

Linda Hammarbäck, Librarian, Biomedical Library, University of Gothenburg, Sweden,
linda.hammarback@ub.gu.se

Helen Sjöblom, Librarian, Biomedical Library, University of Gothenburg, Sweden,
helen.sjoblom@ub.gu.se

Sara Viklund, Librarian, Biomedical Library, University of Gothenburg, Sweden,
sara.viklund@ub.gu.se

Eva Hessman, Librarian, Biomedical Library, University of Gothenburg, Sweden
eva.hessman@ub.gu.se

Doctoral students’ reflection on Generative AI A librarian outlook

Abstract

This practice paper explores doctoral students’ perspectives on the use of generative AI and examines the role of librarians as mediators in this evolving field. Originating from a long-standing doctoral course on information literacy at the Biomedical Library, University of Gothenburg, the study integrates recent advancements in generative AI tools for summarizing and analyzing scientific articles. During the spring semester of 2024, students were introduced to these tools and tasked with incorporating them into their course assignments. This paper analyzes the reflections and insights from these assignments, in the context of their diverse backgrounds from various medical disciplines. The findings provide a nuanced understanding of the students’ views on generative AI and underscore the importance of librarians in facilitating the adoption of new technologies in academic research. The study concludes with reflections on the implications for future course content and the role of librarians in supporting research through technological advancements.

Keywords: generative AI, information literacy, doctoral students, Gartner’s hype cycle

Introduction

In this practice paper we take a closer look at doctoral students’ perspectives on the use of generative AI. We will also discuss the role that we, as librarians, can play as mediators through this rapidly evolving field.

The paper originates from a doctoral course on information literacy offered by librarians at the Biomedical Library at the University of Gothenburg. The course has been held consistently every semester since the 1970s for doctoral students at the Faculty of Medicine. Although the framework of the course remains stable, we make continuous improvements by adding new elements. In recent years, we have applied a broader approach to cover such phenomena, as the ever-changing landscape of publishing, open science, and research waste. We strive to keep the course content up to date and in line with current changes in information technology and scholarly practices. The latest addition to our course is a brief introduction to generative AI tools for summarizing and analyzing scientific articles, which we incorporated into one of our lectures during the spring semester of 2024.

The doctoral students were also tasked with testing generative AI tools in their course assignment. Upon assessing the submitted texts, we found the students' views interesting, and we wanted to gain a better understanding of their perspectives and insights. In this paper we share our analysis of the doctoral students' texts, and we also reflect on the role of librarians in introducing the use of generative AI in research.

The context

In this practice paper we analyze answers and reflections from the written assignments completed by newly admitted doctoral students during the spring of 2024. The students participated in the mandatory course titled *Introduction to Doctoral Studies and Information Retrieval* given at the Faculty of Medicine at the University of Gothenburg. The course had a mix of participants from various departments, e.g. nursing science, odontology, and medicine. It was a diverse group that included students from both clinical and pre-clinical research areas.

During the one-week full-time course, we, as librarians, provide nine hours of lectures and workshops. The library sessions cover scholarly publishing and research context, information retrieval, and information management. This paper focuses on the information management session, where we introduce reference management software and methods for summarizing key elements of the research published in your field. To complete the course, the students need to fulfill two written assignments: In assignment A they need to describe their search strategies, while Assignment B is a quantified overview of their research field. In this paper we use the submissions for assignment B as source for our analysis.

Assignment B consists of two parts. In part one, the students are asked to create a quantified overview of ten relevant articles in their research field. To compile this, students manually add relevant categories for data extraction (e.g. age, gender, outcomes, number of participants, methods). This process is inspired by Griffith university's method for systematic quantitative literature reviews (Griffith University, 2024). Part two was added to the assignment in spring 2024, incorporating generative AI tools for summarizing and analyzing scientific articles. We ask the students to try automated data extraction based on the same ten articles, but this time using one or more AI-tools.

Following the examination, we contacted all 76 course participants via email to request permission to use their submitted assignments for an in-depth analysis and to inform them of our intention to publish an article based on that analysis. We used the assignments from the students who consented (n=57) and downloaded them from the learning platform Canvas. The texts were then pseudonymized by removing all personal information such as email addresses and names. The texts were then uploaded in Taguette for analysis.¹

¹ <https://www.taguette.org/>

We also conducted a follow-up web survey four months after the course, consisting of multiple-choice questions about their current use of AI tools. The questionnaire was sent to all 57 doctoral students who had given us permission to use their text. Of those, 33 responded, resulting in a response rate of 58%.²

The analysis

To analyze the texts, we used a directed qualitative content analysis method (Hsieh and Shannon, 2005). We read and re-read the assignment submissions to gain understanding of the doctoral students' attitudes towards generative AI-tool usage in their research, as well as to identify patterns and correspondences within their reflections. We expected the analysis to show whether the doctoral students considered AI tools:

- a useful method
- a way to create accurate representations through summaries
- something they would continue to use in their studies

Since these presumptions guided our reading, we used a concept-driven process when we developed the themes later used in the resulting content analysis.

In presenting the results of our analysis, we felt the need to rephrase the themes. We wanted them to serve better as headings and to communicate representations of the content. Our themes were hence translated into *Usefulness*, *Trustworthiness* and *Future applications*. In our definition of the themes, we focused on the immediate use for research purposes, i.e. does the tool speed up the steps needed to make sound scientific discovery and analysis? Broader societal aspects turned up in student reflections every now and then, but those are not separately recognized or accounted for in our analysis. *Usefulness* is defined as helpful in speeding up reading and summarizing scientific articles. *Trustworthiness* is defined as providing an accurate representation of article content. *Future application* is the theme where we collected all hopes and concerns for the future development of scientific work and the research process, when or if AI tools are increasingly applied in research practice.

Below we present each theme as well as some excerpts from the student answers, which we exemplify according to three different approaches: Positive, negative and ambivalent.

Usefulness

Some quotes clearly illustrate the opinion that the tested tools can be of great use. Students expressed the immediate usefulness of a tool in terms of excellence.

The potential of these tools is unlimited.

Excellent initial screening process, allowing me to identify articles that demanded more attention.

It's not a replacement for reading everything, but it is a big help for getting started and deciding which papers are most important for my work.

The tool immediately gave me the accurate and summarized information I needed to develop my method which is the first stage of my lab work.

² Assignment B and the follow-up web survey can be accessed at <https://zenodo.org/doi/10.5281/zenodo.13474352>

Negative quotes expressed a lack of trust or experienced fuzziness of summaries and a need to double check the AI tool's interpretation. They also questioned the benefits of using an AI tool, as compared to doing the reading yourself, claiming that better understanding is gained in the traditional way, that is reading the articles yourself.

However, I feel that this filtering function also distances me from the text. When I read articles in their entirety or skim through them based on the original text, I experience that they build upon my understanding and expand my horizon of understanding, and automatically give me a sense of how the article positions itself in relation to other research.

I would still be skeptical to the AI summary and have to double check everything. I wonder, what is the added benefit to this vs. just reading the abstract? Maybe not that much?

Several doctoral students appeared hesitant to take a firm position. While they acknowledged that the tools could be helpful, they also expressed concerns that these tools might distort text comprehension or introduce errors. This ambivalence was reflected in their emphasis on the need for critical thinking and to have a balanced approach to these tools.

However, it's important to note that sometimes the tool occasionally flags the results with "Low confidence, Answer may not be accurate. Double check this source". This emphasizes the need to use it in a critical way and that to better understand, it is necessary to read the whole article more accurately. Furthermore, since it is an AI research tool, it may not be able to answer very specific or complex questions. It is always important to use it with a critical thinking.

Therefore, despite the convenience afforded by summarizing tools, I remain committed to engaging with the full text of each article to ensure a comprehensive understanding of the research presented. This balanced approach allows me to leverage the benefits of automation while prioritizing the thoroughness required for rigorous academic inquiry.

Trustworthiness

Trustworthiness is a complex concept incorporating elements such as accountability, transparency, and accuracy. These elements collectively form a framework for trustworthiness, fundamentally centered on earning and maintaining trust.

Looking at our data the main view regarding trustworthiness is one of ambivalence. Many doctoral students were impressed by the tools output but remained skeptical about how accurately they reflected the content of the articles.

I believe that, similarly to ChatGPT itself, those tools are very useful in research, but must be treated with a limited trust.

At first glance I thought the tools were incredible, but the more I worked with them, the more flaws were discovered.

In some quotes a concern of lack of context emerged, as well as a fear of missing information when only getting answers to the questions one asks.

... diving into them blindly can be misleading, leading to misinterpretation, oversimplification or missing complex details due to the presence of contextual information, complexity and nuanced language of research papers as these programs aren't perfect.

A danger I see with using the programs is that one easily misses pitfalls, such as weaknesses in the study design, when simply asking for the results one is interested in. It is difficult to evaluate for

oneself whether to trust the data provided by the AI, as it does not present information that is not asked for or assess the credibility of the information it provides. I found that I asked many control questions to the answers it gave.

The importance of transparency was emphasized, and there was a desire for more clarity from the tools – such as where in the articles the information was sourced from, and what exactly the analysis was based on.

I also would not really trust it because I did not see any way to verify the information based on the chat messages, such as a link to the relevant text. So I would still find that I have to actually read the article.

Future applications

Regarding future applications, most doctoral students felt cautiously optimistic. Many believed they would use similar generative AI tools during their doctoral studies. They also thought that more advanced models would likely be available in the future.

I believe that this way of processing articles to gain an overview, as Elicit offered, is something I will benefit greatly from during my doctoral studies – but I need to be continuously critical, reading the articles carefully in their entirety before including them in my thesis, and not solely rely on Elicit's summary.

We are in early days of AI and I am sure it will be a helpful tool for both finding and extraction information now and in the future. Who knows where we will be in 5 or 10 years. I don't think anyone could have anticipated what internet has done to/for our society, the next revolutionary technology may be in AI.

Several respondents also mentioned that AI tools cannot and should not completely replace their own reading, and that reading is crucial for developing the ability to express oneself.

Regarding these advancements, I'm approaching them with mixed emotions. On one hand they are exciting for their potential for a new paradigm shift in academic research and the greater good for society. But on the other hand, they also raise concerns, for instance potential decrease in critical reading, writing and ultimately thinking among researchers due to over-reliance on AI.

I think we should also consider potential negative consequences of relying on AI. For example, our ability to express our thoughts might be impacted.

Some doctoral students expressed concern about using AI in the wrong way and emphasized the importance of finding a balance in the interaction between researchers and AI.

I don't think I will be using AI in my research. I would be too afraid of using it in an incorrect way.

Finding the right balance by incorporating both automatization and human expertise is the soundest approach in advancing research.

Follow-up survey

In the follow-up survey conducted four months after the course, there was an overwhelming majority stating that they would continue to use these types of AI tools during their doctoral studies. 73% stated yes, 24% stated maybe. However, 40% had not used the tools at all since the course, and 58% reported using them less than once a week. We also inquired about the reasons behind their continued or discontinued use of AI tools. The three most common reasons for using AI tools in

research were: *It saves me time, It provides an overview of my research, and It helps me find core concepts*. The most prevalent answers to why they choose not to use AI tools were: *I don't find the tools to be reliable and I don't currently read many articles*.

Reflections

Doctoral students seem to be adopting generative AI tools with a sense of excitement and anticipation about their potential benefits. However, this enthusiasm is often accompanied by anxiety and distrust. The follow-up survey revealed that, despite students being confident they would use generative AI in their research, they had not actually done so.

To better understand these various stages of generative AI adoption, we find Gartner's Hype Cycle (figure 1) to be beneficial. The Gartner Hype Cycle is a model that describes the typical progression of new technologies. It begins with the *Technology Trigger* phase, where novel technologies are introduced. The next phase is the *Peak of Inflated Expectations*, characterized by "over-enthusiasm and unrealistic projections". This phase is followed by the *Trough of Disillusionment*, where "the technology does not live up to its overinflated expectations". Moving forward, the *Slope of Enlightenment* emerges – a phase marked by "a true understanding of the technology's applicability, risks, and benefits". Finally, the hype cycle reaches the *Plateau of Productivity*, where "real-world benefits of the technology are demonstrated and accepted" (O'Leary, 2008).




Figure 1. Gartner Hype Cycle. (2024, July 2). Wikimedia Commons. Retrieved August 28, 2024 from https://commons.wikimedia.org/w/index.php?title=File:Gartner_Hype_Cycle.svg&oldid=891097997

When examining the doctoral students' accounts of their interactions with generative AI tools, we can see signs that they are at different stages of the hype cycle. It is also clear that one person can be at multiple stages at the same time. Our analysis shows that most students move between the *Peak of Inflated Expectations* and the *Trough of Disillusionment*.

By linking their experiences to Gartner's Hype Cycle and viewing their behaviors through this lens, we can understand them better. This lens is also useful for us as librarians – by reflecting on our feelings about generative AI through Gartner's Hype Cycle, we can better understand our reactions. Just like the doctoral students, we librarians can be at multiple stages simultaneously. By reflecting on our own positions, we can better meet our users' needs and avoid getting stuck in doomsday or salvation mindsets.

This practice paper highlights the necessity for continuous exploration of the practical applications of generative AI in academic research. It is a rapidly evolving area, both in terms of tools available, practical applications, and regulations governing the use. Currently, generative AI is largely in a pre-regulation phase. Regarding the use of generative AI in writing, both publishers and funding bodies have started to establish guidelines, see for example ICMJE, 2024. However, the legal landscape surrounding the uploading of copyrighted materials, such as publishers' PDF versions of articles, is less clear. Discussions are ongoing, and several publishers are beginning to address this issue in their agreements and libraries need to adapt accordingly. Interpreting the terms of different AI tools regarding copyright material can be challenging. However, Elicit, a tool widely used by the doctoral students in the assignment, explicitly states that it does not use uploaded PDFs to train language models or share them further.³ We are currently considering whether we can proceed with the assignment using AI tools, as this might violate agreements between publishers and the library. This issue was unclear during the study and remains unresolved.

Looking back, we realize we had a unique opportunity to explore generative AI in the doctoral course the spring of 2024. The timing was perfect, as the tools were more accessible, free, and seen as innovative and impressive. Since then, many tools have shifted to more closed models, requiring fees and user registration. We also believe that as the use of generative AI becomes more widespread, people will have higher expectations regarding its capabilities.

In many Swedish universities Academic Language Support units provide guidance on using generative AI in writing, and Pedagogical Support units address concerns related to plagiarism and assessment. In our role as librarians, we have an opportunity to emphasize the role of AI tools in information retrieval and information management. We strongly recommend that the exploration and implementation of generative AI in academic research should be a collaborative journey involving librarians, library users, and other stakeholders. Together with doctoral students and researchers, we can investigate how these tools function and how they can be utilized in academic research. Each perspective contributes a vital piece of the puzzle.

Acknowledgments

This article was made possible by the financial support of *Stiftelsen Apotekare, filosofie doktor Gustaf Bernströms donationsfond*. Thank you to all the doctoral students in the course who were willing to share their reports and let us use them in this practice paper. You were great!

³ <https://support.elicit.com/en/articles/723521>

References

- Griffith University. (2024). *Systematic quantitative literature review*.
<https://www.griffith.edu.au/griffith-sciences/school-environment-science/research/systematic-quantitative-literature-review>
- Hessman, E., Sjöblom, H., Hammarbäck, L., & Viklund, S. (2024). *Dataset for: Doctoral students' reflection on Generative AI: A librarian outlook*. [Dataset] Zenodo.
<https://zenodo.org/doi/10.5281/zenodo.13474352>
- Hsieh, H.-F., & Shannon, S. E. (2005). Three approaches to qualitative content analysis. *Qualitative Health Research*, 15(9), 1277-1288. <https://doi.org/10.1177/1049732305276687>
- ICMJE International Committee of Medical Journal Editors. (2024). *Recommendations for the conduct, reporting, editing, and publication of scholarly work in medical journals*.
<https://www.icmje.org/recommendations/browse/roles-and-responsibilities/defining-the-role-of-authors-and-contributors.html#four>
- Mayring, P. (2015). Qualitative content analysis: Theoretical background and procedures. In: A. Bikner-Ahsbals, C. Knipping, & N. Presmeg (Eds.), *Approaches to qualitative research in mathematics education: Advances in mathematics education* (pp. 365-380). Springer.
https://doi.org/10.1007/978-94-017-9181-6_13
- O'Leary, D. E. (2008). Gartner's hype cycle and information system research issues. *International Journal of Accounting Information Systems*, 9(4), 240-252.
<https://doi.org/10.1016/j.accinf.2008.09.001>