



Adoption of Generative AI by Academic Biomedical Researchers

Dylan Ruediger
Chelsea McCracken
Makala Skinner

October 17, 2024



ITHAKA S+R

Ithaka S+R provides research and strategic guidance to help the academic and cultural communities serve the public good and navigate economic, demographic, and technological change.

Ithaka S+R is part of ITHAKA, a nonprofit with a mission to improve access to knowledge and education for people around the world. We believe education is key to the wellbeing of individuals and society, and we work to make it more effective and affordable.

Copyright 2024 ITHAKA. This work is licensed under a Creative Commons Attribution 4.0 International License. To view a copy of the license, please see <https://creativecommons.org/licenses/by-nc/4.0/>.



Preface

Understanding how biomedical researchers are making use of generative AI is critical to informed decision-making about how to support ethical adoption of the technology and assessing the risks and opportunities it presents to the research enterprise.

Biomedical research has been at the forefront of generative AI-enhanced research. Generative AI's contributions to drug development and protein design are among the most widely celebrated concrete examples of its transformative potential. Biomedicine has also been at the forefront of developing customized, domain-specific large language models (LLMs). It is also a field in which any accelerating effects enabled by generative AI would have immediate impacts on the health of individuals, and for the same reason, where errors created by generative AI have significant potential to cause harm.

Understanding how biomedical researchers are making use of generative AI is critical to informed decision making about how to support ethical adoption of the technology and assessing the risks and opportunities it presents to the research enterprise. However, most studies of the use of generative AI by academic researchers have cast a wide net rather than focusing on adoption in specific disciplines or domains.

Ithaka S+R, with support from the Chan Zuckerberg Initiative (CZI), conducted a survey of biomedical researchers to better understand how researchers in the field think about and use generative AI in their research. The survey yielded a treasure trove of data, and we are excited to share findings from the survey in this report.

Introduction

The commercial release of ChatGPT in November 2022 brought generative AI to the attention of both the general public and most academic researchers. The response within higher education communities has been overwhelming, as reflected in workgroups at universities across the country, uncountable conference sessions and entire conferences dedicated to generative AI, and what seems like an even greater number of opinion pieces, position papers, preprints, and journal articles. New generative AI tools designed for student and faculty use are released every week.¹

At colleges and universities in the United States, discussions about generative AI have been driven by issues related to teaching and learning. The implications of generative AI on the academic research enterprise have received significantly less attention, though they have steadily gained visibility since the early months of 2023, when we began talking with senior administrators about the strategic implications of generative AI for higher education. The literature on generative AI as a research tool has also grown steadily, even as it still lags behind literature on generative AI in teaching and learning.

Generative AI has significant potential to accelerate scientific discovery, and is already making significant impacts on fields such as protein design and drug development.² Its potential to improve search and discovery, synthesize information, and speed the writing process are cross-field use cases that could make engaging with and contributing to the scholarly record less time consuming and more efficient.³ Researchers have

¹ Claire Baytas and Dylan Ruediger, "Generative AI in Higher Education: The Product Landscape," *Ithaca S+R*, 7 March 2024, <https://doi.org/10.18665/sr.320394>.

² Adam Winnifrith, Carlos Outeiral, and Brian L. Hie, "Generative Artificial Intelligence for de Novo Protein Design," *Current Opinion in Structural Biology* 86 (June 1, 2024); Laura Howes, "Generative AI Is Dreaming up New Proteins," *Chemical & Engineering News* 101, no. 12 (April 10, 2023); Yuemin Bian and Xiang-Qun Xie, "Generative Chemistry: Drug Discovery with Deep Learning Generative Models," *Journal of Molecular Modeling* 27, no. 3 (February 4, 2021): 71, <https://doi.org/10.1007/s00894-021-04674-8>.

³ Markus J. Buehler, "Accelerating Scientific Discovery with Generative Knowledge Extraction, Graph-Based Representation, and Multimodal Intelligent Graph Reasoning," *Machine Learning: Science and Technology*, 2024, <https://doi.org/10.1088/2632-2153/ad7228>; Adhari AlZaabi, Amira ALAmri, Halima Albalushi, Ruqaya Aljabri, and

expressed hope that generative AI's ability to find patterns in large datasets could unearth causal relationships that humans would miss and propose novel hypotheses.⁴ Radical frameworks for fully automated scientific discovery via generative AI-powered tools that generate research ideas, write necessary code, conduct experiments, and create visualizations and written outputs to communicate findings are beginning to emerge.⁵

All of these use cases have skeptics, who have a prominent voice in discussions about the use of generative AI in scientific research. The most common criticisms include perceptions that the quality of the outputs from existing generative AI tools is not good enough to be useful in specialized research contexts; concerns about its social or environmental repercussions; and the risks that it poses to research ethics and the integrity of scholarly record. What critics and proponents typically agree on is that there is no way to put generative AI back in its bottle. Its impacts, good or ill, seem almost certain to be consequential, and adoption seems inevitable.⁶

AbdulRahman AalAbdulsalam, "ChatGPT Applications in Academic Research: A Review of Benefits, Concerns, and Recommendations," *bioRxiv*, 18 August 2023, <https://doi.org/10.1101/2023.08.17.553688>; Steven A. Lehr, Aylin Caliskan, Suneragiri Liyanage, and Mahzarin R. Banaji, "ChatGPT as Research Scientist: Probing GPT's Capabilities as a Research Librarian, Research Ethicist, Data Generator and Data Predictor," *arXiv*, 20 June 2024, <https://doi.org/10.48550/arXiv.2406.14765>; Chris Berg, "The Case for Generative AI in Scholarly Practice," SSRN Scholarly Paper, Rochester, NY, 3 April 2023, <https://doi.org/10.2139/ssrn.4407587>.

⁴ Hector Zenil, Jesper Tegnér, Felipe S. Abrahão, Alexander Lavin, Vipin Kumar, Jeremy G. Frey, Adrian Weller, et al, "The Future of Fundamental Science Led by Generative Closed-Loop Artificial Intelligence," *arXiv*, 29 August 2023, <https://doi.org/10.48550/arXiv.2307.07522>.

⁵ Chris Lu, Cong Lu, Robert Tjarko Lange, Jakob Foerster, Jeff Clune, and David Ha, "The AI Scientist: Towards Fully Automated Open-Ended Scientific Discovery," *arXiv*, 15 August 2024, <https://doi.org/10.48550/arXiv.2408.06292>.

⁶ Henrik Skaug Sætra, "Generative AI: Here to Stay, but for Good?" *Technology in Society* 75 (November 1, 2023): 102372, <https://doi.org/10.1016/j.techsoc.2023.102372>; Luke Tredinnick and Claire Laybats, "Black-Box Creativity and Generative Artificial Intelligence," *Business Information Review* 40, no. 3 (September 1, 2023): 98–102, <https://doi.org/10.1177/02663821231195131>; Yogesh K. Dwivedi, Nir Kshetri, Laurie Hughes, Emma Louise Slade, Anand Jeyaraj, Arpan Kumar Kar, Abdullah M. Baabdullah, et al, "'So What If ChatGPT Wrote It?' Multidisciplinary Perspectives on Opportunities, Challenges and Implications of Generative Conversational AI for Research, Practice and Policy," *International Journal of Information Management* 71 (August 1, 2023): 102642, <https://doi.org/10.1016/j.ijinfomgt.2023.102642>.

A handful of studies are beginning to provide glimpses into how researchers are responding to and making use of generative AI. Last September, *Nature* published findings from a global survey of 1,600 researchers from across disciplines that found high levels of excitement and concern, and modest levels of use of generative AI for research purposes.⁷ Noting that “few studies have been published on how researchers are using AI,” *Nature* conducted a second survey—focused on postdocs—designed to elicit specific data about how generative AI was being used. They found that postdocs most often used AI to edit manuscripts, write code, and to find or synthesize journal articles. Among the 31 percent of postdocs who indicated that they used generative AI, 60 percent did so at least weekly.⁸

The broad and diverse domain of biomedical research has emerged as an important leader in generative AI.

The broad and diverse domain of biomedical research has emerged as an important leader in generative AI. Biomedicine has also been at the forefront of developing customized, domain specific large language models (LLMs) like BioMedLM, BioLinkBERT.⁹ Domain specific data are a necessary precondition for fostering ethical adoption of generative AI within biomedical communities. However, most research on AI use has cast a wide net: there is little available literature on discipline or domain specific adoption. To this end, our survey sought to uncover how researchers in the field think about and use generative AI in their

⁷ Richard Van Noorden and Jeffrey M. Perkel, “AI and Science: What 1,600 Researchers Think,” *Nature* 621, no. 7980 (September 27, 2023): 672–75, <https://doi.org/10.1038/d41586-023-02980-0>.

⁸ Linda Nordling, “How ChatGPT Is Transforming the Postdoc Experience,” *Nature* 622 (2023): 655–57, <https://pubmed.ncbi.nlm.nih.gov/37845528/>.

⁹ Patrick Lewis, Myle Ott, Jingfei Du, and Veselin Stoyanov, “Pretrained Language Models for Biomedical and Clinical Tasks: Understanding and Extending the State-of-the-Art,” In *Proceedings of the 3rd Clinical Natural Language Processing Workshop*, edited by Anna Rumshisky, Kirk Roberts, Steven Bethard, and Tristan Naumann, 146–57, Online: Association for Computational Linguistics, 2020, <https://doi.org/10.18653/v1/2020.clinicalnlp-1.17>; Benyou Wang, Qianqian Xie, Jiahuan Pei, Zhihong Chen, Prayag Tiwari, Zhao Li, and Jie Fu, “Pre-Trained Language Models in Biomedical Domain: A Systematic Survey,” *ACM Computing Surveys* 56, no. 3 (October 5, 2023): 1-52, <https://doi.org/10.1145/3611651>.

research. Our findings cover four key areas: a) attitudes about generative AI among biomedical researchers, b) barriers to adoption, c) use cases, and d) support needs.

Our core finding is that generative AI has reached a critical juncture. Biomedical researchers are open to using it for research, and experimenting with ways to do so, but very few are making extensive use of it. Adoption is limited by serious concerns about the accuracy of generative AI's outputs and uncertainty about how to use generative AI productively. Absent continued improvement in the quality of generative AI outputs and the emergence of compelling best practices and models for using generative AI to make researchers more productive, adoption may plateau, at least in the short term.

Several individual findings warrant particular attention:

- Biomedical researchers have a moderate degree of interest in using generative AI in their research. Over 60 percent have experimented with doing so, but most use it sparingly or no longer use it at all.
- There are many barriers to greater adoption of generative AI, but the most significant are concerns about the accuracy of generative AI and a lack of clarity about best practices for using AI effectively and ethically.
- Over half of biomedical researchers expressed strong interest in biomed specific generative AI products, but only 14 percent had used existing biomed specific LLMs or tools.
- Use of generative AI in research is concentrated on scholarly communication tasks such as editing, writing, and accessing and interpreting scholarly literature.
- Many researchers would appreciate more support from funders, publishers, and universities to develop their skills using generative AI in their research.

Our overriding impression is that generative AI is at something of a crossroads. Biomedical researchers are open to using it for research and are experimenting with ways to do so. However, they are doing so cautiously: few are committed users. One upside of the limited commitment that biomedical researchers have to generative AI is that cultural and methodological norms about its use are not yet established.

Our overriding impression is that generative AI is at something of a crossroads. Biomedical researchers are open to using it for research and are experimenting with ways to do so. However, they are doing so cautiously: few are committed users.

Despite the sense that generative AI is a runaway train, there is still time for deliberation, consensus building, and coordinated action by stakeholders to mitigate ethical challenges and maximize its contributions to science. Intervening after norms and practices have been hardened by constant use will be much more difficult than doing so while they are still fluid.

Methods

This report is based on a survey of academic researchers regarding their attitudes towards and use of generative AI in research contexts. The survey was conducted between February 20 and March 29, 2024. The survey instrument was developed by Ithaka S+R with input from CZI. We conducted semi-structured exploratory interviews with five biomedical researchers to help surface themes and identify use cases as part of the process of creating the instrument. The instrument was tested and refined with an additional five cognitive interviews with different individuals.

The primary survey population was faculty members at four-year postsecondary institutions in the United States, and the primary recruitment was via direct email using a customized national contact list of faculty collected and maintained by a third-party vendor.

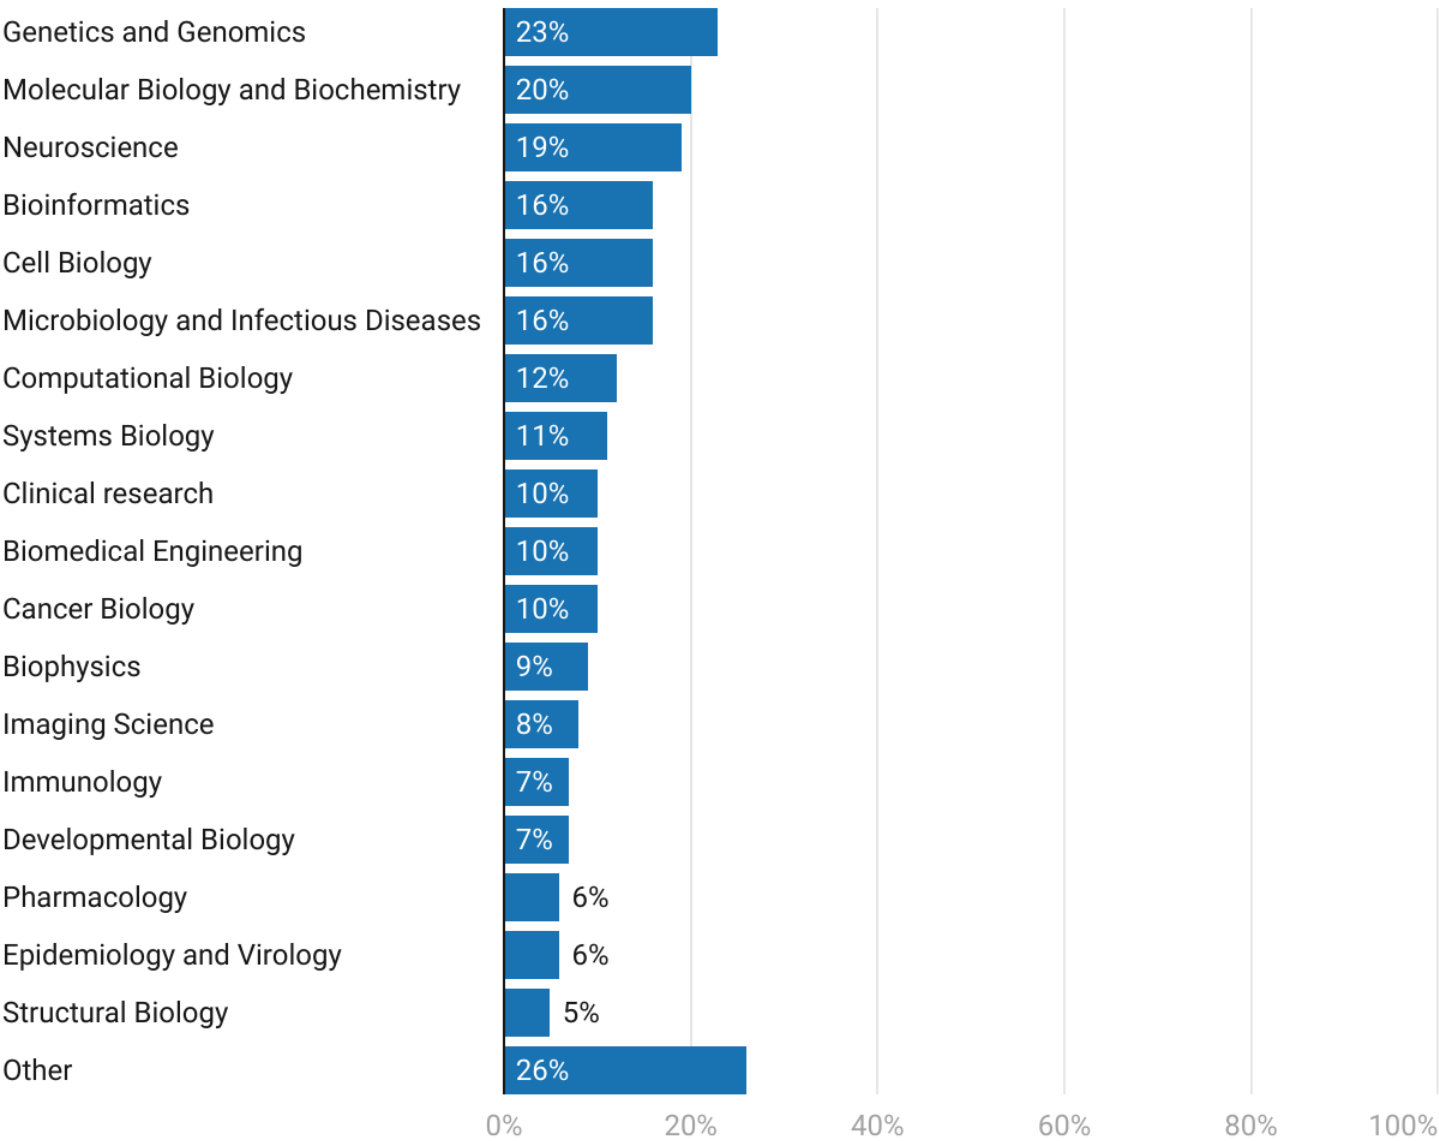
Because our primary goal was to understand how biomedical researchers are using generative AI in their research, 27 percent of the total sample were individuals working in biomedical fields. We used the CIP (Classification of Instructional Programs) code taxonomy developed by the National Center for Education Statistics to define which fields and areas of study we included in our biomedical sample. To yield comparative data, we also solicited responses from 65,875 researchers working in other

disciplines. In total, the survey was fielded to a sample of 90,091 faculty members.

All respondents were asked a series of questions about their attitudes towards and use of generative AI in research contexts. We provided the following definition of generative AI to respondents: “Generative AI” is defined as AI models that have been trained on existing data and can create (“generate”) original content (e.g., text, images, code) from that data, for example, ChatGPT, Midjourney, Google Bard, BioMedLM, BioLinkBERT, Dragon, etc.” Individuals who identified themselves as biomedical researchers were asked to answer several additional questions that addressed how they used generative AI in research contexts in greater detail. At several points, respondents were provided with the opportunity to make open-ended comments. We have included some of these in this report.

We received 2,459 complete, valid, survey responses. Biomedical researchers comprised 770 (31 percent) of respondents. As detailed in Figure 1, respondents represented a wide range of disciplines and fields within the biomedical domain.

Figure 1. “You indicated that you study or conduct research in biomedical and life sciences. In which of the following specific areas do you study or conduct research? (Select all that apply).”¹⁰



Created with Datawrapper

¹⁰ The majority of those describing their specialty as “other” indicated that they were ecologists or in adjacent fields.

The remaining 1,689 respondents represented a wide range of disciplinary backgrounds: 25 percent worked in the social sciences, 17 percent in the humanities, 12 percent in the physical sciences, 9 percent in business, and 6 percent in the visual and performing arts. Our respondents were primarily established academic researchers with a median of 20 years of experience. Seventy percent were tenured or tenure-track. Nearly all (96 percent) worked for academic institutions. Nearly all responses (98 percent) were from individuals whose primary residence is in the United States.

The overall margin for error is 2 percent for the total sample, and 3 percent for the biomedical sample.


Finding Area One: Attitudes Towards Generative AI By Biomedical Researchers

As is the case across disciplines, the workflows of biomedical researchers are beginning to incorporate generative AI. However, interest in adopting generative AI as a research tool exceeds actual usage.

Familiarity with generative AI

The plurality of academic researchers (42 percent) described themselves as moderately familiar with generative AI in general, with 25 percent describing themselves as very or extremely familiar, and 28 percent as slightly familiar. Biomedical researchers' responses were similar to those of the overall survey population. Forty-three percent described themselves as moderately familiar with generative AI, 30 percent as slightly familiar, and 20 percent as either very or extremely familiar with generative AI (Figure 2).


Figure 2. “In general, how familiar are you with generative AI?”



Created with Datawrapper

Biomedical researchers expressed lower levels of familiarity with research applications of generative AI compared with their familiarity of generative AI broadly. Over 20 percent described themselves as not at all familiar with such applications, and an additional 37 percent as only slightly familiar: only 11 percent indicated that they were very or extremely familiar with applying generative AI to their research. Here again, biomedical researchers’ responses are nearly identical to the overall survey population. The gap between general knowledge of generative AI and specific understanding about how to use it in research is a common pattern shared broadly by academic researchers—including those working in biomedical fields (Figure 3).

Figure 3. “In general, how familiar are you with how generative AI can be applied in research settings?”



Created with Datawrapper


Interest in generative AI as a research tool

Levels of interest in using generative AI vary considerably across disciplines. Interest in using generative AI in research among business researchers is outpacing their colleagues in other domains: just 6 percent of business researchers had no interest in using generative AI, and 22 percent were extremely interested in doing so. At the extreme other end of the spectrum, 31 percent of humanities researchers are not interested in

using generative AI, and only 7 percent were extremely interested. Biomedical researchers ranked among the most interested potential adopters. Only 13 percent of biomedical researchers have no interest in using generative AI in their research, and about a third are either extremely or very interested in adopting generative AI into their workflows.

Overall, biomedical researchers' interest levels are comparable to those expressed by social scientists and physical scientists, and well above those of researchers in the performing/visual arts or humanities (Figure 4). Just over six in ten biomedical researchers are at least moderately interested in generative AI, which suggests that there is ample room for growth in adoption over time.

Figure 4. “How interested are you in using generative AI in your own research?”




Created with Datawrapper

Adoption of generative AI

Biomedical researchers' interest in generative AI is also reflected by clear indications that they have experimented with adoption. Sixty-three percent of biomedical researchers have used generative AI in their research (Figure 8). However, much of this use appears to be exploratory and sporadic. When asked if they are currently using generative AI for research, 60 percent of biomedical respondents said no, and only 7 percent indicated that they used it regularly (Figure 5). Clearly, curiosity and experimentation around generative AI is not yet translating into full adoption. Instead, biomedical researchers are proceeding cautiously.

Across sectors, generative AI has often been understood as an inevitable component of the future of work. This may well prove true, but data points such as these are a useful reminder that in biomedical fields (and other academic domains), the overwhelming majority of researchers are far from committed.

Figure 5. “Do you use generative AI in your own research?”



Created with Datawrapper

Finding Area Two: Barriers to Adoption

A wide range of practical issues and ethical dilemmas are contributing to researchers' reluctance to integrate generative AI into their research workflows.

Social and research ethics

The potential harms of generative AI are widely debated, with critics raising important concerns about biases in training data, the degradation of the information and physical environments, and more. Academic researchers are clearly following these conversations, and ethical concerns over the technology are a major factor limiting adoption. For example, roughly six in ten survey respondents expressed moderate to high levels of concern that generative AI would exacerbate existing social biases or inequities. An even greater number considered data security and privacy risks as significant challenges. These kinds of objections to generative AI have a dual character: they can be taken as general skepticism about AI's impact on society or as narrower methodological concerns about using AI as a tool for scholarly research. To put this another way, it is important to differentiate between ethics and *research* ethics, even while acknowledging the considerable overlap between them.

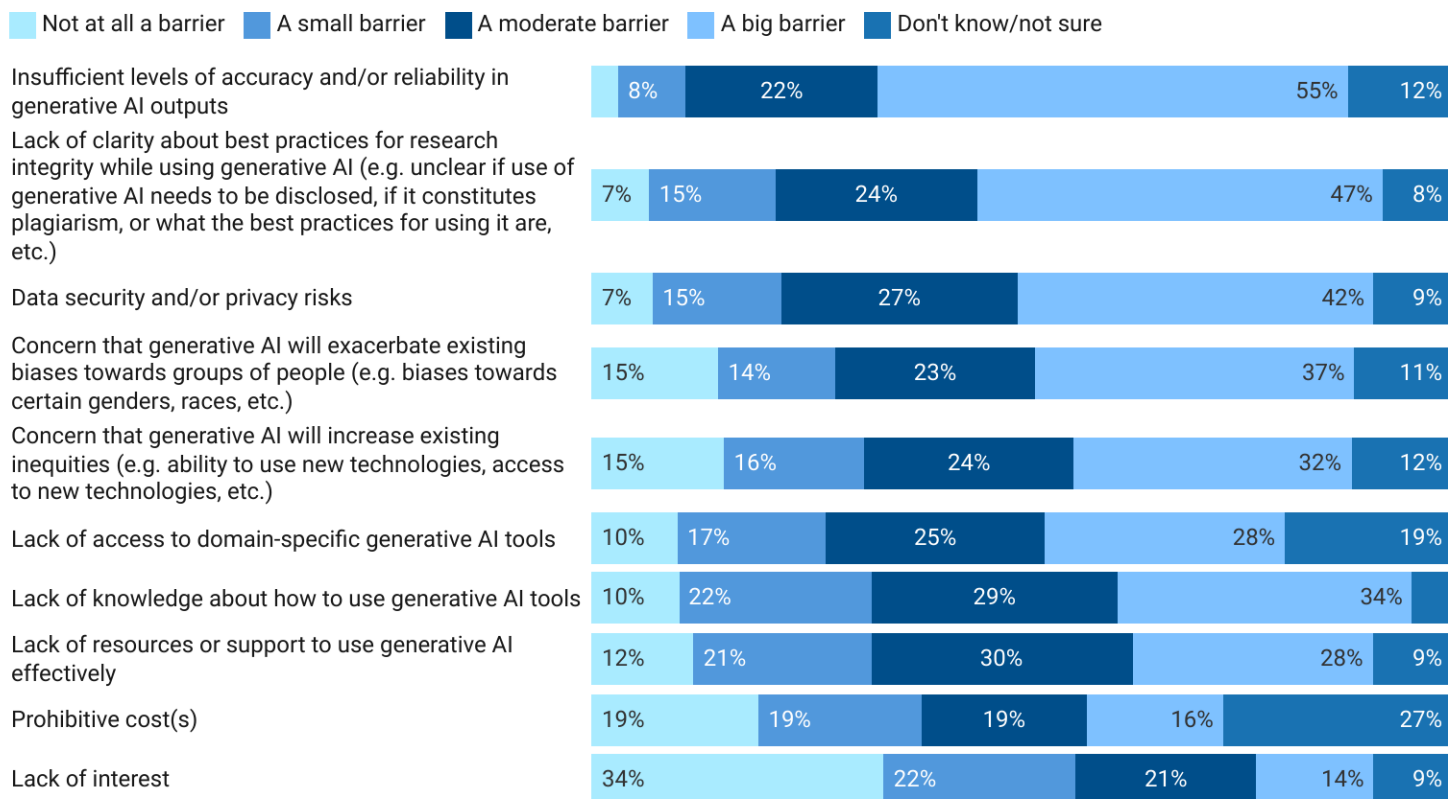
Roughly six in ten survey respondents expressed moderate to high levels of concern that generative AI would exacerbate existing social biases or inequities.

Concerns about research ethics and the integrity of the scholarly record are clearly a factor in many researchers' reluctance to adopt generative AI. At least some of the concerns about biases in LLM training data, for example, are likely due to their potential for distorting research findings. Data security and the protection of privacy are central issues for internal review boards.

It appears that many researchers would be more willing to use generative AI if they felt better equipped to navigate research integrity issues. Forty-five percent of respondents described uncertainty about best practices for research integrity as a major barrier to adoption of generative AI, and another 27 percent said that it was a moderate barrier.

Concerns about research ethics and the integrity of the scholarly record are clearly a factor in many researchers' reluctance to adopt generative AI.

Figure 6. “For you personally, to what extent are the following barriers to incorporating generative AI into your own research?” (Biomed respondents only)



Created with Datawrapper

Biomedical researchers reported similar levels of concern over ethics issues related to social biases, inequalities, and privacy risks as their colleagues in other disciplines. They also reported a need for better clarity about how to use generative AI in ways that conform to research ethics. Similar to the broader researcher population, 47 percent of biomedical researchers described uncertainty about best practices for research integrity as a major challenge, and another 24 percent described it as a moderate barrier to adoption.


Accuracy and reliability

The single biggest barrier to greater adoption of generative AI in research contexts is distrust of the accuracy and reliability of generative AI outputs. Seventy-four percent of researchers cited this as a moderate to large barrier, and just 4 percent did not see it as a barrier at all (Figure 5).¹¹

Accurate, reliable information is the *sine quo non* of scientific research, so it is reassuring to see researchers taking a cautious approach to generative AI while the extent of LLMs' inaccuracies and the value of solutions to them are still subject of considerable debate. However, researchers from different fields had differing opinions about exactly how significant 'hallucinations' and other LLM errors were. On one extreme, a relatively small 38 percent of business researchers—the field that has both the highest interest in generative AI and reported the most use of it—regarded accuracy as a major barrier. In contrast, 59 percent of physical scientists and 57 percent of humanists were also deeply concerned about the accuracy of generative AI, which presumably contributes to their high levels of skepticism about the tools in general (Figure 7). Biomedical researchers had higher than average levels of concern about the accuracy of generative AI and will clearly need to see progress on this front before fully integrating generative AI into their research.

¹¹ Richard Van Noorden and Jeffrey M. Perkel, "AI and Science: What 1,600 Researchers Think," *Nature*, 10 October 2023, <https://www.nature.com/articles/d41586-023-02980-0>, also found high levels of concern about issues related to accuracy, though the wording of their question makes direct comparison impossible.


Figure 7. “For you personally, to what extent are the following barriers to incorporating generative AI into your own research? Insufficient levels of accuracy and/or reliability in general AI outputs.”



Created with Datawrapper

Concerns about accuracy also varied by the age of respondents. The possibility that openness to using generative AI is, in part, generational is discernible in our findings. Older faculty (aged 60+) were less likely to use generative AI, and more likely to cite a lack of familiarity with generative AI as a barrier to adoption than were their younger colleagues. On the issue of accuracy, though, older researchers expressed fewer concerns than younger researchers (Figure 8). This is partially explained by the fact that 60+ respondents were almost twice as likely as respondents aged 18 to 44 to say they didn’t know whether insufficient levels of accuracy in generative AI outputs were a barrier, possibly indicating less knowledge and greater uncertainty about generative AI generally.

Figure 8. “For you personally, to what extent are the following barriers to incorporating generative AI in your own research? Insufficient levels of accuracy and/or reliability in general AI outputs.” Breakdown of concerns about accuracy by age.



Created with Datawrapper

Finding Area Three: Generative AI Use Cases in Biomedical Research

Biomedical researchers have experimented with using generative AI for a range of purposes but do not yet have a clear sense of how to use it effectively.

Specific use cases

Knowing that biomedical researchers are using generative AI is useful information for administrators, funders, scholarly societies, publishers, and other organizations or individuals who are involved in the research enterprise. Leveraging that knowledge in support of scientific discovery depends on understanding how generative AI is being used. To this end, we asked biomedical researchers about whether and how often they have used generative AI to assist with 15 different research tasks spread across the research lifecycle. We also asked them to evaluate the effectiveness of generative AI for each of the 15 use cases.

We found that biomedical researchers are using generative AI for a wide range of purposes ranging from discovery to experimentation to preparing

manuscripts for publication. However, the greatest current use of generative AI is somewhat concentrated at what might be considered as the beginning and end points of a research project: discovering and understanding relevant scholarly literature and editing manuscripts for publication.¹² Administrative tasks only marginally connected to research are another common use case (Figure 9).

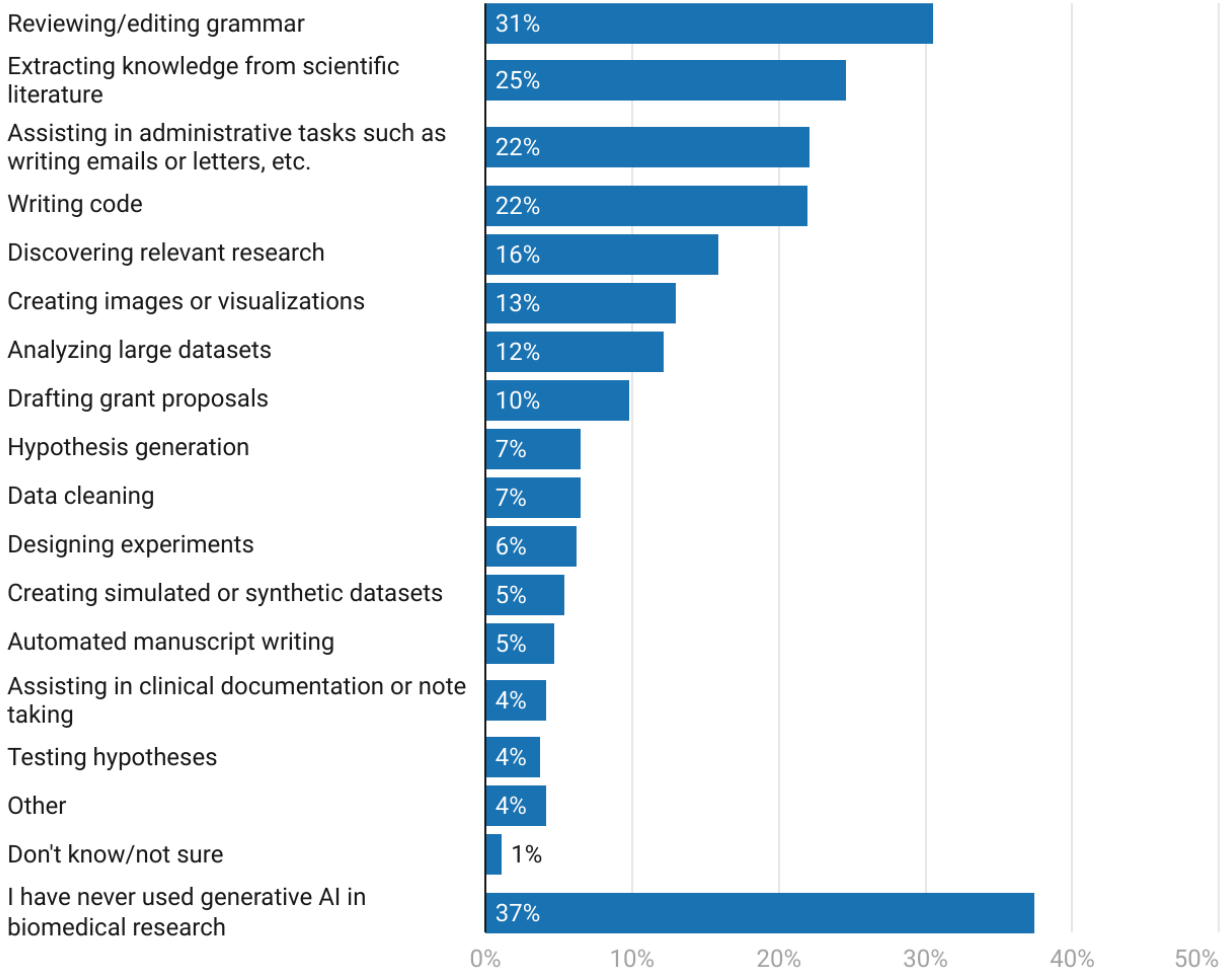
Nearly a third (31 percent) of biomedical researchers have used generative AI as an editorial tool to review and improve grammar, making it the most common use case for generative AI in the field. Other writing related use cases include administrative tasks such as email (22 percent). Relatively few respondents reported more intensive usage of generative AI as a writing aid—just 5 percent indicated having experimented with automated manuscript writing, while 10 percent had used it to draft grant proposals.

Biomedical researchers are also making relatively frequent use of generative AI to discover and interpret scholarly literature.

Biomedical researchers are also making relatively frequent use of generative AI to discover and interpret scholarly literature. Twenty-five percent of biomedical researchers had used generative AI tools to extract knowledge from scientific research, making it the second most widespread use case. Likewise, 16 percent had used it to assist with discovery of relevant resources such as journal articles. Generative AI-assisted discovery and synthesis has been a high priority for large scholarly publishers, who have integrated AI tools onto their platforms over the past year. While we did not gather information about which exact tools researchers are using for discovery, synthesis, or extraction, it seems likely that the rapid release of special purpose search, discovery, and synthesis tools has lowered barriers to using generative AI in this way (Figure 9).

¹² A large-scale survey conducted by Oxford University Press reached the same conclusion; see “Researchers and AI: Survey Findings,” <https://corp.oup.com/news/how-are-researchers-responding-to-ai/>. For an overview of generative AI tools designed for academic research, see Claire Baytas and Dylan Ruediger, “Generative AI in Higher Education: The Product Landscape,” *Ithaka S+R*, 7 March 2024, <https://doi.org/10.18665/sr.320394>.

Figure 9. “In which of the following ways have you used generative AI in your biomedical research?”



Created with Datawrapper

In contrast, generative AI has made only modest inroads into the middle, experimental, phase of research. Only 7 percent of biomedical researchers had generated hypotheses using generative AI, and fewer still had used it to design experiments (6 percent), test hypotheses (4 percent), or generate simulated or synthetic data (5 percent). A somewhat larger number of biomedical researchers have used generative AI to analyze data, but overall, few biomedical researchers are using generative AI to conduct original research (Figure 8).

Biomedical researchers' tendency to use generative AI primarily at the beginning and the end of the research process seems well aligned with how their peers in other disciplines are using it. A fall 2023 survey conducted by Nature found that roughly 15 percent of researchers who used generative AI had used it to create graphics or pictures: 13 percent of biomedical researchers in our survey reported having done so. Both Nature's global sample and our biomedical sample found that approximately 20 percent of researchers were using generative AI for administrative tasks such as email, and roughly one in three researchers had used generative AI to help draft manuscripts.¹³


Frequency of use

It is important to reiterate that the number of biomedical researchers using generative AI for any specific research purpose is dwarfed by the number of non-users, who comprised between 69 percent and 96 percent of respondents for each use case (median = 90). In contrast, the percentage of respondents who regularly use AI for a given task ranges from 15 percent to 1 percent (median = 2) of biomedical researchers. Power users are rare. A few specific tasks show moderate usage: 25 percent of those who used generative AI for reviewing/editing grammar did so sometimes or regularly, as did 17 percent of those using it to extract knowledge from scientific literature, and 16 percent using it to either write code or for administrative tasks. However, many of the tasks we asked about are almost never supported by generative AI, with 90 percent or more of respondents indicating that they never used generative AI for them (Figure 10).¹⁴

¹³ Richard Van Noorden and Jeffrey M. Perkel, "AI and Science: What 1,600 Researchers Think," *Nature* 621, no. 7980 (September 27, 2023): 672–75, <https://doi.org/10.1038/d41586-023-02980-0>.

¹⁴ Nature's survey found that among those who had used generative AI, just 13 percent did so more than once per week.

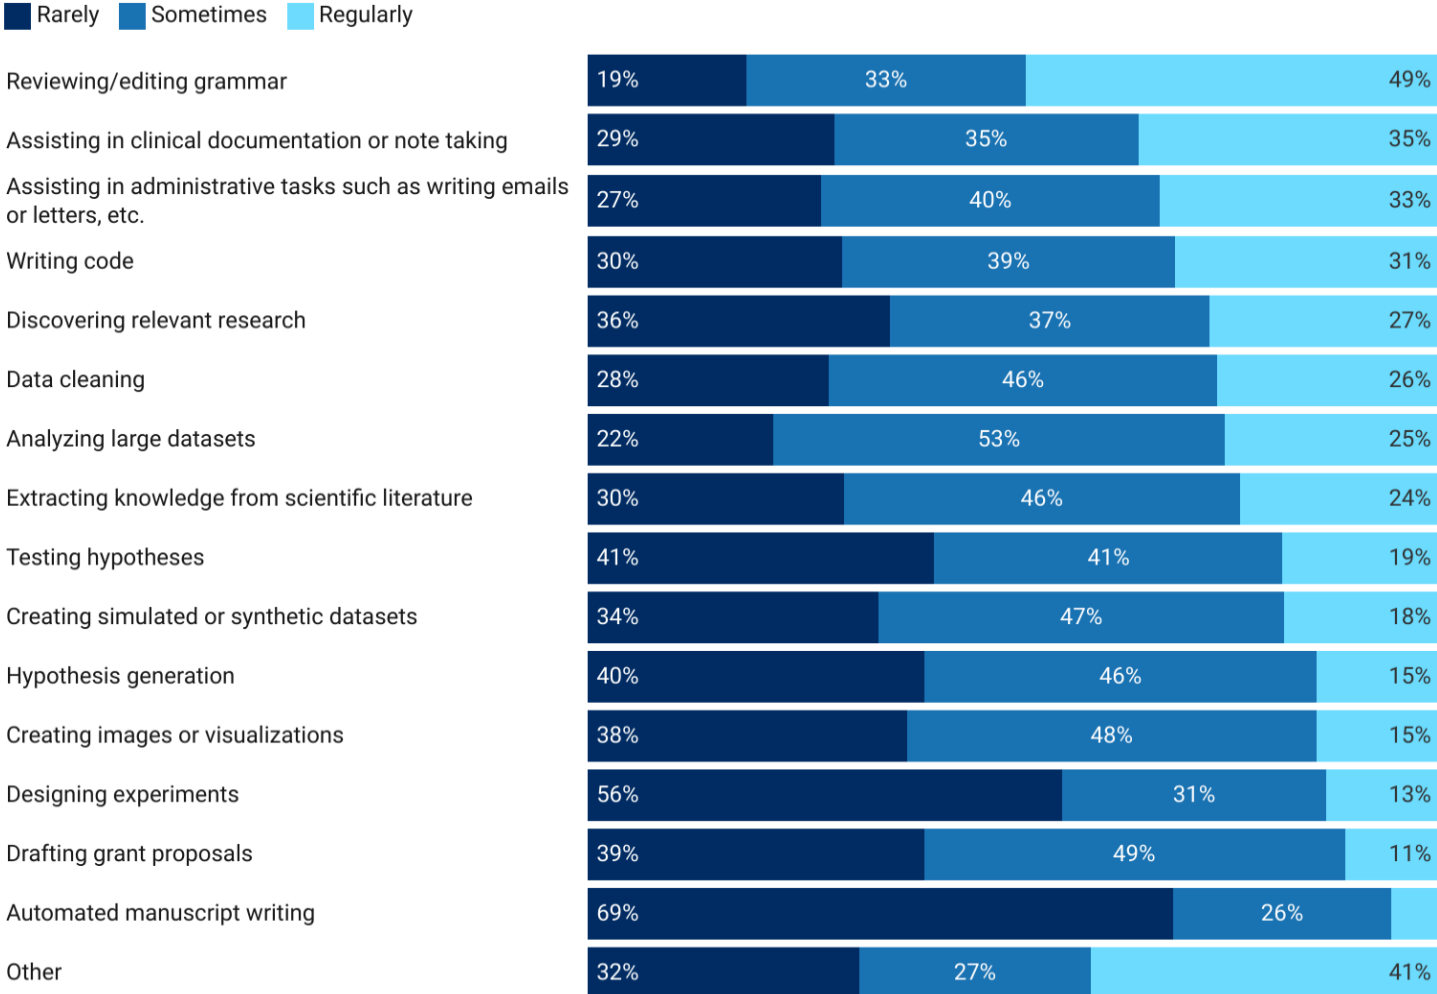
Figure 10. “How frequently do you use generative AI on the following tasks?” (Biomed respondents only)



Created with Datawrapper

There are, however, a few use cases that appear to be developing committed power users. Most notably, among those who had used generative AI for reviewing grammar, 49 percent did so regularly. Around a third of respondents using generative AI to write code, take clinical notes, or for administrative purposes were also regular users. With two exceptions (automated manuscript writing and designing experiments), the majority of biomedical researchers who used generative AI for any specific research task reported doing so either sometimes or regularly (Figure 11).

Figure 11. “You indicated that you have used generative AI on the following research task(s). How frequently do you use generative AI on the following task(s).” (Biomed respondents only)




Created with Datawrapper

Biomedical researchers expressed mixed opinions about the quality of generative AI as a research tool, typically rating their performance as fair to good. The most highly regarded use case for generative AI, reviewing and editing grammar, is also the one that users report using most frequently. Twenty-eight percent of those who used generative AI in this way rated it as producing very good quality results, with another 46 percent rating them as good. In contrast, majorities who had used generative AI for automated manuscript writing, drafting grant proposals, hypothesis generation, and clinical notetaking judged the results poor or

fair. For most tasks, there is no apparent consensus as respondents were split on the quality of results (Figure 12). There are increasing reports from the private sector suggesting that the hype around generative AI has worn off as the limits of the technology become apparent. A similar pattern could be manifesting among biomedical researchers who have tried generative AI: pluralities of researchers rated generative AI's performance as fair to very poor at tasks related to search and discovery, data cleaning and analysis, for example (Figure 12).

Figure 12. “When you use generative AI on the following biomedical research task(s), what was the quality of the work that the generative AI produced?” (Biomed respondents only)



Created with Datawrapper

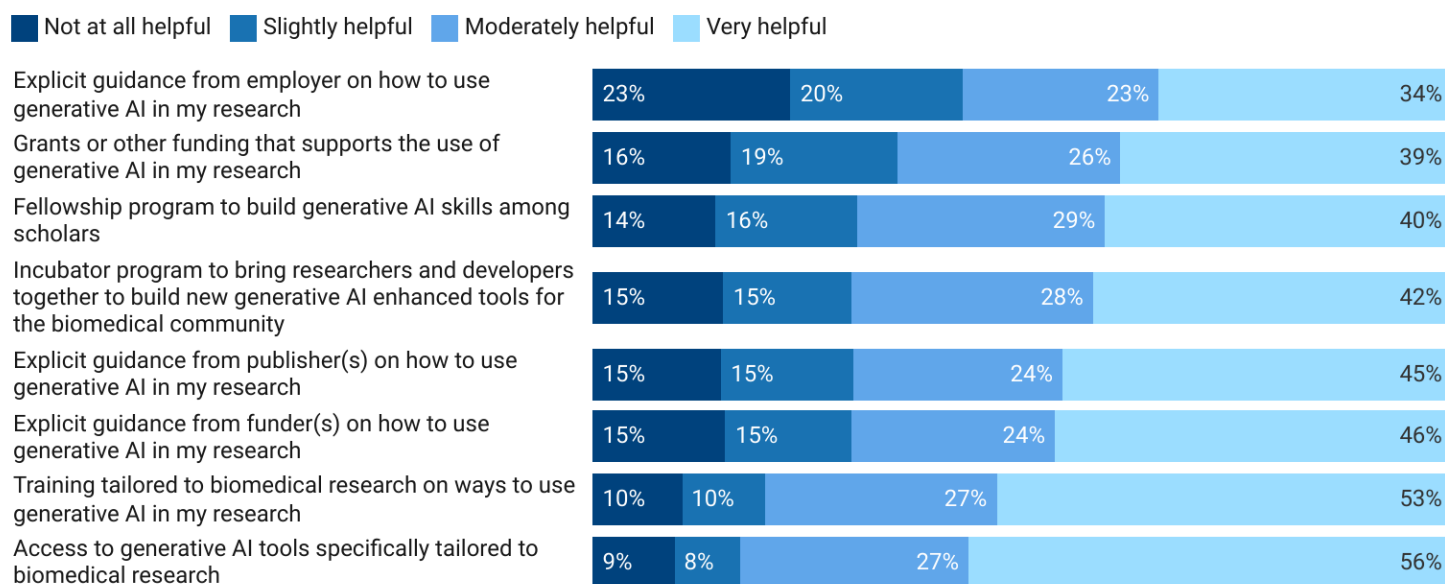
Finding Area Four: Support Needs

Biomedical researchers need support from across the research ecosystem to effectively use generative AI as a research tool.

Categorizing support needs

Biomedical researchers expressed high levels of interest in support, training, and resources, designed to help them better understand and use generative AI in research contexts, another indication that they are curious about and open to greater levels of adoption (Figure 13). In particular, biomedical researchers were interested in access to tools designed specifically for biomedical research and in discipline-specific training about how to integrate generative AI into their research.

Figure 13. “To what extent would the following supports and/or resources be helpful in incorporating generative AI into your biomedical research?”



Created with Datawrapper

The resource that researchers expressed the strongest interest in was biomed-specific generative AI tools: 56 percent of respondents described such tools as very helpful, while an additional 27 percent indicated that such tools would be moderately helpful. One factor driving this demand is the hope that domain-specific tools will provide more accurate, trustworthy, and up-to-date information than is possible using general purpose tools such as ChatGPT.¹⁵

¹⁵ Soumen Pal, Manojit Bhattacharya, Sang-Soo Lee, Chiranjib Chakraborty, “A Domain-Specific Next-Generation Large Language Model (LLM) or ChatGPT Is Required for Biomedical Engineering and Research,” *Annals of Biomedical Engineering* 52, no. 3 (March 2024), <https://doi.org/10.1007/s10439-023-03306-x>; Madhusudan Ghosh, Shrimon Mukherjee, Payel Santra, Girish Na, and Partha Basuchowdhuri, “BLINKtextsubscriptLSTM: BioLinkBERT and LSTM Based Approach for Extraction of PICO Frame from Clinical Trial Text,” in *Proceedings of the 7th Joint International Conference on Data Science & Management of Data (11th ACM IKDD CODS and 29th COMAD)*, 227–31, CODS-COMAD ’24, New York, NY, USA: Association for Computing Machinery, 2024, <https://doi.org/10.1145/3632410.3632442>; Karan Singhal, Shekoofeh Azizi, Tao Tu, S. Sara Mahdavi, Jason Wei, Hyung Won Chung, Nathan Scales, et al., “Large Language Models Encode Clinical Knowledge,” *Nature* 620, no. 7972 (August 2023): 172–80, <https://doi.org/10.1038/s41586-023-06291-2>; Chunyuan Li, Cliff Wong, Sheng Zhang, Naoto Usuyama, Haotian Liu, Jianwei Yang, Tristan Naumann, Hoifung Poon, and Jianfeng Gao, “LLaVA-Med: Training a Large Language-and-Vision Assistant for


The relative value of bespoke and mass market generative AI tools and models for research has been a major topic of conversation among researchers and stakeholders in the research enterprise.¹⁶ Without wading into the core issue of which approach models are best suited to research use, it is worth noting that despite expressing high levels of interest in biomedical-specific generative AI tools, biomedical researchers are making only modest use of BioMedLM, BioLinkBERT, and other biomed-specific generative AI tools or models. Nearly eight out of 10 biomedical researchers who have used generative AI have used general purpose, mass-market tools: just 14 percent reported having used tools designed specifically for use in biomedical research (Figure 14).

Nearly eight out of 10 biomedical researchers who have used generative AI have used general purpose, mass-market tools.

Biomedicine in One Day,” *arXiv*, 1 June 2023, <https://doi.org/10.48550/arXiv.2306.00890>.

¹⁶ Zorik Gekhman, Gal Yona, Roe Aharoni, Matan Eyal, Amir Feder, Roi Reichart, and Jonathan Herzig, “Does Fine-Tuning LLMs on New Knowledge Encourage Hallucinations?” *arXiv*, 13 May 2024, <https://doi.org/10.48550/arXiv.2405.05904>; S. M. Tonmoy, Towhidul Islam, S. M. Mehedi Zaman, Vinija Jain, Anku Rani, Vipula Rawte, Aman Chadha, and Amitava Das, “A Comprehensive Survey of Hallucination Mitigation Techniques in Large Language Models,” *arXiv*, 8 January 2024, <https://doi.org/10.48550/arXiv.2401.01313>; Patrice Béchar and Orlando Marquez Ayala, “Reducing Hallucination in Structured Outputs via Retrieval-Augmented Generation,” *arXiv*, 11 April 2024, <https://doi.org/10.48550/arXiv.2404.08189>; Shubo Tian, Qiao Jin, Lana Yeganova, Po-Ting Lai, Qingqing Zhu, Xiuying Chen, Yifan Yang, et al., “Opportunities and Challenges for ChatGPT and Large Language Models in Biomedicine and Health,” *Briefings in Bioinformatics* 25, no. 1 (January 1, 2024), <https://doi.org/10.1093/bib/bbad493>.

Figure 14. “What kind(s) of generative AI tool(s) have you used in your biomedical research?”



Created with Datawrapper

A second type of support that researchers identified as highly valuable was discipline-specific guidance about best practices for using generative AI in research. Fifty-three percent of biomedical researchers described such help as highly valuable, and 80 percent as at least moderately helpful. It is worth noting that researchers seem to be looking to funders and publishers more than their institutions for explicit guidance on how to use generative AI in their research. Just 57 percent of respondents said that support from their employer would be moderately or very helpful. In contrast, approximately 70 percent of researchers believed that guidance from publishers and funders would be of at least moderate value. Many of their employers would seem to agree: our impression, based on a nonscientific but extensive review of university policies regarding the use of generative AI in academic research, is that institutions are directing researchers to publishers and funders for guidance.¹⁷

¹⁷ For a good overview of publisher policies, see Conner Ganjavi, Michael B. Eppler, Asli Pekcan, Brett Biedermann, Andre Abreu, Gary S. Collins, Inderbir S. Gill, and Giovanni E. Cacciamani, “Publishers’ and Journals’ Instructions to Authors on Use of Generative Artificial Intelligence in Academic and Scientific Publishing: Bibliometric Analysis,” *BMJ* 384 (January 31, 2024) <https://doi.org/10.1136/bmj-2023-077192>.

Discussion

The early impacts of generative AI on scientific practice are most visible in the realm of scholarly communication.

Generative AI has been so ubiquitous that it is sometimes difficult to remember that less than two years ago, the term and the technology were essentially unknown outside of small circles of specialists. Its adoption within higher education has been met with significant resistance from faculty—a survey this spring found that 42 percent of postsecondary instructors prohibit their students from using generative AI, and as documented here, 60 percent of researchers surveyed are not currently using the technology.¹⁸ But what is more remarkable is the openness with which many teachers and researchers have approached generative AI. By early 2024, large majorities of instructors and researchers had experimented to some degree with using generative AI in their work.

The early impacts of generative AI on scientific practice are most visible in the realm of scholarly communication. Major publishers and start-ups have brought dozens of new search and discovery tools designed to streamline the process of locating and synthesizing scholarly literature. As many of these tools are, or will be, integrated as basic features of the discovery platforms run by major publishers, they are likely to soon become so familiar as to be taken for granted. But it is possible that the most far-reaching shifts in scholarly communication practices are the ways that generative AI is being incorporated into academic writing. Recent estimates by Andrew Gray suggest that as many as 1 percent of scholarly articles published in 2023 showed evidence that generative AI had been used during the writing process.¹⁹ With approximately a third of biomedical researchers using generative AI for reviewing and revising written outputs, it is easy to imagine that a similar study of articles published in 2024 will show exponential growth. Though the implications of these changes in scholarly communication practices are not yet clear, they indicate opportunities to accelerate scientific discovery, and raise

¹⁸ Dylan Ruediger, Melissa Blankstein, and Sage Love, "Generative AI and Postsecondary Instructional Practices: Findings from a National Survey of Instructors," *Ithaka S+R*, 20 June 2024, <https://doi.org/10.18665/sr.320892>.

¹⁹ Andrew Gray, "ChatGPT 'Contamination': Estimating the Prevalence of LLMs in the Scholarly Literature," *arXiv*, 25 March 2024, <https://doi.org/10.48550/arXiv.2403.16887>.

difficult questions about whether their effects on research integrity will outweigh those gains.

In many respects, though, generative AI use in biomedical research appears to be at something of a crossroads. Relatively few researchers have become even occasional users. Our survey highlights several key obstacles that may discourage more researchers from transitioning from exploration to adoption including widespread concerns about the accuracy of generative AI outputs, the lack of clear best practices for using generative AI effectively and ethically, and low opinions of the quality of generative AI outputs. Despite the discourse of inevitability that features so prominently in conversations about generative AI, many researchers may remain noncommittal.²⁰ This is particularly likely in relation to experimental applications of generative AI, which, in theory, are the use cases that have the most radical potential to transform scientific inquiry as opposed to scholarly communication.

We are at a unique moment in the reception of generative AI within biomedical communities: researchers are intrigued by the potential value of the technology, and willing to experiment with adopting it, but use cases for the technology have not yet solidified.

We are at a unique moment in the reception of generative AI within biomedical communities: researchers are intrigued by the potential value of the technology, and willing to experiment with adopting it, but use cases for the technology have not yet solidified. Funders, societies, vendors, and other stakeholders have an opportunity to work with researchers to shape the nascent cultural norms and research practices in ways that maximize their scientific potential and mitigate harms.

²⁰ Lauren Leffer, “Too Much Trust in AI Poses Unexpected Threats to the Scientific Process,” *Scientific American*, 1 June 2024, <https://www.scientificamerican.com/article/trust-ai-science-risks/>; “Why Scientists Trust AI Too Much — and What to Do about It,” *Nature* 627, no. 8003 (March 6, 2024): 243–243, <https://doi.org/10.1038/d41586-024-00639-y>.

Four opportunity spaces to address urgent challenges seem particularly important.

Improving the accuracy and quality of generative AI research tools

The primary barrier to greater AI adoption in biomedical research contexts is concern over the accuracy of their outputs. Building trust and transparency into LLMs and creating mechanisms to protect high-quality research data are necessary preconditions to wider adoption and to protect the integrity of the open science environment. Improving accuracy is primarily a technical challenge that is being approached from several directions, including the use of retrieval augmented generation (RAG), knowledge graphs, and specialized large language models.

Responses to our survey indicate that the overall quality of many existing tools are at best marginally optimized for the needs of highly specialized users who have a low tolerance for error. Improved benchmarking of LLMs would enable researchers to make more informed decisions about not only when and how to use them, but about which models are particularly suited for a given task.²¹

Biomedical researchers would benefit from coalition building across the diverse stakeholder communities that collectively benefit from, and bear responsibility for, safeguarding the integrity of the scientific record.

²¹ Benchmarking frameworks are beginning to emerge: See, for example, Liangtai Sun, Yang Han, Zihan Zhao, Da Ma, Zhennan Shen, Baocai Chen, Lu Chen, and Kai Yu, “SciEval: A Multi-Level Large Language Model Evaluation Benchmark for Scientific Research,” *Proceedings of the AAAI Conference on Artificial Intelligence* 38, no. 17 (March 24, 2024): 19053–61, <https://doi.org/10.1609/aaai.v38i17.29872>. For biomedicine specifically, see Qiyuan Chen and Cheng Deng, “Bioinfo-Bench: A Simple Benchmark Framework for LLM Bioinformatics Skills Evaluation,” *bioRxiv*, 21 October 2023, <https://doi.org/10.1101/2023.10.18.563023>; Haoyang Liu and Haohan Wang, “GenoTEX: A Benchmark for Evaluating LLM-Based Exploration of Gene Expression Data in Alignment with Bioinformaticians,” *arXiv*, 21 June 2024, <https://doi.org/10.48550/arXiv.2406.15341>.

Trust also has an important social dimension. Biomedical researchers would benefit from coalition building across the diverse stakeholder communities that collectively benefit from, and bear responsibility for, safeguarding the integrity of the scientific record. This will be especially important if generative AI further destabilizes the larger social, cultural, and political information environments within which the scientific record swims.

Recommendations

- The largest funders have the financial resources to support the development of bespoke foundational models and compute optimized for biomedical research. Such tools have the potential to improve transparency and trust in addition to their potential to improve the quality of generative AI outputs.
- Researchers are beginning to explore the best ways to benchmark the performance of individual generative AI tools against one another and against human created outputs, but there is a clear need to create and validate benchmarking methodologies and standards.
- Scholarly publishers are gatekeepers to the scholarly record: researchers will need confidence that the best and most recent literature is available to generative AI tools used for discovery, synthesis, and research.
- Publishers, preprint servers, and content aggregators will need to develop tools and processes to safeguard the integrity of the scholarly record. They will also need to make use of metadata, perpetual identifiers, and other portable provenance tools to validate the source of their material throughout the transformative processes involved in generative AI and enable rigorous citation.
- Researchers are most often using generative AI for tasks such as searching and synthesizing scholarly material and writing or editing research outputs. These use-cases are areas where domain-specific tools have the potential to have significant impacts on the accuracy and quality of generative AI-assisted research.
- Warning signs about the trustworthiness of the scholarly record have been sounding for some time, and generative AI will only exacerbate them. However, it also creates a new opportunity and level of urgency that should be leveraged in support of frank conversation and sustained collective action among stakeholders focused on strengthening the cultural and technical infrastructure

for trust and integrity across the research enterprise, including consideration of the distribution of responsibility among stakeholders.

Best practices

Researchers do not have a clear sense of how to use generative AI to produce high-quality scholarship that conforms to ethical standards. They would benefit from articulation of best practices, guidelines, and models in this area. To be useful, these best practices will need to reflect disciplinary norms but will also need to be aligned to cross-cutting principles around open science and FAIR sharing. Developing best practices requires consensus building within biomedical research communities that, if successful, will build trust and encourage adoption, while improving the quality of research outputs.

At present, scholarly publishers (and to a lesser degree, funders) have carried much of the weight of articulating research ethics related to generative AI. Their guidelines are useful starting points but would be strengthened by increased participation from research communities, scholarly societies, and other stakeholders about best practices covering the entire research lifecycle. Comprehensive institutional responses from universities to support reskilling by researchers and upgrades to their research infrastructure would have tremendous benefit on supporting best practices at a system level.²²

Recommendations

- Scholarly societies and research communities should host conversations and work to build consensus about how generative AI can contribute to or disrupt disciplinary norms about the ethical conduct of research while there is still a window to shape nascent practices and policies.
- Publishers should be closely observing researchers' practices with generative AI: their position at the center of scholarly

²² Jing Liu and H. V. Jagadish, "Institutional Efforts to Help Academic Researchers Implement Generative AI in Research," *Harvard Data Science Review*, Special Issue 5 (May 31, 2024), <https://doi.org/10.1162/99608f92.2c8e7e81>.

communication gives them a unique opportunity to contribute to the development of a more nuanced vocabulary for describing human/machine collaboration—particularly in regards to reading and writing.

- Universities should be proactive in requiring researchers to conduct research activities in secure AI environments with clear terms of use. This is especially important in biomedical research, which can involve sensitive personal data and is subject to federal regulation. As a corollary, university IT departments, libraries, and other units should commit to providing access to a range of generative AI tools to its campus communities.

Professional development

Biomedical researchers are eager for support services, training, and other resources focused on using generative AI for research. These kinds of professional development opportunities will likely be a primary mechanism through which individual researchers will learn how to adopt best practices for the effective and ethical use of generative AI.

Researchers expressed interest in multiple delivery methods, including help from their institutions, funders, and publishers.

Recommendations

- University units offering training and support to researchers should contextualize generative AI in relation to other forms of AI to help researchers understand which type of AI is best suited to their needs.
- Individual researchers should take the time to identify the specific AI tools best suited instead of using commercial LLMs as a default.
- Universities should identify faculty who are making creative and productive use of generative AI and consider AI Ambassador programs or other means to foster communities of practice within and across departments.
- Evidence suggests that graduate students and postdocs are currently using generative AI more often than faculty. Trainings and workshops designed for or marketed to graduate students and postdocs could have a disproportionate impact on both current and future use of generative AI in research.

- Researchers seek support and guidance from a range of sources. We need an all-hand-on-deck approach to training and best practices involving individual institutions, societies, publishers, funders, and regulators. These efforts will be more useful if additional efforts are made to build consensus across the research enterprise about how to approach generative AI, to create a baseline consistency in norms and integrity standards.

Longitudinal data collection

Our survey captures data about attitudes towards and adoption of generative AI in its first year as a widely available tool. It is a useful snapshot that will become more useful as a longitudinal data point in the future. Because many researchers appear to be on the fence about generative AI, it will be particularly important to understand attitudinal change and adoption over the next several years. Such data will be important for measuring the spread of generative AI across research communities and assessing the effectiveness of initiatives designed to encourage responsible use.

Recommendations

- Biomedical researchers' practices with generative AI are not yet fully formed and could change significantly in the coming years. There will be an ongoing need for surveys of the population, especially those designed to capture longitudinal data.
- The effects of generative AI on scholarly communication and research integrity will need to be closely monitored to assess how, where, and if further intervention is needed.
- Different research communities are likely to use generative AI in ways that reflect pre-existing disciplinary or domain conventions. Surveys similar to this one but focused on different domains can lead to more targeted training and support of researchers and clearer best practices for research conduct.

Acknowledgements

We wish to thank our colleagues Sage Love and Melissa Blankstein for support with the contact list and administration of the survey. Claire Baytas provided insights into the larger landscape of generative AI in higher education throughout the course of the project. Several CZI staff members also played important roles, in particular Kate Hertweck, Adina Abeles, and Denise Robichau. Gary Price has generously shared information about generative AI with us on almost a daily basis. We also wish to thank the individuals who participated in exploratory or cognitive interviews conducted to ideate and refine our survey instrument.