this context LLMs are (or could be) ideally situated to act as sentinels and mediators, tirelessly sifting through, digesting, and cross-referencing hundreds — ultimately thousands — of newly published articles in real time, identifying the most relevant findings, controversial emerging topics, and relationships that would otherwise be impossible for even the most dedicated professional to access. The PaperParser pipeline, which passes from raw scraping to semantic vectorization to generative drafting and stylistic adaptation, is what cutting-edge AI looks like when it's put to use as a finely tuned assistant rather than a wrecking ball. Unlike with extractive methods, abstractive summarization—especially when performed by models tuned to specific scientific domains (e.g., SciLitLLM-14B)—allows for a level of conceptual distillation and reframing that approaches or exceeds what humans can do. LLMs can computationally identify repeating patterns, collectively piece apart independent lines of evidence, and even detect nuanced methodological patterns, inconsistencies or deficits in the literature. The ability to "connect the dots" in ways that were not previously apparent (infer from studies to other studies or to subfields) may constitute a "force multiplier" for hypothesis generation and meta-analytic thinking.

For clinical medicine, the benefits of employing machine-learning models such as those demonstrated in PaperParser are not just hypothetical. Systematic reviews, evidence syntheses, and meta-analyses are the touchstone for translating science into practice, yet they are logistically and cognitively-intensive, requiring many months, if not years to conduct. An LLM-empowered tool could significantly intensify such cycles—by identifying and extracting pertinent studies, synthesizing early conclusions, and indicating key references or controversies to be further examined by human teams. The process of "humanization" is not merely a wrapper around the system output but is itself highly valuable: it customizes outputs to the expectations, culture, and rhetorical pace of the target author. So that they can be something more than cold technical reports: they can be readable, accessible, actionable scholarly products.

Crucially, as the transparency and traceability of LLM architectures improve in the hands of AI researchers, so will the opportunities for dynamic interplay between human inquiry and machine assistance. PaperParser's agent-story pipeline, and the REPL-like workflow it undergirds, navigates a future where researchers could iterate over how to query, guide, and elicit refinements from LLM-generated syntheses while maintaining ownership of interpretations, and leveraging never-before-available scale and speed. This "centaur model" — AI as coach and sparring partner, not rival — opens up rich opportunities for collaborative workflows, collective authorship and ongoing learning between humans and machines.

Although promising and ambitious in its ability to solve the problem of information overload in modern scientific research, the current PaperParser solution also raises several important issues and questions that need to be resolved and answered to realize its full potential as a radical and responsible intervention. First and foremost, we are in need of thorough validation and benchmarking of the AI-produced summaries: although the PaperParser implements state-of-the-art abstractive summarization pipelines—we employ models like SciLitLLM-14B or GPT-based models—its vulnerability to fabricate partial truths or even nonsense by accomplishing high coherence must be counterbalanced by rigorous systems for fact-checking, cross-reference and empirical validation against human generated summaries [28-30]. Secondly, the reliance of the model on proprietary tools, which some are not available or released- like OpenAI's o3-mini-2025-01-31 and GPT-4.1, it calls into question reproducibility and transparency issues. This dependence can prevent independent validation and adoption, particularly within less well-resourced institutions, thereby emphasizing the need for documentation of the tool dependencies and licensing, as well as where possible, their interoperability with well-established open source alternatives.

A further material limitation is the coverage of PaperParser is currently limited to open-access repositories. This means that a lot of critical knowledge – typically found in important, paywalled publications – may therefore be left out of the synthesized knowledge and may critically limit the range of accessible literature that can be legitimately taken into account. Fixing this will not be without technical solutions, including through institutions of access, for legally and ethically managing licensed content, and through collaborative relationships between publishers and academic institutions, as always within strict confines of national and international laws and data privacy regulations (GDPR, etc.). Second, from a user perspective, AI and NLP technologies are developing fast, thus a modular architecture allowing constant updates, continuous maintenance and community contributions are crucial for preventing PaperParser from becoming outdated.

5. Conclusion

PaperParser, a modular cross between automatic literature retrieval and life science summarization, based on the Scrapy framework, widely recognized for its flexible and scalable scrapping engine. This takes care of the smooth acquisition of various scholarly articles, and is combined with SciLitLLM-14B for fine-tuned distillation of the text, producing summaries which can endeavor to be both comprehensive and coherent. These extracted insights are then embedded and indexed in

FAISS whose fast vector intersection computations allow to return contextually relevant items of meaning quickly over large corpora. Crucially, the model uses GPT-4o as a first draft-smith, before a process described as "humanization" where OpenAI's o3-mini-2025-01-31 edits the generated draft to resemble the tone of a science communicator.

Going forward, improvements will focus on increasing the scope of data acquisition—so we will be targeting the licensed repositories and subscription-limited journals that are typically not covered by the automated flows. Furthermore, the integration of machine-relevance algorithms has the potential to prune outputs even more, directing our gaze towards literature that can actually make a difference. Conversely, AI models, including complex architectures like LLMs, can remarkably produce elaborate but incorrect narratives— called "hallucinations". Though the fluency and coherence of these outputs can be impressive, they are also extremely vulnerable to the propagation of misinformation. In academic discussion, the cost of unrestrained hallucination is the corrosion of sound scholarship; in medicine, the stakes are existential ones. Such vulnerabilities needs to be well managed in order to be contained [31]. Without this, the gap between lab-based wisdom and practicing reality could dangerously widen.

Medicine has always been at the junction of disciplines; AI requires an unprecedented level of collaborative will. The development of research collaborations such as that involving informatics, clinical expertise, ethics, and policy will be critical to dealing with the complexity induced by the adoption of AI [32]. In fact, the integration of artificial intelligence into the living body of medical writing is nothing short of revolutionary — a testament to technology's inexorable advance across the intellectual sphere. AI-powered applications, especially when based on advanced NLP, have dramatically reshaped the extremes of medical documentation. Yet, human supervision remains irreplaceable to ensure content quality [33]. There is the threat of factual distortions, subtle biases, and, perhaps most potently of all, a diminution of uniqueness [34,35]. Ethical dilemmas proliferate.

The capability of AI in expediting systematic review and meta-analysis research facilities also points towards this double-edged potential. AI-intensive literature mining and synthesis collapse scales from months to weeks, if not days [36]. In fields afflicted by exponential growth of knowledge, such computational agility is absolutely crucial. Yet the very powers of the new AI bring into sharp relief deep problems with equal force. The web of medical information is too delicate to be left to the churning of the machines; the specter of error, unopposed, could further spread fallacious rumors with alarming efficiency [37]. AI, if wisely employed, has the potential to raise the bar on the accuracy, breadth, and reach of medical writing. Despite this promise is only part of a bargain, the opportunity '*technology as a partner, not a replacement*' justifies the ongoing dialogue across developers, clinicians, and ethicists, centered on a culture of rigor and transparency.

Availability of software code. Our software code is publicly available at the following web URL: https://github.com/AlbertoCorradinPhD/PaperParser .

References

[1] G. L. Hajba and Springerlink (Online Service), *Website Scraping with Python : Using BeautifulSoup and Scrapy*. Berkeley, Ca: Apress, 2018

[2] D. Kouzis-Loukas, *Learning Scrapy : learn the art of efficient web scraping and crawling with Python.* Birmingham, Uk: Packt Publishing, 2016

[3] R. E. Mitchell, *Web scraping with Python : collecting more data from the modern web*. Sebastopol, Ca: O'reilly Media, 2018

[4] S. Li, J. Huang , J. Zhuang, Y. Shi, X. Cai, M. Xu, X. Wang, L. Zhang, G. Ke, H. Cai, "SciLitLLM: How to Adapt LLMs for Scientific Literature Understanding," *arXiv.org*, 2024, https://arxiv.org/abs/2408.15545

[5] R. Bommasani, D. A. Hudson, E. Adeli, R. Altman, S. Arora, S. von Arx, M. S. Bernstein, J. Bohg, A. Bosselut, E. Brunskill, E. Brynjolfsson, S. Buch, D. Card, R. Castellon, N. Chatterji, A. Chen, K. Creel, J. Quincy Davis, D. Demszky, C. Donahue, M. Doumbouya, E. Durmus, S. Ermon, J. Etchemendy, K. Ethayarajh, L. Fei-Fei, C. Finn, T. Gale, L. Gillespie, K. Goel, N. Goodman, S. Grossman, N. Guha, T. Hashimoto, P. Henderson, J. Hewitt, D. E. Ho, J. Hong, K. Hsu, J. Huang, T. Icard, S. Jain, D. Jurafsky, P. Kalluri, S. Karamcheti, G. Keeling, F. Khani, O. Khattab, P. Wei Koh, M. Krass, R. Krishna, R. Kuditipudi, A. Kumar, F. Ladhak, M. Lee, T. Lee, J. Leskovec, I. Levent, X. L. Li, X. Li, T. Ma, A. Malik, C. D. Manning, S. Mirchandani, E. Mitchell, Z. Munyikwa, S. Nair, A. Narayan, D. Narayanan, B. Newman, A. Nie, J. C. Niebles, H. Nilforoshan, J. Nyarko, G. Ogut, L. Orr, I. Papadimitriou, J. S. Park, C. Piech, E. Portelance, C. Potts, A. Raghunathan, R. Reich, H. Ren, F. Rong, Y. Roohani, C. Ruiz, J. Ryan, C. Ré, D. Sadigh, S. Sagawa, K. Santhanam, A. Shih, K. Srinivasan, A. Tamkin, R. Taori, A. W. Thomas, F. Tramèr, R. E. Wang, W. Wang

, B. Wu, J. Wu, Y. Wu, S. M. Xie, M. Yasunaga, J. You, M. Zaharia, M. Zhang, T. Zhang, X. Zhang, Y. Zhang, L. Zheng, K. Zhou, P. Liang, "On the Opportunities and Risks of Foundation Models," *arXiv:2108.07258 [cs]* , vol. 1, no. 1, Aug. 2021, https://arxiv.org/abs/2108.07258

[6] V. Alto, *Building LLM Powered Applications*. Packt Publishing Ltd, 2024

[7] N. A. Khan, K. Osmonaliev, and M. Z. Sarwar, "Pushing the Boundaries of Scientific Research with the use of Artificial Intelligence tools: Navigating Risks and Unleashing Possibilities," *Nepal J Epidemiol.* , vol. 13, no. 1, pp. 1258–1263, Mar. 2023, doi: https://doi.org/10.3126/nje.v13i1.53721

[8] Y. Gao, Y. Xiong, X. Gao, K. Jia, J. Pan, Y. Bi, Y. Dai, J. Sun, M. Wang, H. Wang, "Retrieval-Augmented Generation for Large Language Models: A Survey," *arXiv.org*, Dec. 18, 2023. https://arxiv.org/abs/2312.10997

[9] K. Ellis, M. Nye, Y. Pu, F. Sosa, J. Tenenbaum, A. Solar-Lezama, "Write, Execute, Assess: Program Synthesis with a REPL," *arXiv.org*, 2019. https://arxiv.org/abs/1906.04604

[10] M. Douze, A. Guzhva, C. Deng, J. Johnson, G. Szilvasy, P. Mazaré, M. Lomeli, L. Hosseini, H. Jégou, "The Faiss library," *arXiv.org*, Jan. 16, 2024. https://arxiv.org/abs/2401.08281

[11] OpenAI,"GPT-4 Technical Report," *arXiv (Cornell University)*, Mar. 2023, https://doi.org/10.48550/arxiv.2303.08774

[12] D. Genkina, "AI Prompt Engineering Is Dead - IEEE Spectrum," *spectrum.ieee.org*, Mar. 06, 2024. https://spectrum.ieee.org/prompt-engineering-is-dead

[13] C. Stubben (2024). "tidypmc: Parse Full Text XML Documents from PubMed Central". R package version 2.0

[14] "rpy2: Python-R bridge", https://rpy2.github.io/

[15] N. Jahn (2023). "europepmc: R Interface to the Europe PubMed Central RESTful Web Service". R package version 0.4.3.

[16] M. Belica, "sumy: Module for automatic summarization of text documents and HTML pages.", *PyPI*. https://pypi.org/project/sumy/

[17] I. Atanassova, M. Bertin, and P. Mayr, "Editorial: Mining Scientific Papers: NLP-enhanced Bibliometrics," *Frontiers in Research Metrics and Analytics*, vol. 4, Apr. 2019, doi: https://doi.org/10.3389/frma.2019.00002

[18] Z. Antić, *Python natural language processing cookbook : over 50 recipes to understand, analyze, and generate text for implementing language processing tasks*. Birmingham Packt. Mar. 2021

[19] V. M. Torres-Lopez, G. E. Rovenolt, A. J. Olcese, G. E. Garcia, S. M. Chacko, A. Robinson, E. Gaiser, J. Acosta, A. L. Herman, L. R. Kuohn, M. Leary, A. L Soto, Q. Zhang, S. Fatima, G. J. Falcone, S. Payabvash, R. Sharma, A. F. Struck, K. N. Sheth, M. B. Westover, J. A. Kim, "Development and Validation of a Model to Identify Critical Brain Injuries Using Natural Language Processing of Text Computed Tomography Reports," *JAMA Network Open*, vol. 5, no. 8, pp. e2227109–e2227109, Aug. 2022, doi: https://doi.org/10.1001/jamanetworkopen.2022.27109

[20] T. Wolf, L. Debut, V. Sanh, J. Chaumond, C. Delangue, A. Moi, P. Cistac, T. Rault, R. Louf, M. Funtowicz, J. Davison, S. Shleifer, P. von Platen, C. Ma, Y. Jernite, J. Plu, C. Xu, T. Le Scao, S. Gugger, M. Drame, Q. Lhoest, A.r M. Rush, "HuggingFace's Transformers: State-of-the-art Natural Language Processing," *arXiv:1910.03771*, Feb. 2020, https://arxiv.org/abs/1910.03771

[21] H. Chase, "LangChain," *GitHub*, Oct. 01, 2022. https://github.com/langchain-ai/langchain

[22] F. R. Ward, Z. Yang, A. Jackson, R. Brown, C. Smith, G. Colverd, L. Thomson, R. Douglas, P. Bartak, A. Rowan, "Evaluating Language Model Character Traits," *arXiv.org*, 2024. https://arxiv.org/abs/2410.04272

[23] "XML Path Language (XPath) 3.1," *www.w3.org*. https://www.w3.org/TR/xpath-31/

[24] "Pandas: Python Data Analysis Library," *Pydata.org*, 2018. https://pandas.pydata.org/

[25] M. Reingart, "fpdf: Simple PDF generation for Python," *PyPI*. https://pypi.org/project/fpdf/

[26] PyTorch Team, "torch: Tensors and Dynamic neural networks in Python with strong GPU acceleration," *PyPI*. https://pypi.org/project/torch/

[27] Hugging Face, "Hugging Face – On a mission to solve NLP, one commit at a time.," *huggingface.co*, 2024. https://huggingface.co/

[28] S. Bedi,Y. Liu, L. Orr-Ewing, D. Dash, S. Koyejo, A. Callahan, J. A. Fries, M. Wornow, A. Swaminathan, L. S. Lehmann, H. J. Hong, M. Kashyap, A. R. Chaurasia, N. R. Shah, K. Singh, T. Tazbaz, A. Milstein, M. A. Pfeffer, N. H. Shah, "Testing and Evaluation of Health Care Applications of Large Language Models," *JAMA*, Oct. 2024, https://doi.org/10.1001/jama.2024.21700

[29] M. Khosravi, Z. Zare, S. M. Mojtabaeian, R. Izadi, “Artificial intelligence and decision-making in healthcare: A thematic analysis of a systematic review of Reviews,” *Health Services Research and Managerial Epidemiology*, vol. 11, no. 23333928241234863, 2024, https://doi.org/10.1177/23333928241234863
[30] O. Temsah, S. A. Khan, Y. Chaiah, A. Senjab, K. Alhasan, A. Jamal, F. Aljamaan, K. H. Malki, R. Halwani, J. A. Al-Tawfiq, M. Temsah, A. Al-Eyadhy, “Overview of Early ChatGPT’s Presence in Medical Literature: Insights From a Hybrid Literature Review by ChatGPT and Human Experts,” *Cureus*, vol. 15, no. 4, Apr. 2023, https://doi.org/10.7759/cureus.37281
[31] P. Goktas and A. Grzybowski, “Assessing the Impact of ChatGPT in Dermatology: A Comprehensive Rapid Review,” *Journal of Clinical Medicine*, vol. 13, no. 19, p. 5909, Oct. 2024, https://doi.org/10.3390/jcm13195909
[32] P. Zangrossi, M. Martini, F. Guerrini, P. De-Bonis, G. Spena, “Large language model, AI and scientific research: why ChatGPT is only the beginning,” *Journal of neurosurgical sciences*, Jan. 2024, https://doi.org/10.23736/s0390-5616.23.06171-4
[33] P. Thaichana, M. Z. Oo, G. L. Thorup, C. Chansakaow, S. Arworn, K. Rerkasem, “Integrating Artificial Intelligence in Medical Writing: Balancing Technological Innovation and Human Expertise, with Practical Applications in Lower Extremity Wounds Care,” The International Journal of Lower Extremity Wounds, Jan. 2025, https://doi.org/10.1177/15347346241312814
[34] M. C. Y. Koh, J. N. Ngiam, J. E. L. Oon, L. H.-W. Lum, N. Smitasin, S. Archuleta, “Using ChatGPT for writing hospital inpatient discharge summaries – perspectives from an inpatient infectious diseases service,” *BMC Health Services Research*, vol. 25, no. 1, Feb. 2025, doi: https://doi.org/10.1186/s12913-025-12373-w
[35] A. S. Teoman and E. C. Serefoglu, “Artificial Intelligence-Based Clinical Decision-Making in Erectile Dysfunction: a Narrative Review,” *Current urology reports*, vol. 26, no. 1, p. 22, Nov. 2024, https://doi.org/10.1007/s11934-024-01251-3
[36] B. Fatani, “ChatGPT for Future Medical and Dental Research,” Cureus, Apr. 2023, https://doi.org/10.7759/cureus.37285
[37] H. Mondal, S. Mondal, and A. Juhi, “Responsible Use of Generative Artificial Intelligence for Research and Writing: Summarizing ICMJE Guideline,” *Indian Journal of Orthopaedics*, vol. 58, no. 10, pp. 1504–1505, Aug. 2024, https://doi.org/10.1007/s43465-024-01258-5