

OpenResearcher: Unleashing AI for Accelerated Scientific Research

Yuxiang Zheng^{1,8*} Shichao Sun^{4,8*} Lin Qiu¹ Dongyu Ru¹ Cheng Jiayang⁵ Xuefeng Li^{1,8}
Jifan Lin^{1,8} Binjie Wang^{3,8} Yun Luo⁶ Renjie Pan¹ Yang Xu¹ Qingkai Min⁶
Zizhao Zhang⁷ Yiwen Wang¹ Wenjie Li⁴ Pengfei Liu^{1,2,8†}

¹Shanghai Jiao Tong University ²Shanghai Artificial Intelligence Laboratory ³Fudan University
⁴The Hong Kong Polytechnic University ⁵Hong Kong University of Science and Technology
⁶Westlake University ⁷Tsinghua University ⁸Generative AI Research Lab (GAIR)
catchiz.1@sjtu.edu.cn, pengfei@sjtu.edu.cn

Abstract

The rapid growth of scientific literature imposes significant challenges for researchers endeavoring to stay updated with the latest advancements in their fields and delve into new areas. We introduce OpenResearcher, an innovative platform that leverages Artificial Intelligence (AI) techniques to accelerate the research process by answering diverse questions from researchers. OpenResearcher is built based on Retrieval-Augmented Generation (RAG) to integrate Large Language Models (LLMs) with up-to-date, domain-specific knowledge. Moreover, we develop various tools for OpenResearcher to understand researchers' queries, search from the scientific literature, filter retrieved information, provide accurate and comprehensive answers, and self-refine these answers. OpenResearcher can flexibly use these tools to balance efficiency and effectiveness. As a result, OpenResearcher enables researchers to save time and increase their potential to discover new insights and drive scientific breakthroughs. Demo, video, and code are available at: <https://github.com/GAIR-NLP/OpenResearcher>.

1 Introduction

Global scientific publications are growing annually by about 4%-5% (Pinedo et al., 2024), leading researchers to invest significant time and effort in thoroughly reviewing countless academic papers to find the knowledge that propels their research. This involves daily engagement with a wide range of literature to stay updated with the latest developments in their field, which is essential for maintaining the relevance and innovation of their work.

Recognizing the challenges and inefficiencies inherent in this process, considerable academic efforts have focused on AI-assisted scientific research

(Wang et al., 2023a; Zhai, 2023). They aim to answer the researcher questions from both junior and senior researchers. These questions can be broadly classified into three categories: (1) Scientific Question Answering (Pappas et al., 2020; Ruggeri et al., 2023; Lee et al., 2023; Pramanick et al., 2024), which seeks detailed information or clarification within a specific domain; (2) Scientific Text Summarization (Wang et al., 2022; Ding et al., 2023; Takeshita et al., 2024; Hsu et al., 2024; Zhang et al., 2024), aimed at condensing the latest findings and developments into comprehensive overviews; and (3) Scientific Paper Recommendation (Bai et al., 2019; Kreutz and Schenkel, 2022; Stergiopoulos et al., 2024; Pinedo et al., 2024), which involves suggesting relevant literature and studies based on the researcher's interests or current inquiries. However, academic applications typically focus on a **single** task, lacking a unified solution for all questions, allowing researchers to pose any inquiry freely.

Conversely, recent industry applications, like Perplexity AI,¹ iAsk,² You.com,³ phind,⁴ and SearchGPT,⁵ allow users to inquire about anything beyond specific tasks. They use Retrieval-Augmented Generation (RAG) (Lewis et al., 2020) technique to offer an innovative integration of generative Large Language Model (LLM) with web search capability. The core idea behind them is to offer users not just any answer, but the most accurate and contextually relevant information available. However, the **proprietary** nature of industry applications has hindered their development and may impede academic research in this field.

Besides, both academic and industry applications serve as **passive** assistants, focusing solely on responding to user inquiries rather than engaging

¹<https://www.perplexity.ai/>

²<https://iask.ai/>

³<https://you.com/>

⁴<https://www.phind.com/>

⁵<https://chatgpt.com/search>

*Equal contribution.

†Corresponding author.

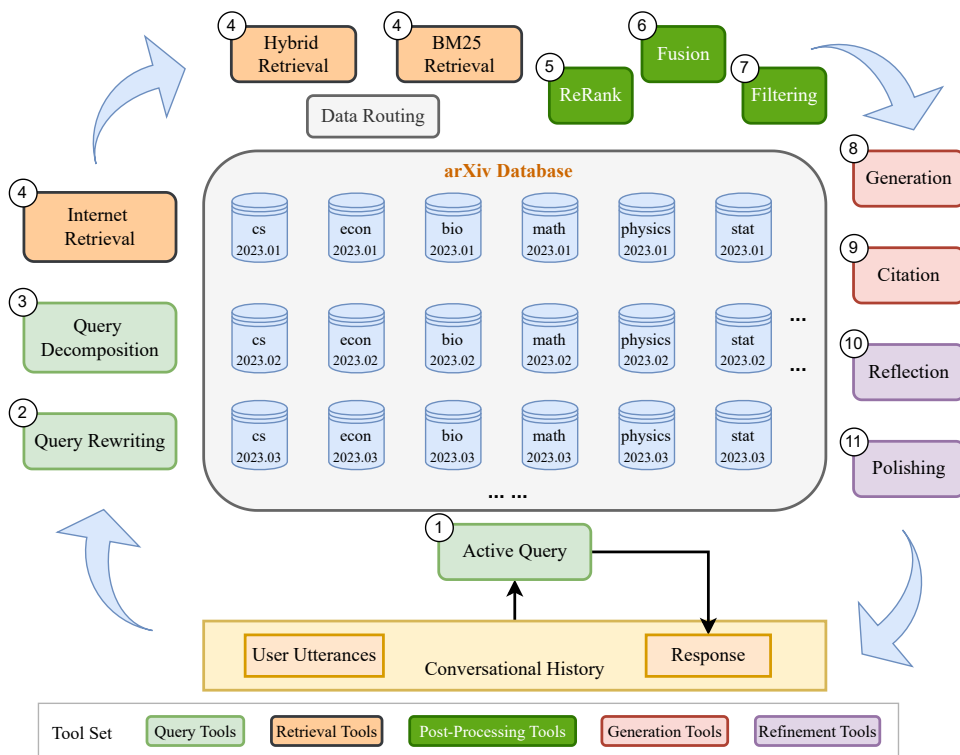


Figure 1: Main Workflow of OpenResearcher.

in active communication. To address these issues in academic and industry contexts, we developed OpenResearcher, an open-source project that harnesses AI to accelerate scientific research. Its main workflow is shown in Figure 1, and its main contributions are as follows:

- **Unified Application** OpenResearcher can address researchers’ diverse questions, such as Scientific Text Summarization, Scientific Paper Recommendation, etc.
- **Open-Source** OpenResearcher is an impressive open-source system to rival the performance of industry applications.
- **Active Assistant** OpenResearcher can connect in the mind or imagination to pose heuristic questions, guiding users to clarify queries for capturing their intent.
- **Retrieval Augmented** OpenResearcher can retrieve from the Internet and arXiv corpus to provide up-to-date, domain-specific, verified knowledge as supporting evidence.
- **Flexible Tool Usage** OpenResearcher can flexibly utilize bespoke tools to build a workflow for a better answer. For example, OpenResearcher adaptively calls a refinement tool

to refine its initial outcomes. This approach helps avoid the computational cost associated with the unnecessary use of some tools.

- **Conversational Interaction** OpenResearcher enables users to engage in deep discussions through conversational follow-up questions.

2 Related Work

2.1 Academic Works

Academic works for scientific research target a specific task, including Scientific Question Answering, Scientific Text Summarization, and Scientific Paper Recommendation.

Scientific Question Answering generates answers for questions within extensive scientific articles. In the early days, cloze-style paper question answering datasets, such as emrQA (Pampari et al., 2018), BioRead (Pappas et al., 2018) and BioMRC (Pappas et al., 2020), are automatically created with the pre-defined question formats (Kwiatkowski et al., 2019). On the other hand, PubMedQA (Jin et al., 2019), BioAsq (Krallinger et al., 2020) and QASPER (Dasigi et al., 2021) involve human annotators in question creation. However, the questions are based only on abstracts. Recently, QASA (Lee et al., 2023) offers advanced questions with annotators reading the entire paper. KIWI (Xu et al.,

2024) uses expert and LLM interactions to refine initial answers into improved long-form answers. SPIQA (Pramanick et al., 2024) expands text question answering to multimodal question answering. **Scientific Text Summarization** aims to condense the long scientific articles into a concise summary. Early works primarily focus on a knowledge graph-centric view (Wang et al., 2022). Recently, Ding et al. (2023) present CocoSciSum, a novel toolkit for controlled summarization of scientific documents, tailored to the scientific community’s needs. Takeshita et al. (2024) introduce ACLSum, an expert-curated dataset for multi-aspect summarization of scientific papers, thoroughly covering challenges, approaches, and outcomes. Hsu et al. (2024) release CHIME, a dataset that hierarchically organizes scientific studies to facilitate the generation of literature reviews. Zhang et al. (2024) introduces MASSW, a comprehensive dataset for summarizing multi-aspects of scientific workflows. **Scientific Paper Recommendation** assists researchers in discovering relevant and suitable scientific information through recommendations. Early approaches (Tanner et al., 2019; Ma and Wang, 2019; Sakib et al., 2020; Manju et al., 2020) in Big Scholarly Data (Khan et al., 2017) have evolved into recently proposed hybrid recommender systems. Pinedo et al. (2024) develop ArZiGo, a web-based prototype system for searching, managing, and recommending scientific articles. Stergiopoulos et al. (2024) present a novel multi-stage recommendation system employing clustering, graph modeling, and deep learning, capable of operating on a large-scale scientific digital library with millions of users and papers.

However, these academic efforts focus on a single function without a unified solution for diverse inquiries and lack a user-friendly web application.

2.2 Industry Research Applications

Recent advancements in LLMs have prompted the industry to explore AI assistants for scientific research, like Perplexity AI, iAsk, You.com, phind, and SearchGPT, designed to handle all kinds of research inquiries in a dialogue. These applications combine chatbot-driven search engines with LLMs, which is academically termed Retrieval Augmented Generation (RAG). These applications also provide citations for the evidence behind their responses. However, the closed-source nature has limited their development and academic research in this area.

3 OpenResearcher

OpenResearcher is designed to leverage AI to speed up the research process by efficiently responding to researchers’ inquiries. As shown in Figure 1, OpenResearcher employs RAG to combine LLMs’ internal knowledge with the latest external information. We design a Data Routing strategy for quick and precise information retrieval that can meet time and domain requirements. Lastly, we have developed multiple tools, including query tools, retrieval tools, post-processing tools, generation tools, and refinement tools. OpenResearcher can flexibly use these tools to customize a workflow for each query.

3.1 Query Tools

A key challenge in retrieval is its dependence on the user’s initial query, which, if imprecise or vague, leads to ineffective results. Junior researchers may struggle to articulate their questions, and scientific terms used across different disciplines add to this complexity. To address this, we have developed tools to help define straightforward questions.

Active Query OpenResearcher enhances a query by adding extra content and context. It asks users to specify their interest area or discipline. It can ensure that generated answers are highly relevant by covering nuances not initially mentioned.

Query Rewriting The users’ queries are always suboptimal for retrieval, especially in real-world scenarios. Besides, the queries are commonly entailed in complex conversational interactions. Therefore, OpenResearcher rewrites the queries for better clarity and effectiveness.

Query Decomposition OpenResearcher decomposes the complex query into a series of sub-queries, improving precision and efficiency for more satisfying responses. Then each sub-query is processed by information retrieval and LLM generation systems accordingly to get the sub-answer.

3.2 Retrieval Tools

OpenResearcher uses advanced retrieval tools to gather comprehensive and accurate information from the Internet and arXiv corpus.

Internet Retrieval OpenResearcher conducts Internet Retrieval through search engines API to collect relevant online information.

Hybrid Retrieval OpenResearcher supports Hybrid Retrieval that employs sparse vector and dense vector representations of both queries and docu-

ments. By leveraging these compact vector embeddings, Hybrid Retrieval can more effectively capture semantic similarities and improve the relevance of retrieved documents.

BM25 Retrieval OpenResearcher conducts BM25 Retrieval, an advanced algorithm used by search engines to rank documents based on their relevance to a query, factoring in term frequency and document length. BM25 stands out for its effectiveness in handling various search queries, making it a widely adopted method in information retrieval.

3.3 Data Routing Strategy

We develop an advanced Data Routing strategy aimed at optimizing the performance of our hybrid retrieval tool. This retrieval tool currently requires substantial processing times to calculate the similarity between a query and all arXiv paper chunks, which can be resource-intensive.

To address this issue, our strategy is to stratify the data based on both temporal and domain-specific information found in the metadata of the arXiv papers. It distributes data across multiple specialized databases, each aligned with a particular time frame and domain. Consequently, the retrieval tool only scans databases relevant to the query, which speeds up the search process and improves result accuracy by concentrating on the applicable data sets.

3.4 Post-Processing Tools

We develop Post-Processing Tools to rerank, fuse, and filter retrieved information, removing noise and redundancy to provide the most pertinent outcomes for the generation of LLMs.

Reranking: OpenResearcher can use a reranking tool to reorder document chunks, prioritizing the most relevant results to condense the retrieval pool.

Fusion: OpenResearcher can use a fusion tool to fuse the retrieved content from the same source into a single paragraph to enhance the context.

Filtering: OpenResearcher can use a filtering tool to filter out redundant and noisy content to preserve the most relevant information.

3.5 Generation Tools

OpenResearcher uses advanced LLMs to produce responses using retrieved information.

Generation OpenResearcher prompts LLMs to utilize retrieved information to generate appropriate responses for user queries.

Citation OpenResearcher can use a citation tool that employs the BM25 matching algorithm to link retrieved information with the response sentences, providing citations for each.

3.6 Refinement Tools

OpenResearcher utilizes LLMs to reflect and polish the initial responses, guaranteeing their accuracy and completeness.

Reflection OpenResearcher prompts LLMs to evaluate the accuracy and completeness of generated responses, meanwhile highlighting grammatical and semantic flaws.

Polishing OpenResearcher instructs LLMs to polish responses according to feedback received.

4 Demonstration

Our web application is built with Streamlit.⁶ Our databases encompass arXiv publications from Jan. 2023 to Jun. 2024, enriched with metadata. This is because most LLMs are trained on pre-2023 data, enabling them to retain this information. This fact also inspires OpenResearcher to answer simple questions without any retrieval, only using LLMs' internal knowledge. We utilize the state-of-the-art GTE-large model (Li et al., 2023) as dense vector and efficient-splade-VI-BT-large (Lassance and Clinchant, 2022) as sparse vector to vectorize our queries and paper chunks. These vectors serve for Hybrid Retrieval, and we use Qdrant⁷ for the vector storage. This Hybrid Retrieval tool extracts the 30 most similar chunks from each selected database. Elasticsearch⁸ supports our implementation of BM25 retriever, which extracts up to 80 chunks. The Bing⁹ API finds 10 relevant outcomes for the Internet Retrieval tool. Besides, we utilize bge-reranker-v2-m3¹⁰ to implement our Reranking tool. This Reranking tool reduces the number of retrieved chunks to 10. Lastly, we use DeepSeek-V2-Chat (DeepSeek-AI et al., 2024) as our backbone LLM to implement all LLM-powered tools, while also supporting various online LLM APIs and locally deployed LLMs through Ollama.¹¹

Figure 2, whose completed screenshot is shown in Figure 3 of Section A, demonstrates the strong

⁶<https://streamlit.io/>

⁷<https://qdrant.tech/>

⁸<https://github.com/elastic/elasticsearch>

⁹<https://www.bing.com/>

¹⁰<https://huggingface.co/BAAI/bge-reranker-large>

¹¹<https://ollama.com/>

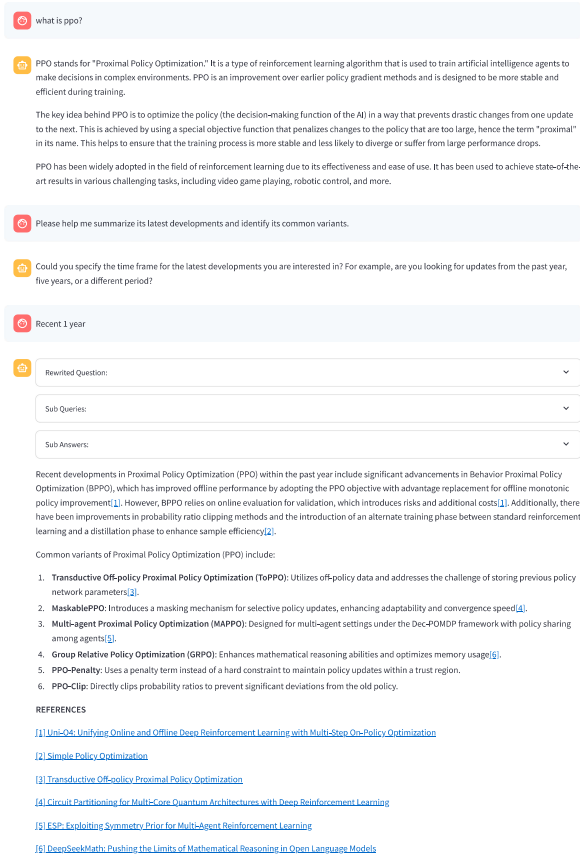


Figure 2: Case between user and OpenResearcher.

capability of OpenResearcher. Firstly, OpenResearcher can flexibly construct a tailored workflow for different queries, including simple queries and complex queries. For simple questions like “What is PPO?”, it directly employs LLMs to produce answers. For more complex queries like “Summarize the recent latest developments and variants of PPO?”, it utilizes multiple tools and provides users with essential details, including active queries, rewritten query, decomposed sub-queries and their sub-answers, retrieved outcomes of each sub-query after post-processing, generated final answer, and citation. This example can showcase its flexibility in handling different queries. With this benefit, our OpenResearcher can speed up responses and reduce computational costs.

Secondly, this figure also shows OpenResearcher can pose questions to users for query clarification. Different from previous passive applications that only answer questions, OpenResearcher utilizes LLMs’ internal knowledge to help users specify their question details. This tool is very crucial for junior students who often struggle to clearly express their questions and confusion.

Thirdly, Figure 2 demonstrates that OpenRe-

searcher supports conversational question answering, enabling users to engage in multi-turn dialogues. This feature allows for continuous and deeper discussions within OpenResearcher.

Lastly, this figure shows our OpenResearcher can enhance the quality and reliability of generated content by retrieving supporting evidence from the Internet and arXiv corpus. Additionally, we have developed a citation tool that links the generated text to the retrieved information, making it easy for researchers to verify the sources and delve deeper by reading the original papers.

5 Experiment

5.1 Evaluation Data

We have collected 109 research questions from more than 20 graduate students, comprising 38 questions on scientific paper recommendation, 38 on scientific text summarization, and 33 on others. These questions arise in their daily scientific research across areas including multimodal, agent, LLM alignment, tool learning, LLM safety, RAG, and others. Answers to these questions are commonly complex and lengthy, requiring graduate students to review many papers. Due to the considerable effort and cost of annotating ground truth answers, we opt to conduct a pairwise comparison instead of providing annotated ground truths.

5.2 Evaluation Applications

Our baseline includes recent industry applications, containing Perplexity AI, iAsk, You.com, and Phind, complemented by a Naive RAG that only utilizes our hybrid retrieval and LLM generation tools. Regarding our OpenResearcher, we remove the Active Query tool to directly obtain the answer. Our OpenResearcher flexibly uses these tools to generate answers without the need to follow the main workflow sequentially.

5.3 Evaluation Metric

In all evaluations, we compared the candidate outcomes from Naive RAG, OpenResearcher, iAsk, You.com, and Phind with those from Perplexity AI. If the candidate outcome outperforms Perplexity AI, it is notated as a “Win”.

We evaluate the generations from the three quality dimensions: (1) **Information Correctness** assesses the factual accuracy of the answers provided by the candidates. It is critical to determine if the information in each output is correct, as inaccuracies

Models	Correctness			Richness			Relevance		
	Win	Tie	Lose	Win	Tie	Lose	Win	Tie	Lose
Ask	2	16	12	12	6	12	2	8	20
You.com	3	21	6	9	5	16	4	13	13
Phind	2	26	2	15	7	8	5	13	12
Naive RAG	1	22	7	14	8	8	5	16	9
OpenResearcher	10	13	7	25	4	1	15	13	2

Table 1: Human Preference compared with Perplexity AI outcome. “Win” means that the current method beats Perplexity AI. More “Win” times means a superior application.

can severely undermine the utility of a QA system. (2) **Information Richness** involves evaluating the depth and scope of the information provided in the answers. Information richness captures whether an answer provides a thorough explanation or context beyond just addressing the question directly. (3) **Information Relevance** judges whether the information presented in the outputs is directly relevant to the question asked. Even if an answer is rich in information and correct, it may not be useful if it does not directly address the query.

5.4 Human Preference

We engaged 12 students with good research experience to conduct the human evaluation. Given the complexity of research questions, we randomly selected 30 questions for human evaluation, ensuring equal coverage of scientific question answering, scientific text summarization, and scientific paper recommendation. For quality control, each instance is annotated by two annotators whose agreement is measured. A third annotator can be involved to resolve disagreements between the two annotators.

The result is shown in Table 1 with an overall agreement of 90.67%. Our OpenResearcher achieves superior information correctness, relevance, and richness compared to all other applications. OpenResearcher significantly outperforms Perplexity AI with more “Win” than “Lose”. Specifically, compared to Naive RAG, OpenResearcher demonstrates better performance in all metrics. This suggests that our various tools significantly enhance the quality of the answers.

5.5 LLM Preference

Inspired by the widespread use of GPT-4 series for pairwise comparison (Zheng et al., 2023; Wang et al., 2023b; Sun et al., 2024) and their different preferences compared to humans (Li et al., 2024), we also utilize GPT-4o for LLM preference evaluation. We evaluate based on two criteria: infor-

Models	Richness			Relevance		
	Win	Tie	Lose	Win	Tie	Lose
iAsk	42	0	67	38	0	71
You.com	15	0	94	16	0	93
Phind	52	1	56	54	0	55
Naive RAG	41	1	67	57	0	52
OpenResearcher	62	2	45	74	0	35

Table 2: GPT-4o Preference Results compared with Perplexity AI outcome.

mation richness and relevance, since GPT-4o struggles to verify information accuracy without external knowledge. Despite the availability of citation papers, their length and quantity exceed LLMs’ capacity to confirm factuality.

The results are shown in Table 2. This supplemental LLM evaluation further demonstrates our system’s powerful performance. These results show our OpenResearcher achieves the best information relevance and richness among all applications. Furthermore, OpenResearcher surpasses Naive RAG in both metrics, demonstrating its superior performance due to our design.

6 Conclusion

We introduce OpenResearcher, an active AI assistant to accelerate the research process, catering to a broad spectrum of inquiries from researchers. OpenResearcher employs Retrieval-Augmented Generation (RAG) to enhance LLMs with the latest, verified, and domain-specific knowledge. It interacts with users to clarify their queries. Moreover, we have developed various tools for OpenResearcher to understand researchers’ queries, search from the scientific literature, filter retrieved information, provide accurate and comprehensive answers, and refine these answers. OpenResearcher can use these tools flexibly to build a pipeline that delivers accurate and comprehensive answers, outperforming those from industry applications, as

judged by human and GPT-4o.

Ethical Considerations

OpenResearcher integrates LLMs and search engines, known as retrieval-augmented generation (RAG), to accelerate scientific research. Despite being instructed to ground the generated responses in retrieved knowledge from scientific publications, LLMs may still generate hallucinations. Consequently, users are advised to verify crucial information derived from our LLM-based features.

Acknowledgments

The authors would like to thank the anonymous reviewers for their suggestions and feedback on the work. This work was partially funded by the National Natural Science Foundation of China (62476168), Qingyuan Research Project.

References

Xiaomei Bai, Mengyang Wang, Ivan Lee, Zhuo Yang, Xiangjie Kong, and Feng Xia. 2019. Scientific paper recommendation: A survey. *Ieee Access*, 7:9324–9339.

Pradeep Dasigi, Kyle Lo, Iz Beltagy, Arman Cohan, Noah A. Smith, and Matt Gardner. 2021. [A dataset of information-seeking questions and answers anchored in research papers](#). In *Proceedings of the 2021 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 4599–4610, Online. Association for Computational Linguistics.

DeepSeek-AI, Aixin Liu, Bei Feng, Bin Wang, Bingxuan Wang, Bo Liu, Chenggang Zhao, Chengqi Deng, Chong Ruan, Damai Dai, Daya Guo, Dejian Yang, Deli Chen, Dongjie Ji, Erhang Li, Fangyun Lin, Fuli Luo, Guangbo Hao, Guanting Chen, Guowei Li, H. Zhang, Hanwei Xu, Hao Yang, Haowei Zhang, Honghui Ding, Huajian Xin, Huazuo Gao, Hui Li, Hui Qu, J. L. Cai, Jian Liang, Jianzhong Guo, Jiaqi Ni, Jiashi Li, Jin Chen, Jingyang Yuan, Junjie Qiu, Junxiao Song, Kai Dong, Kaige Gao, Kang Guan, Lean Wang, Lecong Zhang, Lei Xu, Leyi Xia, Liang Zhao, Liyue Zhang, Meng Li, Miaojun Wang, Mingchuan Zhang, Minghua Zhang, Minghui Tang, Mingming Li, Ning Tian, Panpan Huang, Peiyi Wang, Peng Zhang, Qihao Zhu, Qinyu Chen, Qishi Du, R. J. Chen, R. L. Jin, Ruiqi Ge, Ruizhe Pan, Runxin Xu, Ruyi Chen, S. S. Li, Shanghao Lu, Shangyan Zhou, Shanhuang Chen, Shaoqing Wu, Shengfeng Ye, Shirong Ma, Shiyu Wang, Shuang Zhou, Shuiping Yu, Shunfeng Zhou, Size Zheng, T. Wang, Tian Pei, Tian Yuan, Tianyu Sun, W. L. Xiao, Wangding Zeng, Wei An, Wen Liu, Wenfeng Liang, Wenjun Gao, Wentao Zhang, X. Q. Li, Xiangyue Jin, Xianzu Wang, Xiao Bi, Xiaodong Liu, Xiaohan Wang,

Xiaojin Shen, Xiaokang Chen, Xiaosha Chen, Xiaotao Nie, Xiaowen Sun, Xiaoxiang Wang, Xin Liu, Xin Xie, Xingkai Yu, Xinnan Song, Xinyi Zhou, Xinyu Yang, Xuan Lu, Xuecheng Su, Y. Wu, Y. K. Li, Y. X. Wei, Y. X. Zhu, Yanhong Xu, Yanping Huang, Yao Li, Yao Zhao, Yaofeng Sun, Yaohui Li, Yaohui Wang, Yi Zheng, Yichao Zhang, Yiliang Xiong, Yilong Zhao, Ying He, Ying Tang, Yishi Piao, Yixin Dong, Yixuan Tan, Yiyuan Liu, Yongji Wang, Yongqiang Guo, Yuchen Zhu, Yuduan Wang, Yuheng Zou, Yukun Zha, Yunxian Ma, Yuting Yan, Yuxiang You, Yuxuan Liu, Z. Z. Ren, Zehui Ren, Zhangli Sha, Zhe Fu, Zhen Huang, Zhen Zhang, Zhenda Xie, Zhewen Hao, Zhihong Shao, Zhiniu Wen, Zhipeng Xu, Zhongyu Zhang, Zhuoshu Li, Zihan Wang, Zihui Gu, Zilin Li, and Ziwei Xie. 2024. [Deepseek-v2: A strong, economical, and efficient mixture-of-experts language model](#). *Preprint*, arXiv:2405.04434.

Yixi Ding, Yanxia Qin, Qian Liu, and Min-Yen Kan. 2023. [CocoSciSum: A scientific summarization toolkit with compositional controllability](#). In *Proceedings of the 2023 Conference on Empirical Methods in Natural Language Processing: System Demonstrations*, pages 518–526, Singapore. Association for Computational Linguistics.

Chao-Chun Hsu, Erin Bransom, Jenna Sparks, Bailey Kuehl, Chenhao Tan, David Wadden, Lucy Lu Wang, and Aakanksha Naik. 2024. [Chime: Llm-assisted hierarchical organization of scientific studies for literature review support](#). *arXiv preprint arXiv:2407.16148*.

Qiao Jin, Bhuwan Dhingra, Zhengping Liu, William Cohen, and Xinghua Lu. 2019. [PubMedQA: A dataset for biomedical research question answering](#). In *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing (EMNLP-IJCNLP)*, pages 2567–2577, Hong Kong, China. Association for Computational Linguistics.

Samiya Khan, Xiufeng Liu, Kashish A Shakil, and Mansaf Alam. 2017. [A survey on scholarly data: From big data perspective](#). *Information Processing & Management*, 53(4):923–944.

Martin Krallinger, Anastasia Krithara, Anastasios Nentidis, Georgios Paliouras, and Marta Villegas. 2020. [Bioasq at clef2020: Large-scale biomedical semantic indexing and question answering](#). In *Advances in Information Retrieval: 42nd European Conference on IR Research, ECIR 2020, Lisbon, Portugal, April 14–17, 2020, Proceedings, Part II 42*, pages 550–556. Springer.

Christin Katharina Kreutz and Ralf Schenkel. 2022. [Scientific paper recommendation systems: a literature review of recent publications](#). *Preprint*, arXiv:2201.00682.

Tom Kwiatkowski, Jennimaria Palomaki, Olivia Redfield, Michael Collins, Ankur Parikh, Chris Alberti,

- Danielle Epstein, Illia Polosukhin, Jacob Devlin, Kenton Lee, Kristina Toutanova, Llion Jones, Matthew Kelcey, Ming-Wei Chang, Andrew M. Dai, Jakob Uszkoreit, Quoc Le, and Slav Petrov. 2019. [Natural questions: A benchmark for question answering research](#). *Transactions of the Association for Computational Linguistics*, 7:452–466.
- Carlos Lassance and Stéphane Clinchant. 2022. [An efficiency study for splade models](#). In *Proceedings of the 45th International ACM SIGIR Conference on Research and Development in Information Retrieval, SIGIR '22*, page 2220–2226, New York, NY, USA. Association for Computing Machinery.
- Yoonjoo Lee, Kyungjae Lee, Sunghyun Park, Dasol Hwang, Jaehyeon Kim, Hong-in Lee, and Moontae Lee. 2023. Qasa: advanced question answering on scientific articles. In *Proceedings of the 40th International Conference on Machine Learning, ICML'23*. JMLR.org.
- Patrick Lewis, Ethan Perez, Aleksandra Piktus, Fabio Petroni, Vladimir Karpukhin, Naman Goyal, Heinrich Küttler, Mike Lewis, Wen-tau Yih, Tim Rocktäschel, Sebastian Riedel, and Douwe Kiela. 2020. [Retrieval-augmented generation for knowledge-intensive nlp tasks](#). In *Advances in Neural Information Processing Systems*, volume 33, pages 9459–9474. Curran Associates, Inc.
- Junlong Li, Fan Zhou, Shichao Sun, Yikai Zhang, Hai Zhao, and Pengfei Liu. 2024. Dissecting human and llm preferences. *arXiv preprint arXiv:2402.11296*.
- Zehan Li, Xin Zhang, Yanzhao Zhang, Dingkun Long, Pengjun Xie, and Meishan Zhang. 2023. Towards general text embeddings with multi-stage contrastive learning. *arXiv preprint arXiv:2308.03281*.
- Xiao Ma and Ranran Wang. 2019. Personalized scientific paper recommendation based on heterogeneous graph representation. *IEEE Access*, 7:79887–79894.
- G Manju, P Abhinaya, MR Hemalatha, GG Manju, et al. 2020. Cold start problem alleviation in a research paper recommendation system using the random walk approach on a heterogeneous user-paper graph. *International Journal of Intelligent Information Technologies (IJIT)*, 16(2):24–48.
- Anusri Pampari, Preethi Raghavan, Jennifer Liang, and Jian Peng. 2018. [emrQA: A large corpus for question answering on electronic medical records](#). In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing*, pages 2357–2368, Brussels, Belgium. Association for Computational Linguistics.
- Dimitris Pappas, Ion Androutsopoulos, and Haris Papageorgiou. 2018. [BioRead: A new dataset for biomedical reading comprehension](#). In *Proceedings of the Eleventh International Conference on Language Resources and Evaluation (LREC 2018)*, Miyazaki, Japan. European Language Resources Association (ELRA).
- Dimitris Pappas, Petros Stavropoulos, Ion Androutsopoulos, and Ryan McDonald. 2020. [BioMRC: A dataset for biomedical machine reading comprehension](#). In *Proceedings of the 19th SIGBioMed Workshop on Biomedical Language Processing*, pages 140–149, Online. Association for Computational Linguistics.
- Iratxe Pinedo, Mikel Larrañaga, and Ana Arruarte. 2024. Arzigo: A recommendation system for scientific articles. *Information Systems*, 122:102367.
- Shraman Pramanick, Rama Chellappa, and Subhashini Venugopalan. 2024. Spiqa: A dataset for multi-modal question answering on scientific papers. *arXiv preprint arXiv:2407.09413*.
- Federico Ruggeri, Mohsen Mesgar, and Iryna Gurevych. 2023. [A dataset of argumentative dialogues on scientific papers](#). In *Proceedings of the 61st Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 7684–7699, Toronto, Canada. Association for Computational Linguistics.
- Nazmus Sakib, Rodina Binti Ahmad, and Khalid Haruna. 2020. A collaborative approach toward scientific paper recommendation using citation context. *IEEE Access*, 8:51246–51255.
- Vaios Stergiopoulos, Michael Vassilakopoulos, Eleni Tousidou, and Antonio Corral. 2024. An academic recommender system on large citation data based on clustering, graph modeling and deep learning. *Knowledge and Information Systems*, pages 1–34.
- Shichao Sun, Ruifeng Yuan, Ziqiang Cao, Wenjie Li, and Pengfei Liu. 2024. [Prompt chaining or stepwise prompt? refinement in text summarization](#). *Preprint, arXiv:2406.00507*.
- Sotaro Takeshita, Tommaso Green, Ines Reinig, Kai Eckert, and Simone Ponzetto. 2024. [ACLSum: A new dataset for aspect-based summarization of scientific publications](#). In *Proceedings of the 2024 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies (Volume 1: Long Papers)*, pages 6660–6675, Mexico City, Mexico. Association for Computational Linguistics.
- William Tanner, Esra Akbas, and Mir Hasan. 2019. Paper recommendation based on citation relation. In *2019 IEEE international conference on big data (big data)*, pages 3053–3059. IEEE.
- Hanchen Wang, Tianfan Fu, Yuanqi Du, Wenhao Gao, Kexin Huang, Ziming Liu, Payal Chandak, Shengchao Liu, Peter Van Katwyk, Andreea Deac, et al. 2023a. Scientific discovery in the age of artificial intelligence. *Nature*, 620(7972):47–60.
- Pancheng Wang, Shasha Li, Kunyuan Pang, Liangliang He, Dong Li, Jintao Tang, and Ting Wang. 2022. [Multi-document scientific summarization from a knowledge graph-centric view](#). In *Proceedings of*

the 29th International Conference on Computational Linguistics, pages 6222–6233, Gyeongju, Republic of Korea. International Committee on Computational Linguistics.

Yidong Wang, Zhuohao Yu, Zhengran Zeng, Linyi Yang, Cunxiang Wang, Hao Chen, Chaoya Jiang, Rui Xie, Jindong Wang, Xing Xie, et al. 2023b. Pandalm: An automatic evaluation benchmark for llm instruction tuning optimization. *arXiv preprint arXiv:2306.05087*.

Fangyuan Xu, Kyle Lo, Luca Soldaini, Bailey Kuehl, Eunsol Choi, and David Wadden. 2024. Kiwi: A dataset of knowledge-intensive writing instructions for answering research questions. *arXiv preprint arXiv:2403.03866*.


Xiaoming Zhai. 2023. [Chatgpt for next generation science learning](#). *XRDS*, 29(3):42–46.

Xingjian Zhang, Yutong Xie, Jin Huang, Jinge Ma, Zhaoying Pan, Qijia Liu, Ziyang Xiong, Tolga Ergen, Dongsub Shim, Honglak Lee, and Qiaozhu Mei. 2024. [Massw: A new dataset and benchmark tasks for ai-assisted scientific workflows](#). *Preprint*, arXiv:2406.06357.

Lianmin Zheng, Wei-Lin Chiang, Ying Sheng, Siyuan Zhuang, Zhonghao Wu, Yonghao Zhuang, Zi Lin, Zhuohan Li, Dacheng Li, Eric Xing, et al. 2023. Judging llm-as-a-judge with mt-bench and chatbot arena. *arXiv preprint arXiv:2306.05685*.

A Completed Case

Ask anything you want to know!



[GABRIEL](#)

Example

What is PPO?

In multimodal pretraining, the...

Search the paper and tell about...

Clear Chat History

OpenResearcher

Welcome to *OpenResearcher*, an advanced Scientific Research Assistant designed to provide a helpful answer to a research query. With access to the arXiv corpus, *OpenResearcher* is can provide you with the latest scientific insights. Explore the frontiers of science with *OpenResearcher*—where answers await.

what is ppo?

PPO stands for "Proximal Policy Optimization." It is a type of reinforcement learning algorithm that is used to train artificial intelligence agents to make decisions in complex environments. PPO is designed to optimize the performance of a policy, which is a set of rules or guidelines that an agent follows to make decisions.

PPO is an improvement over earlier policy optimization algorithms, such as Trust Region Policy Optimization (TRPO), because it is more stable and easier to implement. PPO works by iteratively updating the policy in small steps, ensuring that each update does not deviate too far from the previous policy. This helps to prevent the agent from making drastic changes to its behavior that could lead to instability or poor performance.

PPO has been shown to be effective in a wide range of applications, including video games, robotics, and other complex control tasks. It is a popular choice for researchers and practitioners in the field of reinforcement learning because of its robustness and ease of use.

Please help me summarize its latest developments and identify its common variants.

Could you specify the time range for the latest developments you are interested in? For example, are you looking for updates from the past year, five years, or another specific period?

recent 1 year

Revised Question:

Can you provide a summary of the latest developments in Proximal Policy Optimization (PPO) within the past year, and also identify its common variants?

Sub Query:

```

1 : "Retrieve information on recent developments in Proximal Policy Optimization (PPO) within the past year."
2 : "Identify common variants of Proximal Policy Optimization (PPO)."
```

Sub Answers:

2311.03351	2311.03350	2402.02774	2401.16025	2310.17490	2402.15246	2312.00111	2401.17976	2405.09662	2403.08943	2305.08477	2306.17178	2307.36336	2402.03111
------------	------------	------------	------------	------------	------------	------------	------------	------------	------------	------------	------------	------------	------------

RETRIEVED INFO:

id: 2311.03351 Recently, the work on Behavior Proximal Policy Optimization (BPPO) has made significant advancements by adopting the PPO (ppo) objective with advantage replacement to achieve offline monotonic policy improvement. This has resulted in improved offline performance compared to one-step offline algorithms (onstepRL). However, BPPO heavily relies on online evaluation to validate the improvement of the behavior policy, introducing risks and additional costs.

SUB ANSWER:

Recent developments in Proximal Policy Optimization (PPO) within the past year include several significant advancements:

- Behavior Proximal Policy Optimization (BPPO):** This approach has made notable progress by adopting the PPO objective with advantage replacement to achieve offline monotonic policy improvement. BPPO has shown improved offline performance compared to one-step offline algorithms. However, it relies heavily on online evaluation to validate the improvement of the behavior policy, which introduces risks and additional costs.
- Novel Probability Ratio Clipping Methods or Objective Functions:** Researchers have designed new probability ratio clipping methods or objective functions to enhance the performance of PPO. These include contributions from Queeney et al. (2021), Garg et al. (2021), and Markowitz & Staley (2023).
- Alternate Training Phases:** There have been developments in alternating training phases between standard reinforcement learning and distillation phases to improve sample efficiency. This approach has been explored by Cobbe et al. (2021) and Wang et al. (2023).

These advancements highlight the continued evolution and refinement of PPO methods to address various challenges in reinforcement learning.

RETRIEVED INFO:

id: 2401.17976 MaskablePPO is a novel variant of the Proximal Policy Optimization (PPO) algorithm. It introduces a masking mechanism to enhance its capabilities by extending the core PPO algorithm with a mask that selectively determines which parts of the policy should undergo updates during training. This fine-grained control enhances adaptability and control during the training process and achieves faster convergence.

SUB ANSWER:

Common variants of Proximal Policy Optimization (PPO) include:

- PPO-Clip:** This is the most popular variant, proposed by Schulman et al. in 2017. It optimizes the actor by maximizing the surrogate objective at each rollout. PPO-Clip clips the probability ratio in the objective function to prevent it from becoming too large or too small, which effectively limits the change in the policy at each update.
- PPO-Penalty:** This variant approximately solves a KL-constrained update like TRPO, but penalizes the KL-divergence in the objective function instead of making it a hard constraint. It automatically adjusts the penalty coefficient over the course of training to ensure it is scaled appropriately.
- MaskablePPO:** This novel variant introduces a masking mechanism to enhance its capabilities by selectively determining which parts of the policy should undergo updates during training. This fine-grained control enhances adaptability and control during the training process and achieves faster convergence.
- Multi-agent Proximal Policy Optimization (MAPPO):** This variant is specialized for multi-agent settings. In the default formulation of MAPPO under the Decentralized Partially Observable Markov Decision Process (Dec-POMDP) framework, the agents employ the trick of policy sharing.
- Group Relative Policy Optimization (GRPO):** This variant enhances mathematical reasoning abilities while concurrently optimizing the memory usage of PPO.
- Direct Preference Optimization (DPO) and Contrastive Preference Optimization (CPO):** These variants have been shown to be more stable, computationally efficient, and effective on various NLP tasks.
- PPO-RPE:** This variant uses relative Pearson (RPE) divergence to design the threshold adaptively. In PPO-RPE, the relative density ratio, which can be formed with symmetry, replaces the raw density ratio. This symmetry allows for easier estimation of the error scale from the center, enabling the threshold to be adapted for the estimated error scale.

These variants of PPO aim to improve stability, efficiency, and performance across various applications and settings.

Recent developments in Proximal Policy Optimization (PPO) within the past year include several significant advancements:

- Behavior Proximal Policy Optimization (BPPO):** This approach has made notable progress by adopting the PPO objective with advantage replacement to achieve offline monotonic policy improvement^[1]. BPPO has shown improved offline performance compared to one-step offline algorithms. However, it relies heavily on online evaluation to validate the improvement of the behavior policy, which introduces risks and additional costs^[1].
- Novel Probability Ratio Clipping Methods or Objective Functions:** Researchers have designed new probability ratio clipping methods or objective functions to enhance the performance of PPO. These include contributions from Queeney et al. (2021), Garg et al. (2021), and Markowitz & Staley (2023).
- Alternate Training Phases:** There have been developments in alternating training phases between standard reinforcement learning and distillation phases to improve sample efficiency. This approach has been explored by Cobbe et al. (2021) and Wang et al. (2023).

Common variants of Proximal Policy Optimization (PPO) include:

- PPO-Clip:** This is the most popular variant, proposed by Schulman et al. in 2017. It optimizes the actor by maximizing the surrogate objective at each rollout. PPO-Clip clips the probability ratio in the objective function to prevent it from becoming too large or too small, which effectively limits the change in the policy at each update.
- PPO-Penalty:** This variant approximately solves a KL-constrained update like TRPO, but penalizes the KL-divergence in the objective function instead of making it a hard constraint. It automatically adjusts the penalty coefficient over the course of training to ensure it is scaled appropriately.
- MaskablePPO:** This novel variant introduces a masking mechanism to enhance its capabilities by selectively determining which parts of the policy should undergo updates during training^[2]. This fine-grained control enhances adaptability and control during the training process and achieves faster convergence^[2].
- Multi-agent Proximal Policy Optimization (MAPPO):** This variant is specialized for multi-agent settings^[3]. In the default formulation of MAPPO under the Decentralized Partially Observable Markov Decision Process (Dec-POMDP) framework, the agents employ the trick of policy sharing^[3].
- Group Relative Policy Optimization (GRPO):** This variant enhances mathematical reasoning abilities while concurrently optimizing the memory usage of PPO^[4].
- Direct Preference Optimization (DPO) and Contrastive Preference Optimization (CPO):** These variants have been shown to be more stable, computationally efficient, and effective on various NLP tasks^[5].
- PPO-RPE:** This variant uses relative Pearson (RPE) divergence to design the threshold adaptively^[6]. In PPO-RPE, the relative density ratio, which can be formed with symmetry, replaces the raw density ratio^[6]. This symmetry allows for easier estimation of the error scale from the center, enabling the threshold to be adapted for the estimated error scale^[6].

These advancements and variants highlight the continued evolution and refinement of PPO methods to address various challenges in reinforcement learning.

REFERENCES

[1] [Use-Or-Modify-Online-and-Offline-Deep-Reinforcement-Learning-with-Multi-Step-On-Policy-Optimization](#)

[2] [Control-Partitioning-for-Multi-Core-Quantum-Architectures-with-Deep-Reinforcement-Learning](#)

[3] [ESF: Enabling Symmetry Prior for Multi-Agent Reinforcement Learning](#)

[4] [DeepSeekMath: Pushing the Limits of Mathematical Reasoning in Open Language Models](#)

[5] [Authorship Style Transfer with Policy Optimization](#)

[6] [Proximal Policy Optimization with Adaptive Threshold for Symmetric Relative Density Ratio](#)

What do you want to know? I will give you an answer. >

Figure 3: Screenshot showing the completed case in Figure 2.