

TIB AIssistant: a Platform for AI-Supported Research Across Research Life Cycles

Allard Oelen¹, Sören Auer^{1,2}

¹TIB – Leibniz Information Centre for Science and Technology, Hannover, Germany

²L3S Research Center, Leibniz University of Hannover, Hannover, Germany

Abstract

The rapidly growing popularity of adopting Artificial Intelligence (AI), and specifically Large Language Models (LLMs), is having a widespread impact throughout society, including the academic domain. AI-supported research has the potential to support researchers with tasks across the entire research life cycle. In this work, we demonstrate the TIB AIssistant, an AI-supported research platform providing support throughout the research life cycle. The AIssistant consists of a collection of assistants, each responsible for a specific research task. In addition, tools are provided to give access to external scholarly services. Generated data is stored in the assets and can be exported as an RO-Crate bundle to provide transparency and enhance reproducibility of the research project. We demonstrate the AIssistant's main functionalities by means of a sequential walk-through of assistants, interacting with each other to generate sections for a draft research paper. In the end, with the AIssistant, we lay the foundation for a larger agenda of providing a community-maintained platform for AI-supported research.

Keywords

AI Assistant, AI-Assisted Research, Scholarly Assistant, Scholarly AI Platform

1. Introduction

The recent advancements of Artificial Intelligence (AI), in particular generative AI, such as Large Language Models (LLMs), have a profound impact on our society already [1]. Among other fields, scholarly research is one of the areas where AI already has, and will likely even more in the future, change how work is conducted [2]. However, to fully leverage AI for research, the researcher needs to be aware of the many different approaches and needs to have experience with guiding the LLM in such a way that it produces outputs that are useful for their research. Especially due to the growing number of tools, techniques, and approaches, it can be overwhelming and challenging to effectively leverage AI for scholarly research.

To address these challenges, we present the TIB AIssistant (i.e., AI-assistant). The TIB AIssistant is a domain-agnostic AI-supported research platform that helps researchers, by means of AI, in various steps of the research life cycle. This work demonstrates a concrete implementation of our vision of the TIB AIssistant, which is published as a vision paper [3]. While the vision paper outlines a high-level list of design principles and considerations, this article provides additional value by introducing a concrete implementation of the AIssistant platform, a workflow of different assistants, and a demonstration of the user interactions with the platform. We aim to provide support across the research life cycle by means of different agents (hereafter called *assistants*). This ranges from providing guidance during ideation, helping to share research questions, writing up articles, to publishing a paper. The assistants take in certain inputs and generate specific outputs. If the assistants are used in sequence, generated output is stored and are used as input for the next assistant. However, it is also possible to use the assistants in isolation, i.e., to perform a single task. Assistants can use a predefined set of tools in any arbitrary order to accomplish their tasks. Most of these tools are external services, called via REST endpoints, and are called when deemed necessary by the LLM. For example, this makes it possible for assistants to find related work via actual scholarly search platforms, or to find articles based on certain

ISWC 2025 Companion Volume, November 2–6, 2025, Nara, Japan

✉ allard.oelen@tib.eu (A. Oelen); auer@tib.eu (S. Auer)

ORCID 0000-0001-9924-9153 (A. Oelen); 0000-0002-0698-2864 (S. Auer)



© 2025 Copyright for this paper by its authors. Use permitted under Creative Commons License Attribution 4.0 International (CC BY 4.0).

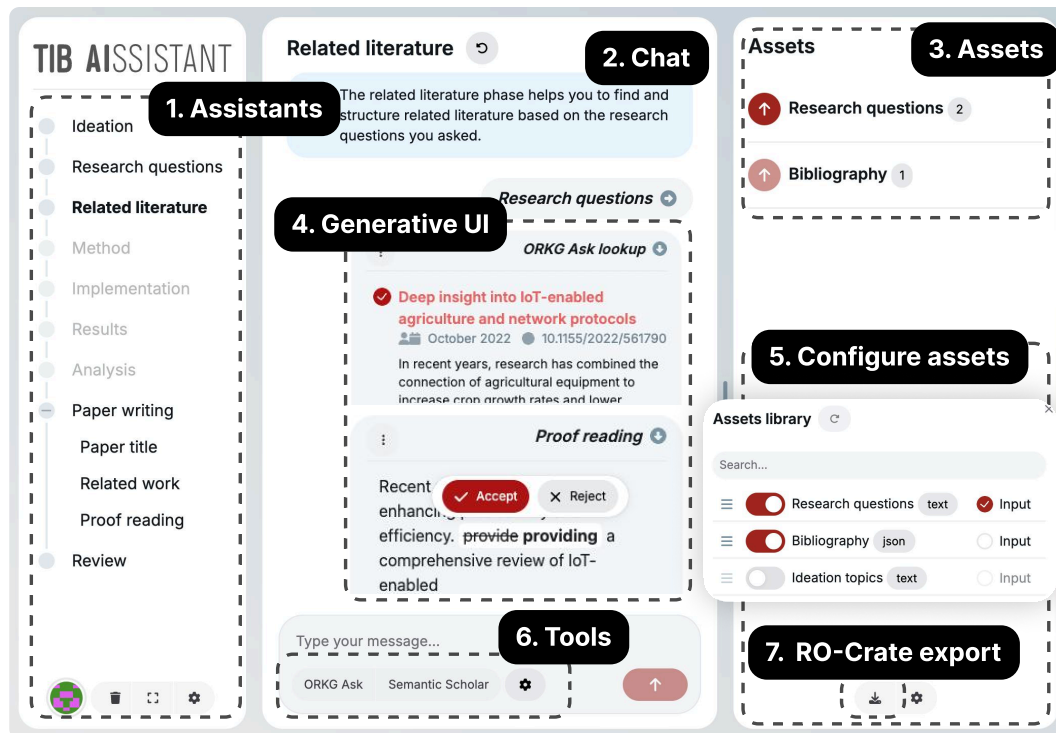


Figure 1: Screenshot of the TIB AIssistant interface. 1. The assistants that can either be used sequentially or individually. 2. Chat window with user and assistant messages, including tool responses. 3. Assets related to the selected assistant. 4. Generative UI selected based on invoked tools (minified example showing two tools: literature search and proofreading). 5. Configure assets; a similar modal is available for assistants and tools. 6. Tools related to the selected assistant. 7. Export assets button to generate a RO-Crate bundle.

identifiers (e.g., DOIs or ORCIDs). Transparency is a crucial aspect of AI-supported research, meaning that provenance data needs to be published alongside the research output to ensure full openness about the approach. We implement the RO-Crate [4] approach and use additional ontologies to record how knowledge is created and by whom.

In the end, the AIssistant serves as a library of assistants and tools for academic research, and provides a single platform in which they are integrated. In this work, we demonstrate how the platform operates and show how a researcher can use various assistants, starting from a research idea formulation and fetching related work, to drafting paper sections.

2. Related Work

The industrial revolution and afterward the digital revolution had a substantial impact on all aspects of life [5]. The rise of AI is considered to be potentially even more impactful than the previous revolutions [6]. Now with the rise of Generative AI, and specifically Large Language Models growing in size, this impact of AI is rapidly accelerated [7]. For example, LLMs are expected to have a significant impact on the labor market [8]. Furthermore, LLMs have an impact on many other fields, including medicine [9], finance [10], and law [11]. LLMs have a significant impact on different aspects of scholarly research as well, for example, for generating research ideas [12], literature reviews [13], and paper writing [14]. To effectively use LLMs for such purposes, prompt engineering is of crucial importance [15]. The ability to write effective prompts becomes an important skill for professionals seeking to leverage AI for their work [16]. To this end, with the TIB AIssistant, we provide a library of prompts (i.e., the assistants) that have been manually curated and can be integrated directly inside research workflows.

There are many individual approaches to leveraging LLMs for research. Some of these approaches aim to create a fully automated research life cycle using AI. This includes the work of Sakana AI by

means of the AI Scientist, which fully automates the research life cycle by iteratively executing life cycle phases [17]. There are several shortcomings of their approach, for example, the limited ability to find related literature, and issues during code generation leading to failing experiments [18]. While fully automated approaches are an interesting research avenue and are surely going to be improved in the near future, we take a different approach with the AIAssistant. We believe that researchers must be at the center of doing research, and therefore, AI should be an assistive technology, leaving the human researcher in control of the process.

3. Approach

The AIAssistant’s approach consists of three main components: assistants, tools, and assets. We will discuss those concepts on a high level to lay the foundation for the platform.

Assistants are agents that are designed to perform a single task. For example, formulating research questions or coming up with a paper title. Assistants consist of a name, system prompt, LLM model, what assets they consume and generate, and the tools that can be called. The LLM model selection determines which model to use to accomplish the task. Simpler tasks can be performed by smaller models, reducing the required resources and waiting time.

Tools are generally external web services that are called via REST endpoints. The tools have a name, a textual description, an input schema, and an execution function. The description and input schema are provided to the LLM to provide instructions on when the tool should be called and how it should be called. This is handled via the *Tool Calling* mechanism [19]. Each assistant specifies the tools that are callable, which limits the space of possible tools and thus improves the tool selection accuracy. The output of tools is either sent to the chat as text or handled by displaying specialized UI components. This follows the paradigm of *Generative UI*, where UI components are shown on demand, depending on the task at hand. In section 4, we demonstrate several implemented generative UI components.

Assets are the data store that assistants can access. This differs from the conversation’s history (i.e., context window) as it is structured and only contains data that the user selects. Assets have a name and a data type (e.g., text, object, JSON, etc.). Assistants can have required input assets, which either consist of previously generated assets or are manually provided by the user. Reproducibility is one of the cornerstones of science. Therefore, being transparent about how AI was used to support the researcher’s work is a crucial aspect of our approach. The assets are central in our approach, as they serve as output of the generated knowledge. We provide an *Export assets* functionality, which compresses assets into an archive file, containing individual files for each asset, and a metadata file which provides provenance data for the files. We use the RO-Crate [4] for the metadata, which uses JSON-LD and Schema.org annotations to package research artifacts. During the export process, the user can provide additional provenance data (i.e., their name and the license of the data). In addition, we use SPAR ontologies, such as DOCO and DEO [20], to provide more fine-grained semantic annotations for exported data where appropriate. We envision that the generated RO-Crate bundles will be published alongside the research articles to provide transparency of the process. The semantic representations of the generated knowledge provide a more machine-actionable way to access and analyze the knowledge.

4. Implementation and Demonstration

A screenshot of the prototype user interface is displayed in Figure 1. The previously discussed key components are depicted in this screenshot. The three main components, Assistants, Assets, and Tools, are configurable by users. By default, the selected assistant decides what assets and tools are activated, but users can modify this selection to fit their use cases. We will demonstrate here via a system walk-through how a research life cycle using the AIAssistant can look like. Afterward, we discuss the technical details of our prototype.

Table 1

Implemented domain-agnostic assistants following a sequential pipeline. For each assistant, the input and output assets and tools are listed. Certain assistants use input assets from previously generated output assets.

Name	Description	Input assets	Output assets	Tools	Generative UI
Ideation	Come up with research ideas based on user-provided input	-	Ideation topics	Crossref, ORCID, PDF URL, Unpaywall	-
Research questions	Define research questions based on research topics	Ideation topics	Research questions	-	-
Related literature	Find related literature based on research questions	Research questions	Bibliography	ORKG Ask, Semantic Scholar	Literature search component
Paper writing: title	Write a suitable paper title based on provided content	Ideation topics, research questions	Paper: title	-	-
Paper writing: related work	Write a related work section using a bibliography	Research questions, bibliography	Paper: related work	-	-
Paper writing: proofreading	Proofread the text and propose changes	Paper	Paper	-	Track changes component
Review	Review the paper content similarly as a peer reviewer	Paper	Paper	-	-

4.1. System Walk-Through

In our walk-through, we adopt the stance of a researcher who is interested in intertwining Semantic Web and AI research. We will follow the life cycle as depicted in the sidebar of Figure 1. For demonstration purposes, we only show domain-agnostic assistants, domain-specific ones are disabled, as can be seen in the screenshot. A list of assistants is provided in Table 1. A demonstration video is available online.¹

Firstly, the researcher wants to generate ideas related to their topic of interest in the *Ideation* assistant. This assistant asks the user to provide content they are interested in. In our example, the researcher provides a DOI from related papers and an ORCID, which will respectively fetch related content from Crossref and list articles from ORCID. After refining the ideas via the chat, the researcher adds the final topics to the ideation asset. Next, the researcher goes to the *Research questions* assistant and inputs the previously generated ideation topics from the assets. After iterative refinement in the chat component, the final research questions are added to the assets.

As a next step, the researcher aims to find related work via the *Related literature* assistant. There, the research questions are provided to the chat and are used to find related work via ORKG Ask [21]. The user manually reviews the suggested literature and adds relevant articles to the bibliography. This is done via the generative UI approach, where a dedicated literature search component is displayed in the chat window. This component is displayed only for tools that support literature search. It lists articles, allows users to select articles, and provides pagination to explore the listed items further.

Afterward, the paper writing assistants are used: *Paper title*, *Related work*, and *Proofread*. The first two assistants use the ideation topics and research questions. The related work assistant also uses the bibliography to connect related work to their citations. The final text and bibliography can be exported to LaTeX and imported to an Overleaf project as a draft for writing the actual manuscript. The proofread assistant uses another type of generative UI component, providing a dedicated interface where users can accept and reject proposed changes via a track-changes-like interface. As a final step of the walk-through, the authored text is reviewed via the *Review* assistant.

4.2. Technologies

The AIAssistant is implemented in Typescript and React using the Next.js framework and is published as open source software under a permissive MIT license.² We use Next.js both as frontend and backend, with the backend primarily serving as a wrapper for making calls to the LLMs. We use the OpenAI API, and specifically the GPT-4o mini model, which provides a good balance between costs and performance.

¹<https://doi.org/10.5446/71179>

²<https://gitlab.com/TIBHannover/orkg/tib-aissant/web-app>

The Vercel AI SDK ensures that model providers and individual models can be changed when necessary. Local browser storage via IndexedDB is used to store assets and chat history. For authentication, we use ORKG accounts with enable single sign-on. The service can only be used when the user is signed in to ensure the number of daily used tokens can be effectively limited as a means of cost management. In the future, we will implement a Bring Your Own Key (BYOK) approach, allowing users to utilize their own API keys to access the platform without imposed token limits. To encourage community development of both the platform and content, we have published development documentation online.³

5. Discussion and Conclusion

The TIB AIAssistant is not a direct alternative or replacement of existing LLM user interfaces. The AIAssistant is tailored toward research and provides various advantages compared to regular LLM interfaces. First and foremost, the integration of assistants provides researchers with a library of curated prompts, which enables them to use LLMs within their research without the need to craft these prompts themselves. LLMs themselves provide almost limitless possibilities for research, but without curated lists of use cases (i.e., prompts), for many researchers, these benefits remain out of reach. Furthermore, the integrated tool library focuses on research-oriented tools and services, which are not typically provided by regular LLM interfaces. Tools can utilize dedicated user interface components, via the generative UI paradigm, to offer additional interaction possibilities beyond the conversational (i.e., text-only) approach of regular LLM interfaces. In the end, we envision the AIAssistant as a platform that integrates many different prompts, models, and tools from a wide variety of research domains, created and maintained by the research communities themselves. Alongside the open source nature of the AIAssistant, the collaboration of research communities to support different domains is a key aspect of our approach. This sets the AIAssistant apart from existing approaches and other commercial initiatives. To facilitate community collaboration, we have provided developer documentation and an easy-to-use format for defining assistants to encourage community contributions. At the current stage, we have provided a framework in which the previously mentioned aspects can be integrated. We demonstrated how different assistants and prompts can be used to perform research tasks, however, the actual implementation of a large-scale domain-specific prompt and tool library is left for future work. To summarize, although the individual tasks within the assistant can indeed also be executed by existing LLM interfaces, the AIAssistant platform is tailored toward scholarly research. The integration of different scholarly tools and domain-specific prompts into a single platform makes it possible for any researcher, regardless of the AI proficiency, to leverage AI in their workflows.

In the current implementation, all assistants use the same model (GPT 4o-mini). The technology stack supports switching models and providers. For now, we do not want to be limited by the capabilities of the model, e.g., the ability to select the right tool for tool calling, or the capability to generate research ideas for arbitrary domains. However, we envision that assistants will use a model tailored to the task at hand. This means that smaller, locally hosted, or fine-tuned models are most likely viable options for various assistants. Using such approaches reduces the models' required resources, benefiting both costs and environmental impact, and is therefore an important research direction for future work. Another aspect of future work is the evaluation of our approach. This will focus both on specific assistants and how well they are able to accomplish their respective objectives, as well as evaluating the platform from a usability perspective. Experimenting with these smaller models is out of scope for this work.

To conclude, in this work, we demonstrated what an AI-supported research assistant looks like, implemented in the TIB AIAssistant. With this, we lay the foundation for a comprehensive community-maintained platform for researchers who want to leverage AI for their work. We demonstrated how different assistants work together by storing intermediary results in the assets storage. Also, we showed how external tools play an essential role in providing a research-oriented platform. Finally, we used a RO-Crate export functionality to provide machine-readable provenance data of the generated content. In the future, we plan to further simplify the process of adding assistants and tools to foster

³<https://tibhannover.gitlab.io/orkg/tib-aissistant/web-app/storybook>

further community involvement. To accomplish this, we are experimenting with no-code solutions, where graphical user interfaces allow users to configure the platform further. Additionally, we plan to experiment with different models and add additional assistants and tools to support more complete research life cycles for various domains.

Acknowledgments

We thank our colleague Mohamad Yaser Jaradeh for his valuable comments in reviewing this paper. Additionally, we want to thank the entire AIAssistant team for their contributions to the platform, including research and development efforts. This work was co-funded by NFDI4DataScience (ID: 460234259) and by the TIB Leibniz Information Centre for Science and Technology.

Declaration on Generative AI

During the preparation of this work, the authors used ChatGPT and Grammarly in order to: Paraphrase and reword, improve writing style, and Grammar and spelling check. Also, the authors used OpenAI TTS for generating the voice-over in the demo video. After using these services, the authors reviewed and edited the content as needed and take full responsibility for the publication's content.

References

- [1] R. Bommasani, D. A. Hudson, E. Adeli, R. Altman, S. Arora, S. von Arx, M. S. Bernstein, J. Bohg, A. Bosselut, E. Brunskill, et al., On the opportunities and risks of foundation models, arXiv preprint arXiv:2108.07258 (2021). doi:10.48550/arXiv.2108.07258.
- [2] C. A. Bail, Can generative ai improve social science?, Proceedings of the National Academy of Sciences 121 (2024) e2314021121. doi:10.1073/pnas.2314021121.
- [3] S. Auer, A. Oelen, M. Y. Jaradeh, M. Khalid, F. Keya, S. K. Gaddipati, J. D'Souza, L. Schlüter, A. Alasti, G. Rabby, A. Jiomekong, O. Karras, Towards ai-supported research: a vision of the tib aassistant, 5th International Workshop on Scientific Knowledge: Representation, Discovery, and Assessment (Sci-K), November 2-6, 2025, Nara, Japan (2025).
- [4] S. Soiland-Reyes, P. Sefton, M. Crosas, L. J. Castro, F. Coppens, J. M. Fernández, D. Garijo, B. Grüning, M. L. Rosa, S. Leo, E. O. Carragáin, M. Portier, A. Trisovic, R.-C. Community, P. Groth, C. Goble, Packaging research artefacts with ro-crate, Data Science 5 (2022) 97–138. doi:10.3233/DS-210053.
- [5] M. Xu, J. M. David, S. H. Kim, et al., The fourth industrial revolution: Opportunities and challenges, International journal of financial research 9 (2018) 90–95. doi:10.5430/ijfr.v9n2p90.
- [6] S. Makridakis, The forthcoming artificial intelligence (ai) revolution: Its impact on society and firms, Futures 90 (2017) 46–60. doi:10.1016/j.futures.2017.03.006.
- [7] A. Tamkin, M. Brundage, J. Clark, D. Ganguli, Understanding the capabilities, limitations, and societal impact of large language models, arXiv preprint arXiv:2102.02503 (2021). doi:10.48550/arXiv.2102.02503.
- [8] T. Eloundou, S. Manning, P. Mishkin, D. Rock, Gpts are gpts: Labor market impact potential of llms, Science 384 (2024) 1306–1308. doi:10.1126/science.adj0998.
- [9] J. Clusmann, F. R. Kolbinger, H. S. Muti, Z. I. Carrero, J.-N. Eckardt, N. G. Laleh, C. M. L. Löffler, S.-C. Schwarzkopf, M. Unger, G. P. Veldhuizen, et al., The future landscape of large language models in medicine, Communications medicine 3 (2023) 141. doi:10.1038/s43856-023-00370-1.
- [10] Y. Li, S. Wang, H. Ding, H. Chen, Large language models in finance: A survey, in: Proceedings of the Fourth ACM International Conference on AI in Finance, ICAIF '23, Association for Computing Machinery, New York, NY, USA, 2023, p. 374–382. doi:10.1145/3604237.3626869.

- [11] Z. Z. Chen, J. Ma, X. Zhang, N. Hao, A. Yan, A. Nourbakhsh, X. Yang, J. McAuley, L. Petzold, W. Y. Wang, A survey on large language models for critical societal domains: Finance, healthcare, and law, arXiv preprint arXiv:2405.01769 (2024). doi:10.48550/arXiv.2405.01769.
- [12] C. Si, D. Yang, T. Hashimoto, Can llms generate novel research ideas? a large-scale human study with 100+ nlp researchers, arXiv preprint arXiv:2409.04109 (2024). doi:10.48550/arXiv.2409.04109.
- [13] F. Bolanos, A. Salatino, F. Osborne, E. Motta, Artificial intelligence for literature reviews: Opportunities and challenges, Artificial Intelligence Review 57 (2024) 259. doi:10.1007/s10462-024-10902-3.
- [14] W. Liang, Y. Zhang, Z. Wu, H. Lepp, W. Ji, X. Zhao, H. Cao, S. Liu, S. He, Z. Huang, et al., Mapping the increasing use of llms in scientific papers, arXiv preprint arXiv:2404.01268 (2024). doi:10.48550/arXiv.2404.01268.
- [15] J. White, Q. Fu, S. Hays, M. Sandborn, C. Olea, H. Gilbert, A. Elnashar, J. Spencer-Smith, D. C. Schmidt, A prompt pattern catalog to enhance prompt engineering with chatgpt, arXiv preprint arXiv:2302.11382 (2023). doi:10.48550/arXiv.2302.11382.
- [16] B. Meskó, Prompt engineering as an important emerging skill for medical professionals: tutorial, Journal of medical Internet research 25 (2023) e50638. doi:10.2196/50638.
- [17] C. Lu, C. Lu, R. T. Lange, J. Foerster, J. Clune, D. Ha, The ai scientist: Towards fully automated open-ended scientific discovery, arXiv preprint arXiv:2408.06292 (2024). doi:10.48550/arXiv.2408.06292.
- [18] J. Beel, M.-Y. Kan, M. Baumgart, Evaluating sakana’s ai scientist for autonomous research: Wishful thinking or an emerging reality towards’ artificial research intelligence’(ari)?, arXiv preprint arXiv:2502.14297 (2025). doi:10.48550/arXiv.2502.14297.
- [19] T. Schick, J. Dwivedi-Yu, R. Dessi, R. Raileanu, M. Lomeli, E. Hambro, L. Zettlemoyer, N. Cancedda, T. Scialom, Toolformer: Language models can teach themselves to use tools, in: A. Oh, T. Naumann, A. Globerson, K. Saenko, M. Hardt, S. Levine (Eds.), Advances in Neural Information Processing Systems, volume 36, Curran Associates, Inc., 2023, pp. 68539–68551.
- [20] S. Peroni, D. Shotton, The spar ontologies, in: D. Vrandečić, K. Bontcheva, M. C. Suárez-Figueroa, V. Presutti, I. Celino, M. Sabou, L.-A. Kaffee, E. Simperl (Eds.), The Semantic Web – ISWC 2018, Springer International Publishing, Cham, 2018, pp. 119–136. doi:10.1007/978-3-030-00668-6_8.
- [21] A. Oelen, M. Y. Jaradeh, S. Auer, Orkg ask: a neuro-symbolic scholarly search and exploration system, arXiv preprint arXiv:2412.04977 (2024). doi:10.48550/arXiv.2412.04977.