



Generative AI in Academic Research: Perspectives and Cultural Norms

Cornell University Task Force
15 Dec 2023

| | |
|---|-----------|
| Executive Summary | 2 |
| Introduction | 3 |
| Framework for Using Generative AI in Research | 4 |
| Generative AI Use across Research Stages | 7 |
| A. Research Conception and Execution Stage | 8 |
| GenAI for Literature Review | 8 |
| GenAI for Research Infrastructure | 9 |
| GenAI for Data Collection and Generation | 9 |
| GenAI for Ideation & Hypothesis Generation | 10 |
| B. Research Dissemination Stage | 11 |
| C. Research Translation Stage | 13 |
| D. Research Funding & Funding Agreement Compliance Stage | 15 |
| Perspectives and Cultural Norms | 18 |
| Considerations for the Cornell Research Community | 19 |
| Considerations for Cornell Leadership | 21 |
| Appendix 0. Prompts on GenAI in Research (Discussion Starters or Frequently Asked Questions) | 22 |
| A. Research Ideation and Execution Stage | 22 |
| B. Research Dissemination Stage | 23 |
| C. Research Translation Stage | 24 |
| D. Research Funding and Funding Agreement Compliance Stage | 25 |
| Appendix 1. Existing Community Publication Policies | 28 |
| Appendix 2. References Consulted or Cited | 29 |
| Appendix 3. Task Force Charge | 32 |

Executive Summary

This report offers perspectives and practical guidelines to the Cornell community, specifically on the use of Generative Artificial Intelligence (GenAI) in the practice and dissemination of academic research. As emphasized in the charge to a Cornell task force representing input across all campuses, the report aims to establish the initial set of perspectives and cultural norms for Cornell researchers, research team leaders, and research administration staff. It is meant as internal advice rather than a set of binding rules. As GenAI policies and guardrails are rapidly evolving, we stress the importance of staying current with the latest developments, and updating procedures and rules governing the use of GenAI tools in research thoughtfully over time. This report was developed within the same 12-month period that GenAI became available to a much wider number of researchers (and citizens) than AI specialists who help create such tools. While the Cornell community is the intended audience, this report is publicly available as a resource for other research communities to use or adapt. No endorsement of specific tools is implied, but specific examples are referenced to illustrate concepts.

Recognizing many potential benefits and risks of GenAI tools, we address the use of GenAI at four stages of the research process: (i) research conception and execution stage, (ii) research dissemination stage, (iii) research translation stage, and (iv) research funding and funding agreement compliance stage. We further outline coordinating duties by researchers that apply across these stages: duty of discretion, duty of verification, and duty of disclosure; identify categories of GenAI use in research; and illustrate how these duties apply to specific categories and situations in the research process. We emphasize the importance of clearly defining individual and collective/communal responsibilities for meeting these duties throughout the research process. We conclude by offering a set of guidelines for Cornell researchers in varied faculty, staff, and student roles, as well as considerations for Cornell leadership. It is important that Cornell offers its research community access to appropriate GenAI tools and resources, particularly to improve our “AI literacy” regarding the limits of the appropriate use of specific public and commercial GenAI tools and the risks involved in their use for academic research. It is equally important that researchers have Cornell-facilitated access to licensed GenAI tools with privacy/confidentiality provisions, and thus important that Cornell researchers from varied communities understand the value, limitations, and trade-offs of using such tools in research.

The report also contains responses to anticipated questions about best practices and use cases for each of the four stages of research (Appendix 0) that may serve as discussion starters for research communities. Finally, we offer a summary of existing community publication policies regarding the use of GenAI in research from funders, journals, professional societies, and peers, which we surveyed as part of the preparation of this report (Appendix 1); references consulted and cited that include a list of recommended resources (Appendix 2); and task force charge (Appendix 3). Notably, the task force included Cornell faculty and staff quite familiar with GenAI tools and uses, and the task force elected to not use GenAI in drafting the structure, text, or figures of this report.

Introduction

Generative Artificial Intelligence (GenAI) offers transformative capabilities, but we must strike a balance between exploring the potential of these tools and ensuring that research meets standards of veracity, validity, originality, and reproducibility. Briefly, GenAI has the capacity to generate new content (new text or images or audio), typically by computer-generated pattern recognition gleaned from access to large volumes of prior examples. These prior examples, data collectively called training sets or training data, can be provided by the GenAI user or provided by others, with or without their explicit awareness or consent.

This exciting capability to spark new ideas from prior knowledge, perhaps now connected in unexpected new ways, can now be accessed by the masses via online tools and for-fee apps. These users include the masses of academic researchers with a shared sense of research integrity, but with widely varied experience in computer programming and in cultural norms for creation, authorship, and invention. Many of these tools, whether “free to the user” or fee-based, have been released by for-profit companies that maintain the model details as proprietary (i.e., do not disclose details of the trained models or the training data sets that serve as the basis for the GenAI output). Open-source approaches for GenAI development can be a counterpoint that provides a more transparent toolset, but are not an automatic panacea to responsible development or use of such tools in academic research. **Thus, we need to develop common ground and guardrails that prioritize research integrity, accelerate innovation, address obvious issues like data privacy and security, and reflect on non-obvious issues like how practices and expertise of research communities will evolve for better and worse.** No one policy can cover the range of research carried out at a university, from archives to surveys to experimental labs to pure math and the visual arts. GenAI capabilities and affordances are also changing from month to month, only a year into publicly available and initially free versions, but for the cost of providing your own email address or mobile phone number to a for-profit company. External to Cornell, GenAI policies and guardrails are rapidly evolving, and procedures and rules governing the use of GenAI tools in the research enterprise should be regularly updated to stay current with the latest developments. Internal to Cornell, aligning practices for GenAI use with existing policies on research data and with our institutional values will also remain a work in progress for years to come.

GenAI has many potential benefits for researchers at all stages of their research career, for administrative staff who provide key support to but do not participate in research and translation activities directly, and for other users in the course of conducting and administering research. These include:

- **Abstractions.** Many systems for data analysis and document retrieval have been available only to those with substantial programming experience. Likewise, systems for creative audiovisual generation have been limited to specialists or those with years of content production experience. GenAI tools can provide powerful results with interfaces accessible to anyone.
- **Efficiency.** Even when researchers, administrative staff, and other users have the capability to perform a task, GenAI may be able to produce comparable output in dramatically less time, allowing users to focus on more difficult, human interaction-intensive, and/or rewarding tasks.
- **Scale.** GenAI may allow users to perform a task such as coding/annotating documents or generating infographic images on a larger scale than would be possible with manual effort. This ability may allow users to explore larger or broader data sets that were previously limited.

At the same time, we have identified several concerns that apply across the wide range of research disciplines at Cornell and represented internationally. These include:

- **Blind spots and potential bias.** All current generative language models are entirely defined by their training data, and thus perpetuate the omissions and biases of that training data. The model behind ChatGPT has no access to information that was not presented during training, and can only access that information through learned combinations of parameters. There is no explicit and verifiable representation of data or text encoded in the model. Other GenAI systems, such as GPT4, may increasingly have access to web searches and databases to retrieve verifiable sources, and the underlying language model may be able to interact with text returned from those searches (in the same way it interacts with user queries), but the model itself still has no "knowledge" beyond the statistical regularities of its training data.
- **Validation and responsibility.** There is a risk that systems are good enough that users will become trusting and complacent, but bad enough that there are serious problems that have profound consequences. A system that produces seemingly plausible answers, yet is prone to false and biased information, can cause researchers to lower their guard. Therefore, we emphasize the crucial role of researcher validation of research outputs produced with the help of GenAI tools. Responsibility is an area that is particularly sensitive to discipline-specific variation. Most fields have tacit understandings of roles and responsibilities (e.g., principal investigator, Ph.D. student, corresponding author), which may differ substantially from those in even closely related fields.
- **Transparency and documentation.** Guidelines for the use of GenAI in research vary greatly across journals, funding agencies, and professional associations, from blanket bans to restrictions on certain outputs to permissions with disclosure (most with the emphasis that AI-generated outputs should be original, i.e., reflecting the authors' actual work). Laws and regulations for patents, copyright and data use are evolving and vary among countries. As the policies regarding the use of GenAI continue to evolve, maintaining documentation and reporting transparency will remain critical to ensure the reproducibility and replicability of research findings produced with the help of GenAI tools.
- **Data privacy and protection.** GenAI tools should not be assumed *a priori* to be private or secure. Users must understand the potential risks associated with inputting sensitive, private, confidential, or proprietary data into these tools, and that doing so may violate legal or contractual requirements, or expectations for privacy.
- **Resource utilization tradeoffs.** Because GenAI users perceive the output of such tools to be a "free good" or at least generated "in the cloud" even when user fee-based, as with research travel impacts the resource utilization of this computational output can be out of sight and out of mind. However, the magnitude of computational resources operating on Earth to create models based on large volumes of training data, and the electricity use and potential cooling water use associated with such computational processing, can be in tension with values associated with sustainability. A recent study posted on arXiv and currently under peer review attempts to quantify this resource tradeoff in terms of carbon dioxide emissions, equating the carbon cost for you to generate a single image using a specific energy-intensive GenAI model to that required to fully charge your mobile phone (Luccioni et al., 2023 preprint). (CITE: <https://arxiv.org/pdf/2311.16863.pdf>). Certainly, other research-related activities may contribute more significantly on a per-instance basis (e.g., a research group flying to present at an overseas conference). However, Cornell's public commitment to climate action and our individual sense of responsibility for our own resource use choices benefit from our shared awareness that the use phase (or inference phase) of GenAI can be estimated, is non-zero and will remain so without concerted effort, and naturally scales with access to computational resources.

Framework for Using Generative AI in Research

The epochal developments of the past five years in GenAI have enabled systems to *generate* complex, recognizable outputs such as text, programming code, images, and voices. AI as a field has been around for decades, but the output of systems has often been narrow, binary predictions: whether an email is spam, or whether a transaction is likely to be fraudulent. GenAI offers dramatic

new capabilities, generating output in response to *prompts* (i.e., questions, requests, instructions) from the user. See inset: *What are generative uses of AI?*

GenAI provides the user a sense of power in its apparent intellectual assistance on demand, which unsurprisingly also vests the user with a need to take responsibility. Academic research groups and projects often include multiple users with different stages of contribution, different degrees of experience and leadership, and different responsibilities to research integrity and translation of research results to societal impact. Thus, we begin with the following general framework describing categories of **uses** of GenAI in research and categories of **duties** that researchers may have.

There are many levels of potential **uses** in research, ranging from surface level adjustments to applications that blur the boundary of authorship. At one extreme, we might consider systems that simulate a copyeditor, correcting spelling or grammar, which are already integrated in many word processing systems. At the other extreme might be a system that acts as a ghost writer, converting vague descriptions and specifications into polished final presentations. In between, systems might act more like research assistants, collecting and collating information, writing short computer programs, or filling in paragraph bodies from thesis statements. Other uses might be more like reviewers or editors, enabling researchers to "bounce ideas" or summarize a passage of text to ensure that it reads correctly.

These uses imply corresponding **duties** by researchers. Most high-performance GenAI systems are currently available as third-party (i.e., company product, not university-managed resource) cloud (i.e., using remotely located computers) applications, so there is a researcher duty of *discretion* in what data should be uploaded. GenAI, while usually convincing and fluent (at least in English), is often factually incorrect or lacking in attribution, so *verification* is another key duty to ensure accuracy and validity of research outputs. Researchers may also have a duty to provide transparency and *disclosure* to identify how and where GenAI contributed. Finally, we need clear lines of individual and collective *responsibility* to ensure that the other duties are actually executed.

For the remainder of this report we will identify specific situations in the research process, and describe how they relate to these *categories of use* and what duties we believe apply in academic research – and are consistent with Cornell shared values. **In all of these research stages, we consider GenAI to be a useful research tool that can and should be explored and used to great scholarly advantage.** As with all tools, the user is responsible for understanding how to use such tools wisely. As with all academic research, the responsibilities are shared, but the research leader – called principal investigator in some fields and contexts, and lead author, corresponding author, or lead inventor in others – is considered

WHAT ARE GENERATIVE USES OF AI?

To define what this report considers to be generative uses of AI tools (GenAI tools), a key distinction is whether the tool produces outputs that contain concepts and interpretations that the researcher or author did not supply.

Under this distinction, neither a spell checker nor a tool that fixes grammar and improves sentence syntax is considered GenAI, as long as it does not alter concepts and interpretations that the researcher or author supplied. Conversely, a large language model (LLM) that writes a draft of the related work section is strongly generative, even if it received the relevant papers (supplied from the user or scraped from the web) as input. Similarly, a tool that creates an image or diagram in response to a prompt should also be thought of as generative if the prompt contains substantial ambiguity.

Even in this attempt to describe a distinction between technical improvement and new creation, it becomes clear that the boundary of what constitutes GenAI cannot be cleanly delineated. We expect that greater understanding of GenAI tools and their impact on research will shift interpretations of assistance versus creation. Note also that this definition hinges on the use of the tool, not on the tool itself, implying that the same tool used in one context may count as generative, while it is not generative in other contexts.

responsible for communicating expectations to their research colleagues and students, and ultimately bearing consequences of intentional or incidental errors in tool use.

Generative AI Use across Research Stages

We considered four stages of research, each of which may receive different emphasis among Cornell's impressive breadth of research and scholarship areas. **Figure 1** illustrates these four stages where GenAI can be used to great advantage, with appropriate sense of duty by the researcher(s).




Figure 1: Stages of the academic research life cycle in which uses of GenAI can be envisioned, with considerations of use and duty varying by research stage and by research discipline. While such boundaries across research stages are porous, this framing helps distinguish the tradeoffs and responsibilities of use in context. Image component attributions: Appendix 2.

These four stages that can be considered in the life cycle of research include:

- A. **Research Conception and Execution Stage:** Includes ideation by the individual and research team, prior to any public dissemination of ideas or research results.
- B. **Research Dissemination Stage:** Includes public sharing of research ideas and results, including peer-reviewed journal publications, manuscripts and books, and other creative works.
- C. **Research Translation Stage:** Includes reducing research findings or results to practice, which may be in the form of patented inventions or copyrights, for products or processes or policies.
- D. **Research Funding and Funding Agreement Compliance Stage:** Includes proposals seeking funding of research plans, as well as compliance with expectations of sponsors or the US government policies relevant to Cornell as an institution of higher education and a research university.

A. Research Conception and Execution Stage

In this section we discuss uses of GenAI for the "internal" research process prior to the preparation of public documents. Research conception and execution includes literature review, research infrastructure, research ideation, and hypothesis generation.

GenAI for Literature Review

The volume of published research literature and data have been expanding exponentially, accelerating with technology advances such as movable type and publishing changes such as electronic journal proliferation and public databases. It is widely assumed that there are pockets of information in distinct fields that, if combined, could lead to breakthroughs. But despite the volume of published information, those serendipitous connections are infrequent because fields are mutually inaccessible due to technical language, and no one from either field knows to look for the other (e.g., epidemiologists and aerosol physicists). In fact, interdisciplinary research often espouses the mixing of existing information in new ways, implying that GenAI systems which can keep track of vastly more information than any person may find connections that might be missed entirely by humans. Acting as "state-of-the-art information retrieval" systems (Extance, 2018), they go beyond conventional databases, such as Google Scholar and PubMed, by being able to retrieve, synthesize, visualize, and summarize massive amounts of existing knowledge (e.g., Semantic Scholar, Scopus AI, Microsoft Academic, Iris.ai, Scite, Consensus). As such, they have the potential to help overcome the problem of "undiscovered public knowledge" (Swanson, 1986; Davies, 1989), which may exist within the published literature, and break through disciplinary silos, facilitating the discovery of relevant research across diverse academic disciplines.

In this context, suggested practices for using GenAI in the literature review phase of research conception are that:

- GenAI can be used to triage, organize, summarize, and quickly get directionally oriented, in the context of an exponentially growing base of reported claims and established knowledge.
- GenAI can be used to assist with drafting literature reviews, although researchers should fact-check and be aware of incomplete, biased, or even false GenAI outputs. In some cases, it can help to provide GenAI with explicit prompt text to try to guard against the use of fake references (e.g., Dowling & Lucey, 2023), although it still does not guarantee accurate results.
- Subject to authorship, citation, and fact-checking considerations, it can be helpful to use GenAI to ideate and iterate on the quality of a literature review. Examples include (a) refining the review to include both prior research and its connection to the new research

WHAT ABOUT REVIEW ARTICLES?

In practice, these suggestions on use of GenAI in literature review during research conception and ideation could also be applied to those authoring a review article (i.e., containing no new research data or findings, but summarizing past work and identifying new ways forward for the field). However, review articles are currently thought in most communities to contribute an original slant or perspective of the authors; a "ghost written" review article that is conceived by GenAI and then simply fact-checked or copy edited by the authors seems a distortion of that intent, and this possibility should be considered carefully by the authors to confirm that it still expresses their own insights.

idea, (b) rewriting the style of the literature review, and (c) refining the literature review to emphasize the contribution of the new research, such as is relevant to other gaps in literature, uncertainties, or even market sentiments.

Duty of verification. The reliability and quality of AI-powered literature review tools are limited by the databases they search, which can affect the comprehensiveness and accuracy of the results. Therefore, it is advisable to use these tools in conjunction with other methods. Another major concern when using these tools is plagiarism, as they can produce verbatim copies of existing work without proper attribution or introduce ideas and results from actual published work but provide incorrect or missing citations. To minimize the risk of unintentional plagiarism, it is best to start with original text and then use GenAI assistance to refine it, in line with the distinction between AI-assisted and AI-written text (van Dis, Bollen, Zuidema et al., 2023). This will help ensure that AI-generated text is original and reflects the authors' own work, as also emphasized in journals and professional societies that permit the use of GenAI tools (but note that some journals prohibit the use of GenAI for any part of the research process entirely; see **Appendix 1** for a summary of existing community publication policies, noting that such policies are subject to change by those communities and publishers). Finally, depending on the extent of GenAI assistance with information search and literature review production (specifically, when it is used beyond grammatical polishing of author-written text), researchers may have a *duty of disclosure* for this research stage.

GenAI for Research Infrastructure

One of the more benign possible uses of AI is in improving workflows and research processes. Collecting and processing data often involve custom software, using complicated APIs (application programming interface, or software with a specific function) that may be poorly documented. Code generation tools such as Copilot have become powerful and successful, leading to significant improvements in users' ability to create software to collate and analyze data. Other ways might involve using GenAI to help construct or critique survey questions or interview templates. In each case, the AI is not involved in producing or recording data, but in building the infrastructure that is itself used to produce data.

A second category of infrastructure might include code or language generation for presentation of research results. APIs for generating figures, such as matplotlib or ggplot2, are notoriously complicated, with innumerable options for modifying the appearance and layout of graphics. Code generation may help in producing programs to generate graphics from data sets, without being directly involved in the construction of data sets themselves. Similarly, language models might assist in generating alt-text for image accessibility.

Duty of verification. As with any other use of GenAI, infrastructure-building uses require careful checking to ensure that outputs are correct. There should be clear responsibility for who will do this verification. We see less need to disclose the use of GenAI in these "back office" contexts relative to other uses, though still with care for potential implications at later research stages.

GenAI for Data Collection and Generation

A subtle but important distinction is when we move from using *GenAI to help develop tools that we use in research* to using *GenAI as a tool for research*, specifically for data collection and generation. In principle, the potential for data collection is enormous. GenAI can be used to help construct data

sets from unstructured data, such as descriptions of patents, job vacancies, SEC (US Security and Exchange Commission) filings, banker speeches, etc. GenAI tools can also be used to synthesize information coming from text, or images. They can be employed to self-assess (predict accuracy) and augment which coding tasks are conducted by human iteration. Advantages for data collection and generation using GenAI as a tool for research include:

- Collecting and organizing data. Consider the Cornell Lab of Ornithology's example of eBird as one data-rich source: Through this global application platform, birdwatchers have submitted a large amount of bird observations that have already informed development of species distribution models (Sullivan et al., 2014).
- Generating data out of unstructured information.
- Summarizing data coming from various sources. Data related to human clinical trials or patient outcomes hold different and important data privacy concerns, but the collection and organization/cleaning of such data is a key step in inference for patient-centered health outcomes (Waitman et al., 2023).
- Scaling up data collection with GenAI by conducting faster and less resource-intensive experiments.

The challenges of using GenAI tools for data collection and generation primarily relate to the *duties of verification and disclosure*:

- Issues with performance and accuracy: Large language models like ChatGPT are currently not fundamentally trained to speak accurately or stay faithful to some ground truth.
- The reliance on large amounts of data may be challenging, and the needed data may not always be available.
- Bias (King and Zenil, 2023): AI is traditionally trained on data that has been processed by humans. Example: In using ML to categorize different types of astronomical images, humans might need to feed the system with a series of images they have already categorized and labeled. This would allow the system to learn the differences between the images. However, those doing the labeling might have different levels of competence, make mistakes and so on. GenAI could be used to detect and to some extent redress such biases.
- Attribution: Data sources may not always be tracked. There is a need for ensuring correct attribution of data sources.

Given these challenges, the use of GenAI tools for data generation and collection must be carefully documented and disclosed to facilitate research assessment, transparency, and reproducibility.

GenAI for Ideation & Hypothesis Generation

While the use of GenAI for idea generation is under early consideration by most academic researchers, it is important to weigh its strengths and weaknesses in the early phases of research. If we think of the idea generation process as a creative process (as opposed to fact-checking or verification), then complementing ideation with GenAI can potentially offset human weaknesses, such as comparatively poorer memory recall versus recognition processes and narrower breadth of knowledge bases. In this sense, GenAI can complement individual researchers during the idea-generation process and democratize access to research assistants. On the other hand, scientific knowledge relies on the ability to reason rationally, do abstract modeling and make logical inferences. However, these abilities are handled poorly by statistical machine learning (ML). Humans do not need a large amount of data or observations to generate a hypothesis, while statistical ML

relies on vast amounts of data. As a consequence, computers are still unable to formulate impactful research questions, design proper experiments, and understand and describe their limitations.

Furthermore, assessing the scientific value of a hypothesis requires in-depth, domain-specific, and subject-matter expertise. An example is the potential of “Language Based Discovery” (LBD) as the possibility to create entirely new, plausible and scientifically non-trivial hypotheses by combining findings or assertions across multiple documents. If one article asserts that “A affects B” and another that “B affects C,” then “A affects C” is a natural hypothesis. The challenge is for LBD to identify which assertions of the type “A affects C” are novel, scientifically plausible, non-trivial and sufficiently interesting that a scientist would find them worthy of study (Smalheiser et al. 2023). Whereas GenAI does well in identifying and retrieving potential data constructs, researcher domain expertise remains critical for determining the quality of output (Dowling and Lucey, 2023).

While considering the possibilities of GenAI-human collaboration for research ideation, it is essential to emphasize *the duty of discretion* to prevent the leakage of proprietary, sensitive, and confidential information into public information space. Furthermore, since using GenAI tools is an evolving space, academics should learn more about GenAI technologies and stay abreast of potentially useful ways for hypothesis generation. For example, as food for thought, one of the GenAI prompts used by Dowling and Lucey (2023) for idea generation: “You [the GenAI tool] created this research idea, and I’d like you to improve it. Could you see if there is an additional article that can be added, to improve the research idea. Can you also talk about the novel contribution of the idea. Please keep it to about 100 words.” Moreover, as we humans gain experience and familiarity with new tools, we do well to be observant to the expectation that they can also change how we conduct research and interact with researchers at this ideation and hypothesis generation stage – in ways that are not always easy to identify *a priori*.

Finally, we note that research execution includes expectations of responsible conduct of research, which for some studies and disciplines includes prior approval of data use and management, animal welfare and procedures, and human subjects. Use of GenAI in research will likely augment considerations of these approvals per expectations of sponsors or federal agencies through research integrity review processes of the university. Those considerations related to research compliance are expanded in Section D. Next, we consider the stage where research of any type is disseminated through public disclosure including peer-reviewed publications.

B. Research Dissemination Stage

GenAI offers new affordances that support both positive and negative outcomes for research dissemination (Nordling, 2003). On the positive side there is the potential to level the playing field for non-native speakers of English; to provide writing assistance resulting in improved clarity; and for new tools that aid in more equitable discovery of related work (improving on common practices of searching for well-known authors, for example). On the negative side there are serious and reasonable concerns such as erroneous information being disseminated because of inadequate verification; easier plagiarism (either intentional or accidental); lack of appropriate attribution because current LLM-based tools are unable to indicate the source of information; bias and ideological influence; and inappropriate use of GenAI as a lazy peer review tool. Additionally, we must be aware that careless use of GenAI may entrench biases in scholarly communication and dissemination, such as reinforcing the positions of prominent scholars and preferring sources in English as the dominant language of the initial trained models available to the general public. However, future GenAI tools may also provide new interventions to oppose existing biases that are entrenched in current practice. As such, the following is less focused on specific GenAI tools available

today, but more on general recommendations for the responsible use of GenAI tools in research dissemination that upholds research integrity as a principal value at Cornell.

In this section and research stage, we do not discuss questions of copyright, confidentiality or intellectual property (see Section C), but instead focus on the conceptual impact that GenAI can have on producing research output. Following from the definition of GenAI from above, a key distinction is whether the tool **produces output for dissemination that contains concepts and interpretations that the author did not supply**. From this perspective, GenAI tools that fill in concepts or interpretations for the authors can fundamentally change the research methodology, they can provide authors with “shortcuts” that lead to a breakdown of rigor, and they can introduce bias. This makes it imperative that users of GenAI tools are well aware of the limitations of these tools, and that it is clear who is responsible for the integrity of the research and the output that is produced.

Below we discuss these issues in more detail and provide a minimal set of norms that we recommend across all disciplines. However, we recognize that the methodology and standards are differentially impacted by GenAI across disciplines (e.g., humanities as well as engineering), and that community norms around the use of GenAI may be stricter than what is outlined below.

Authorship: We posit that GenAI tools do not deserve author credit and cannot take on author responsibility. This means that authors of research outputs, not any GenAI tools used in the process, carry the responsibility for checking the correctness of any statements. Authors must be aware that GenAI tools can and do produce erroneous results including “hallucinated” citations. The content will be viewed as statements made by the authors. Indeed, there are emerging concerns on impact to scientific publishing with which publishers and AI ethicists are now grappling (Conroy 2023), but the responsibility of authentic authorship is a component of research integrity that will continue to rest with the human authors.

Impact on Concepts and Interpretations: Researchers need to be aware that GenAI tools can have a substantial impact on the research output depending on how they are used to fill in concepts and add interpretations. If the impact is substantial, we recommend that the use of GenAI is disclosed and detailed so that readers are aware of its potential impact. What constitutes substantial impact depends on the type of publication (e.g., journal articles, books, talks, reports, reviews, research proposals) and community norms in the respective discipline. An example that is probably considered to have a substantial impact in any discipline is the use of GenAI to draft a related work section.

Impact on Methodology: Writing and other dissemination activities typically cannot be separated from conducting the research, and the act of writing is explicitly part of the research methodology in some disciplines. A key concern is that the use of GenAI as a “shortcut” can lead to a degradation of methodological rigor. If the use of GenAI tools can be viewed as part of the research methodology, then we recommend disclosure so that readers can assess the rigor of the methodology. Indeed, there may be collective impact on methodology at the scale of the research community’s practices. Whether GenAI becomes a tool that sharpens our minds or a blunt instrument that dulls them is a question that Cornell (and other communities of research scholars) must address actively over time. Historically, human imagination sees most tools as helpful implements to move on to harder problems and more creative discovery and analysis, if one masters the tool instead of the other way around. But we can also recognize from past experiences that zeal for rapid development of exciting new research-enabling capabilities – especially when these provide

competitive advantage over peers that can relate to economic competition or even national security – can shift even the best intentioned individuals to start to behave collectively as a group that focuses sharply on the benefits without openly discussing the costs and trade-offs.

Potential for Bias: Just as authors need to be aware of human biases in their sources, authors using GenAI tools need to be aware that these tools have biases and may reflect specific ideologies. It is the authors' responsibility to counteract or avoid these biases, and authors cannot offload responsibility for bias in their work on the AI system. For example, use of a GenAI tool to create a hospital scene might result in an image in which the nurses are female and the doctors are male. Changing the prompt could address this bias. Another issue is that GenAI tools may reflect a particular ideology, or they may perpetuate historical biases because GenAI tools are trained on historical data. This may be compounded by particular algorithms such as citation analysis which has an inherent time lag, and might further bias recommendations back in time or towards a dominant group or language.

Acceptable Use: There are many different forms and venues of research dissemination: journal articles, books, talks, reports, reviews, research proposals, etc. What is acceptable use of GenAI in one form of communication is not necessarily acceptable in other forms, and authors must adhere to community standards. To take an extreme example, having a GenAI tool draft a peer review from scratch runs counter to the idea of peer review and has an extremely high impact on the review, even if the author checks and edits the review. This is likely unacceptable in most communities. Even within communities, different publication venues (e.g., journal, conference) may have different policies, and authors must check and follow these if more stringent than what is outlined here.

AI Literacy to support Research Integrity: Rigorous and ambitious use of GenAI tools requires a good understanding of the strengths and weaknesses of these tools. Furthermore, as GenAI has become part of the integrity of research and its dissemination, then research leaders such as principal investigators and faculty supervising student research should now make the appropriate use of GenAI part of their mentoring. In particular, part of their mentoring is to communicate the standards and the norms in their specific fields to the researchers and students they lead – just as they mentor on other topics of research conduct (e.g., plagiarism, co-authorship, privacy regulations).

Regulations: Any use of GenAI tools needs to be compliant with regulations (e.g., copyright law, privacy regulations such as HIPAA and FERPA, confidentiality agreements, and intellectual property). In particular, users must be aware that use of GenAI tools may disclose sensitive information to a third party, which may be in violation of regulations and confidentiality norms (Lauer et al. 2023; Conroy 2023). This extension to implications for subsequent research translation to policies, processes, and products of all types is discussed further in Section C.

C. Research Translation Stage

The use of GenAI in any stage of the research process may impact the translation, protection and licensing of intellectual property (IP), commercialization of technology, open-source release of software and other uses of the research output downstream. Interpretations of laws and new regulations regarding GenAI are major topics for governments in many countries including the US. There may be new government agencies and international organizations created for AI regulation and coordination in the near future. In fact, while the European Union recently announced new regulations on artificial intelligence, current understanding is that most of this EU policy effort to create these “first rules” preceded widespread use of GenAI (European Union Commission 2023).

We can draw no immediate conclusions on how the EU's risk-based approach may impact GenAI development and uses specifically. The nature of the impact of GenAI is still evolving and may change in coming years, with legislation and guidelines expected to lag the use of GenAI and its shorter term implications, which may be inadequately addressed under current laws and regulations. The following are important areas for researchers to consider for translation when they use GenAI in their research process:

Inventorship and Patentability: Recent US case law has held that inventors of patents must be human beings under the US Patent Act. Documentation of human contribution and disclosure of the nature of GenAI utilization are essential for patent eligibility. Key information needs to be carefully documented, such as:

- Specific GenAI tools used and rationale for their use;
- Detailed input into and output of the GenAI tool;
- Whether the outputs lead to any aspects of the conception of the invention;
- Contributions of individual inventors in the inventive idea, and how they directed and refined the GenAI output; and
- For research done in teams, delineate the role of GenAI for each inventor.

Copyright and Authorship: Under current US copyright law, copyright can protect only material that is the product of human creativity and authors of copyrighted materials must be human beings. When incorporating GenAI-generated content, the authors should:

- Clearly document the boundary between human-created and GenAI-created content with clear annotations.
- According to guidance published by US Copyright Office, if copyright registration is sought, the nature and extent of the use of GenAI, if containing more than *de minimis* AI-generated material, must be disclosed with clarifications of what part of the work was created by the researchers and what part of the work was created by the GenAI.
- Specifically for computational algorithms and code, where research code can be further translated to wider use through copyright and various licensing types including open-source licensing, considerations attach at a time of active discussion. We note emerging considerations of copyright infringement, not only for creative works such as songs but also for computational code. For example, it is possible that code generated by a LLM reflects code reproduced verbatim from the LLM training set unbeknownst to the user. When such code is part of a research outcome that may be made available to licensees (even open-source licensees), it is possible but not yet well understood how use of such code, even when unintentionally plagiarized from other original sources, may violate copyright or invalidate licenses.

Commercialization and Fair Use: For research that leads to commercialization and publications with financial benefits, to mitigate risks of potential infringement claims, the inventors and authors should:

- Prioritize the use of GenAI tools that are trained using only public domain works, which is a small but growing area of development. For example, the recently announced AI Alliance coalition that includes Cornell as a founding member and anchored by two for-profit companies, IBM and Meta, advocates for development of open-source software tools including those enabling GenAI (<https://thealliance.ai/>; Lin, B., 2023).

- Understand that the commercial intent can significantly impact fair use considerations. Consult with relevant university offices, such as the Center for Technology Licensing or General Counsel's office, when there are questions.
- Stay informed of ongoing litigation that may influence the use of copyrighted materials in GenAI training data set. There are pending class action copyright suits by authors against entities owning GenAI tools for training without compensation to the authors.

Data Privacy and Protection: For data that researchers enter into GenAI themselves, it is important that researchers follow Cornell Policy 4.21 on Research Data Retention. Note that this is an existing policy and practice, simply extended to GenAI. Private, confidential, or proprietary data owned or controlled by Cornell may have certain contractual or legal restrictions or limitations, such as those to sponsors or collaborators, that would preclude their use in GenAI research projects, and it is a researcher's responsibility to verify/determine whether any such data sets have restrictions on such use before inputting them into public-facing GenAI and ensuring compliance with any restrictions mandated by contract, law, or governing body (e.g., IRB, IAUCC). Any use of patient or human derived data should be disclosed to such governing body during the approval process and any such data set should only be used in research projects upon the explicit approval of the relevant governing body on campus.

Training: Specific to this stage of research translation, it is recommended that the university provide ongoing workshops on campus or through online platforms, and offer training materials through websites and other distribution channels, on topics related to the use of GenAI and its impact on patent rights, copyrighted materials, commercialization, open-source release and other uses to aid the researchers in understanding their rights, obligations, best practices and landscapes of relevant laws and regulations. Indeed, Cornell includes faculty and staff experts that can facilitate and co-develop such resources as part of their scholarly practice.

D. Research Funding & Funding Agreement Compliance Stage

During the research funding and funding-agreement compliance stage, there are many potential applications of GenAI. For example, these tools can be leveraged to assist in the writing of technical, science-related information for a proposal to a sponsor or a donor, such as the technical scope and anticipated impact. On the non-technical side, they can also be used to draw appropriate data from multiple data sources to develop information for a biosketch, a report of Current and Pending Support, and other documentation relevant at this stage of the research process.

Work conducted during the Research Funding and Funding Agreement Compliance stage is poised to benefit from the use of GenAI tools, for example, due to efficiency improvements and reductions in the time taken to produce previously time-consuming work. However, the use of these tools also comes with risks. GenAI may produce outputs that include incorrect or incomplete information. These tools also may lack sufficient security and privacy protections, both in the tools themselves, and in any plug-ins and add-ons to the tools.

Note that we and many federal agency sponsors refer to the person of primary responsibility in research as the *PI*, or principal investigator, and pronounced pee-I, for shorthand. We acknowledge that this term is common for research in the sciences and engineering with cultures of team-based research and that other fields have a tradition of independent scholarship and authorship even when enrolled as graduate students.

Responsibility: As with the earlier stages of research, users of GenAI hold some burden of responsibility (or *duty*) in the Research Funding and Funding Agreement Compliance stage. In this stage, however, it is common to attach the primary responsibility of compliance to the leader of the research effort. For example, the accuracy of any information contained in a proposal for funding is ultimately the responsibility of the PI, and so if the PI uses GenAI in the development of materials for that proposal, they must review the information in those materials and correct any omissions, errors, or otherwise inaccurate information. The PI must also understand that although resources (for example, research administration staff professionals in Cornell departments, colleges/schools, or research & innovation office units) are available to help them during this stage of the research process, these resources cannot certify to the accuracy of much of the information provided to them, and therefore cannot be expected to identify mistakes in that information, such as those generated by GenAI. The PI must also understand that they are responsible for the activities of students and research personnel working on funded projects under their supervision or mentorship, and for ensuring the appropriate use of GenAI tools by those individuals. See **Appendix 0**, Prompts on Gen AI in Research for suggested discussion starters.

During this stage, individuals may desire to input information into GenAI to assist in the production of their research proposals, reports to sponsors, or even the public dissemination and translation stage documents that may have specific restrictions placed by the sponsors. Because some of this information will be highly sensitive, such as unpublished technical information or private funding data, users of GenAI tools must understand their responsibility for protecting the privacy and security of any information they input into these tools, and must seek approval to do so from the owner (such as the PI) of any such information. In fact, even in the peer review of sponsored research proposals (e.g., faculty serving on review panels for NSF or study sections for NIH), the use of GenAI may not be allowable by the sponsor (NSF Notice to Research Community, Dec 2023).

In this stage of the research process, it is also important for those who are responsible for making decisions regarding the use of sponsored funds to consider whether, and under what circumstances, it is appropriate to charge the use of GenAI tools to a research account, and to ensure their awareness with each sponsor's requirements. Although some sponsors are clear on whether and how funds may be applied to the use of GenAI, others are not.

Guidance and Training: The nature of these tools, their potential applications, and the associated benefits and pitfalls will continue to develop and change over time, and thus, so will the appropriate guidance on how to use them. Although information and guidance should be shared with users about the risks of the use of GenAI and about which tools to avoid, it is also important to share information and training on how users can make use of these tools, how to navigate security and privacy concerns with confidence, and to provide access to tools that have been vetted and found to be aligned with the university's expectations for security and privacy.

In this context, we suggest the following considerations as resources developed by and for the research community, including staff professionals experienced in research integrity, information systems, and user experience.

- Broad communications and outreach about GenAI in responsible conduct of research.

These communications should include guidance and resources on the use of GenAI, as well as information about training, what tools to use or avoid, and references to offices and units that are available to provide support. When appropriate, this outreach should be shared by central offices and posted to central web pages – such as the recently developed [Artificial Intelligence website](#) hosted by Cornell Information Technology (CIT) that links to Cornell's GenAI in Education report and other resources – rather than from individual units or departments, to create consistent understanding and information access across campus. Providing this type of outreach from central offices can help ensure that the university as a whole is looking to the same resources; that inquiries and concerns come to the appropriate offices; that approaches, advice, and guidance given are consistent across units; and that gaps in accessibility of information and learning are kept to a minimum.

As with training on the use of, for example, animals, human participants, or biological agents in research, centralized training should be provided on the use of GenAI. This training should not only focus on risks and concerns, but on how to get the most out of these types of tools, and how to use them better. “Hackathons as Training” should make it enjoyable for researchers to gain new skills, while also contributing to the safe and responsible use of GenAI.

- Guidance on navigating mistakes made and security breaches should be communicated university-wide. It is important to acknowledge that with these new tools comes some anxiety about making mistakes in using them appropriately or even safely. To an extent, inadvertent

WHAT ABOUT RESEARCH INVOLVING HUMAN PARTICIPANTS?

Applications of GenAI involving human subjects, including but not limited to healthcare research, would be expected to include Institutional Review Board (IRB) review. Research in which GenAI may be a useful tool may also include involving patient-related data. It will take time to consider how ethics and societal impact can be considered thoughtfully and reviewed/approved by appropriate boards when GenAI tools are used or developed in such proposed studies. This consideration is already an active area of research for some at Cornell (e.g., Center for Precision Nutrition and Health; see references cited) and is connected naturally to the faculty-staff IRBs that serve all of Cornell's campuses as part of research integrity and assurance. Part of that research integrity process connects to existing Cornell Policy 4.21 on Research Data Retention (see section: Perspectives & Cultural Norms). Additional resources will benefit researchers who propose such research with unique data protection and ethical/societal considerations, which will continue to meet federal and review body guidelines.

mistakes present opportunities for education and training. However, it is also important that any mistakes that lead to security, privacy, or other concerns are handled correctly and in a timely manner. Information should be shared university-wide about Cornell's expectations and processes with regard to what to do in the case of a potential security or other risk related to the use of GenAI, and which responsible offices should be notified.

- Additional tools and resources should be developed to provide guidance. It would be beneficial to researchers and administrative staff alike to develop a GenAI-enabled tool (e.g., a form of a chatbot) that would respond to common inquiries about the use of GenAI in research. For example, “Can I use GenAI to edit my scope of work?” This tool could be populated with responses to common questions, so that consistent answers could be communicated broadly – even while appreciating that perspectives and cultural norms and even sponsor requirements and expectations may be changing fluidly in the coming years. Because such a tool would be automated and would provide immediate access to answers to these types of common questions, it would both reduce wait times and delays associated with other means of gathering this information, and reduce administrative workload in responding to these types of common requests. Similarly, resources that facilitate awareness of resource use (e.g., estimated carbon dioxide emissions associated with tool use; see Section A) can be made available at the Artificial Intelligence website and/or developed by Cornellians whose research and translation focus includes sustainability practices (e.g., Cornell Atkinson Center for Sustainability).

Perspectives and Cultural Norms

Having framed the use of GenAI in research across the stages of research above, we here summarize the perspectives that can inform our cultural norms. The widespread availability of GenAI tools offers new opportunities of creativity and efficiency, and as with any new tool depends on humans for responsible and ethical deployment in research and society. Thus, it is important that Cornell anticipates that researchers can and should use such tools appropriately, facilitates researcher access to appropriate GenAI tools and to resources to improve researchers' “AI literacy.” It is also important that we develop shared understanding of the limits of appropriate use of specific publicly available and commercial GenAI tools, as well as the tradeoffs or risks involved in their use.

While these perspectives and cultural norms will vary reasonably among different research communities, and likely vary over time in the coming years, we offer the following summary considerations. These are considerations of both opportunity (ambitious use that may create new knowledge, insights, and impact for the world) and accountability or responsibility (duty grounded in research integrity of individuals, research teams, and institutions including Cornell). We consider these to be peer-to-peer guidelines, not a suggestion of any formalized university policy. However, we remind our fellow Cornellians that two existing policies naturally extend to use of GenAI tools in research:

1. [As noted in the University Privacy Statement](#), Cornell strives to honor the [Privacy Principles](#): Notice, Choice, Accountability for Onward Transfer, Security, Data Integrity and Purpose Limitation, Access, and Recourse. This is noted on Cornell's Artificial Intelligence website, along with preliminary guidelines of accountability that are discussed in this report in the context of researcher duties and research integrity.
2. [Cornell Policy 4.21 on Research Data Retention](#). Private, confidential, or proprietary data owned or controlled by Cornell may have certain contractual or legal restrictions or limitations, such as those to sponsors or collaborators, that would preclude their use in GenAI research projects. It is the researcher's responsibility to verify/determine whether any such data sets have restrictions on such use before inputting them into public-facing GenAI and ensuring governing body compliance (e.g., IRB, IAUCC). Relatedly, any use of patient or human derived data should be disclosed to such governing body during the approval process, and any such

data set should only be used in research projects upon the explicit approval of the relevant governing body on campus.

Considerations for the Cornell Research Community

We as colleagues encourage faculty, research and administrative staff, and students to help develop the norms, technology, and public reflection on GenAI use in research, to both shape and stay current on these uses and scholarly practices. These five areas of consideration for the Cornell research community are summarized below, as part of responsible experimentation.

HELP DEVELOP THE NORMS, TECHNOLOGY, and PUBLIC LITERACY around GenAI.

- Actively develop the norms and best practices around the use of GenAI in their disciplines.
- Develop GenAI technology that is particularly suited for research (e.g., improved attribution). GenAI development for academic use should not be left solely to for-profit companies.
- Engage in GenAI public literacy efforts to foster responsible and ethical use of GenAI tools. Using at least one of these tools is enormously helpful to being part of that conversation and process, and many are freely and publicly available with associated caveats on risks of use. **Table 1** provides examples of currently available GenAI tools that can be accessed (denoted as “free” to indicate no financial charge to the user). We emphasize user awareness and appropriate caution: only publicly available data should be included, and the user should assume that any entry of information by the user can be absorbed into that tool’s training set and potentially exposed to others.

STAY UP-TO-DATE with GenAI Uses and Practices

- Each research subcommunity (whether a faculty member’s research group, a department, interdisciplinary research center or institute, or college/school as those researchers see fit) gather relevant information on relevant policies by professional associations, journals and funding institutions to stay up-to-date with evolving policies, practices, and requirements in your field. **Appendix 0** may serve as a discussion starter.
- Train in how to use GenAI tools in a safe, effective, and productive manner in research and innovation. Develop expertise in the potential limitations and risks of GenAI tools.

| Free* GenAI Tool | How to find it | What it can do |
|------------------------|--|---|
| ChatGPT | Go to https://chat.openai.com/ in any browser | Generates text. The free version of ChatGPT uses the OpenAI GPT-3.5 model. |
| Copilot with Bing Chat | Go to https://bing.com/chat . In the Microsoft Edge browser, you can open up Bing Chat in a sidebar that gives you additional functionality, by clicking on the icon in the top right corner. For help installing Edge you may need to reach out to your IT support. | Generates text and images. With the sidebar version in Edge, allows you to ask questions about specific web pages or PDFs in the browser window. The free version of Bing Chat that is available now uses the latest OpenAI models, GPT-4 for text and DALL-E 3 for images. |
| Google Bard | Go to https://bard.google.com/ in any browser. | Generates text. It now uses a more powerful model called Gemini. |
| Stable Diffusion | Go to https://stablediffusionweb.com in any browser. | Generates images using the open Stable Diffusion XL model. |
| Runway ML | Go to https://runwayml.com in any browser. | Generates video from text or still images using the Runway Gen-2 model. |

Table 1: A sample of free GenAI tools as of December 2023, where *free indicates no financial fee charged to the user. None of these tools are endorsed by Cornell, and those using these tools should note that these tools are **only for use with public data**.

Further, when acting as well-informed academic researchers with access to this research tool among others, consider the individual and shared **duties of verification, disclosure, and discretion** across the stages of research ideation and execution, public disclosure, translation, and funding expectations:

Duty OF VERIFICATION

- DO verify the accuracy and validity of GenAI outputs. The responsibility for research accuracy remains with researchers.
- DO check for unintentional plagiarism. GenAI can produce verbatim copies of existing work, or more subtly, introduce ideas and results from other sources but provide incorrect or missing citations.

Duty OF DISCLOSURE

- DO keep documentation and provide disclosure of GenAI use in all aspects of a research process, in accordance with the principles of research reproducibility, research transparency, authorship and inventorship.

Duty of DISCRETION

- X DO NOT assume that GenAI is private. GenAI systems run on training examples, and user input and behavior are a prime source. Even if organizations that provide GenAI tools do not currently claim to use data in this way, there is no guarantee that they may not in the future.
- X DO NOT share confidential, sensitive, proprietary, and export-controlled information with publicly available GenAI tools.
- X DO NOT assume that GenAI output is already considered part of the public domain (e.g., not legally encumbered by copyright). GenAI tools can "memorize" their training data and

repeat it with a level of verbatim accuracy that violates copyright. Even material that is not copyrighted may produce liability for corporate partners in sponsored research, if it is derived from data generated by a competitor.

Considerations for Cornell Leadership

We also provide considerations for Cornell leadership, particularly for aspects of GenAI preparedness and facilitated use in research and innovation that can be implemented collectively across Cornell's colleges, schools, and campuses.

- Develop a knowledge base module, perhaps as part of responsible research conduct training resources, for rigorous, ethical and responsible use of GenAI in research and related activities. Users of GenAI tools need to understand their strengths and weaknesses, as well as regulations around privacy, special data considerations such as personally identifiable, human subject, or proprietary commercial data, and confidentiality and commercialization.
- Consider procurement of Cornell-licensed GenAI tools with data and privacy protection as facilities for research, as well as for administrative and teaching uses. Text generation and chat, program code generation, streamlined processes, and image generation would likely all be of value.
- Consider development or co-development of GenAI tools that are particularly suited for academic research use cases, including use cases in research administration services.
- Identify relevant central offices responsible for providing university-wide communications, guidance, outreach, and training to all GenAI users on various aspects of uses. To the extent that it is possible and relevant, information on the use of GenAI should be shared from central locations to encourage consistent access and understanding across the university, and to avoid siloed, inconsistent, or incorrect information.
- In support of Cornell's public engagement mission, recognize Cornell efforts that improve GenAI public literacy beyond the university-affiliated community.
- Consider periodic updates to Cornell guidance, through a task force or other appropriate mechanisms, given the rapidly changing landscape of generative AI tools, uses, and considerations in academic research and translation of research outcomes.

Appendix 0. Prompts on GenAI in Research (Discussion Starters or Frequently Asked Questions)

We further consider best practices and use cases in response to questions (prompts) for each of the research stages. During the task force, we simply used these prompts to stimulate early conversations and perspectives among Cornell colleagues from different fields of research and at different stages of research. The responses to such questions provided below are not prescriptive or complete, but share our initial, collective responses to such prompts as a diverse group of faculty and staff began shared discussion of this topic.

These same questions could be used within Cornell research group discussions or at department faculty meetings. For researchers to generate familiarity or insights into how these tools work to generate text or images or audio, these same prompts could be entered into multiple GenAI-enabled programs, or entered multiple times in the same GenAI-enabled program, or variations on these prompts.

A. Research Ideation and Execution Stage

When using a tool such as ChatGPT to generate research ideas for a research project sponsored by NSF, how do the researcher and principal investigator decide on which information and ideas to enter and “share” with ChatGPT?

Any information entered into public versions of ChatGPT involves sending that data to a third-party that is under limited confidentiality and privacy restrictions (if any) with end users and not party to agreements with the NSF via data entry by other users, and the information that is entered can eventually become public. As such, where the use of ChatGPT for research idea generation does not currently violate any known general NSF policies, users should also be sure not to violate any other agreements that may exist relative to their funding sponsorship agreements, such as confidentiality, intellectual property, and entity identification clauses (e.g., mentioning NSF in any input data may be discouraged to the extent that it is in conflict with an agreement).

When using a tool such as ChatGPT to brainstorm solutions for research sponsored by a company (e.g., Samsung, Boeing, Johnson and Johnson, Google), how do the researcher and principal investigator decide what information about the project can be entered and shared?

Again, any information entered into public versions of ChatGPT involves sending that data to a third-party that is under limited confidentiality and privacy restrictions (if any) with end users and not party to agreements with corporate sponsors via data entry by other users. Furthermore, any information that is entered can eventually become public. As such, where the use of ChatGPT for brainstorming solutions does not currently violate any known general policies, users should also be sure not to violate any other agreements that may exist relative to their funding sponsorship agreements, especially including but not limited to, confidentiality and intellectual property clauses.

When using generative AI tools to summarize the literature for the introduction or discussion of a peer-reviewed article, how should the researcher and corresponding author attribute or disclose this section of a manuscript or thesis?

In general, authors and/or principal investigators have ultimate responsibility for works (including their accuracy), and furthermore, summaries should not violate plagiarism rules and regulations. Citation style guides and support websites (e.g., for APA, Chicago, Harvard) are being updated to reflect proper citations for verbatim output from generative AI and other uses. As a general practice, authors should be transparent and fully disclose uses of generative AI technologies, consistent with publication outlet, department, or area policies.

What are the conditions, if any, that a researcher should not use GenAI to generate research ideas? Examples may vary among research fields, sources of information included in a prompt including FERPA or HIPAA data, and collaborating or sponsoring organizations.

There are no general rules that prohibit the user of generative AI to generate research ideas. However, because inputs into public generative AI platforms are not confidential and data can also become public, sensitive information and individual data should never be entered for any phase of a research project, pursuant to personal data identification, re-identification, and/or chain of custody FERPA and HIPAA requirements.

In the process of research publication development, can tools such as Bard or ChatGPT be used: To summarize responses to online surveys of income level of state residents? To summarize preclinical research animal model histology? To summarize patients' blood oxygen levels in a registered NIH study? How can the differences among these use cases be distinguished in the responsible conduct of research?

The summary of online survey data and/or other data sets can be assisted by Bard or ChatGPT to the extent that the researcher 1) does not enter confidential or data protected by other laws (e.g., HIPAA), 2) does not violate a broader agreement (e.g., between researcher, institution, host, and/or funder), and 3) has responsibility for the accuracy of the summary.

B. Research Dissemination Stage

When a figure for a publication or presentation or patent disclosure is generated by AI tools (e.g., Midjourney; DeepAI), how should the principal investigator (who is typically the corresponding author or communicating inventor) verify the accuracy and intellectual ownership over the data or content of that image?

We see a distinction between cases where the author supplies all semantics and uses a tool for layout and rendering, and when a tool is used to introduce semantics such as organization of ideas or the generation of structure. Where the author supplies all semantics the use is akin to PowerPoint style suggestions or an automated layout tool, and acknowledgement is generally unnecessary unless publication policies require it. However, when GenAI introduces new semantics then we advise acknowledging its use, and checking the output carefully for accuracy. Norms around intellectual ownership are in flux at this time, and we caution authors that this poses a substantial risk.

Can images used in publications and theses be created wholly by generative AI? How does this expectation change if the GenAI-drafted images are edited by the authors? How does that vary among research disciplines?

Uses include generation of a cover image for a book or presentation, or images similar to clipart. We advise that authors should generally acknowledge the use of GenAI in this case, and they should carefully check images for bias and accuracy. Authors should be aware that, in the US and many other jurisdictions, it is not possible to claim copyright in non-human created works even if any human additions/edits may be copyrightable. There are unresolved legal questions regarding possible copyright infringement both as a result of the training of GenAI programs on works under copyright, and as a result of output that might closely resemble works under copyright (see, for example, the US Congressional Research Service, “Generative Artificial Intelligence and Copyright Law”, Updated September 29, 2023, <https://crsreports.congress.gov/product/pdf/LSB/LSB10922>). There is also significant variability in the acceptability of GenAI-generated images based on the publication venue.

When and how should the corresponding author inform a journal of manuscript elements created by GenAI, if not explicitly required to disclose by the publisher and when not obviously using or studying GenAI? Examples may be a proposed cover image, a graphical depiction of a new method, a graph containing research-generated data, a set of test data, generation of derivative data, etc.

GenAI technology is increasingly being built into services that provide grammar checks, polish language, and proofreading. General purpose tools such as ChapGPT are also effective for these tasks. It is not usual to acknowledge the use of checking and suggestion tools. GenAI tools, like human proofreaders who may not understand the subject matter in detail, can suggest changes that change the intended meaning so authors must still verify suggestions with care. Commercial checking and suggestion tools are being extended with GenAI features to draft entire sections or articles, or summarize texts, so authors should consider when their use crosses the line into generative use as defined above.

When and how should research group leaders (e.g., faculty) communicate these expectations of appropriate/ethical/responsible use of GenAI in research to researchers who are undergraduate students? Graduate students? Postdoctoral researchers? Other research staff?

Educating about the responsible use of GenAI should become part of the regular training on research methodology and norms of the respective discipline. This includes training that research leaders provide, but it also is a responsibility of Cornell to educate faculty and students on the affordances and pitfalls of GenAI tools.

C. Research Translation Stage

If an invention is reduced to practice in part by use of generative AI, how should the inventors document and inform others when considering a disclosure of invention or copyright?

Any use of GenAI in the conception and reduction to practice of an invention or generation of copyrighted materials should be carefully documented and disclosed to the Center For Technology Licensing by the inventors/authors.

For example, as to the conception and reduction to practice:

- What was the GenAI tool used?
- What were the inputs to the GenAI? Do you have rights to the data used for input?

- What were the outputs of the GenAI?
- How did the outputs of the GenAI, if at all, lead to the conception of any aspect of the invention?
- Were there any difficulties encountered in using the GenAI to yield the outputs desired and, if so, which researchers updated use of the GenAI, model and/or data to yield the desired outputs?
- Which researchers substantively contributed to/controlled the development of the input and output corresponding to the invention?

If the research outcome is open-source licensable and/or posted on an open-source repository (e.g., code or algorithm or app), should and how should the researcher disclose use of GenAI in creation of the “open source” item?

Disclosure of the use of a specific GenAI tool and possibly even the origins of, and the rights to use, the input data used will likely be viewed as the standard for ethical behavior over time. Currently there are no hard and fast rules.

If the research outcome is a creative work (e.g., book, play, sculpture, musical score, multimedia exhibit) that used GenAI in the creation of that work, how should the researcher disclose that contribution in discussions of copyright?

According to Guidance published by USCO, if copyright registration is sought, the nature and extent of the use of GenAI, if containing more than *de minimis* AI-generated material, must be disclosed, and what part of the work was created by the researchers and what part of the work was created by the GenAI.

How should the researcher inform themselves of the uncompensated contributions of others to the GenAI output used in their own creative and/or copyrighted work or invention? How does this responsibility depend on whether the researcher derives personal financial benefit (e.g., royalties on published book) from the research outcome?

There are pending class action copyright suits by authors against entities owning GenAI tools. In those suits, GenAI tools are alleged to utilize existing copyrighted works for training without compensation to the authors. Commercial purpose is an important factor in the determination of fair use. For research that leads to commercialization and publications with financial benefits, it will be safer to use GenAI tools that are trained using only public domain works. For data that the researchers put into GenAI themselves, it’s important that they make sure they have the rights to do so regardless of whether they expect financial benefits from the output of GenAI.

D. Research Funding and Funding Agreement Compliance Stage

If grant proposal information related to science (technical scope) and non-science (biosketch, current & pending funding) components are generated by generative AI, who is responsible for editing them before submission to a potential sponsor? Who is responsible if there are omissions or errors in those work products?

The PI is responsible for the accuracy of information related to the science, as well as for omissions and errors in that information. On the non-science side, the PI is, again, primarily responsible for the information contained in their proposal. There are resources available to help them (such as research administration staff), but those resources cannot certify the accuracy of the information

provided to them, or identify mistakes in information provided to them that occurred as a result of the use of GenAI.

It is also important for the PI and their less experienced collaborators (mentees, supervised students) to discuss concerns about inputting information into GenAI tools. This information can be highly sensitive (unpublished technical information, for example) or personal to an individual (Current & Pending funding that must be disclosed to employers and sponsors but not to peers or the general public). To an extent, whether information is considered sensitive may depend on the context of the use of GenAI or of the research field itself. Consensus of this task force was that the PI is responsible for the security of his or her research information, but that anyone who intended to input information into GenAI would need to seek approval to do so from the owner of that information (such as the PI).

Should the costs of Generative AI be charged to a research account, assuming this is not disallowed by the corresponding funding agreement (i.e., not disallowed by a research sponsor)?

The appropriate source of funds for GenAI in research may depend on how the GenAI is being used. If such use is categorized in such a way that other things falling under the same type of use could be charged to a research account (e.g., software services), then it is plausible that the use of GenAI may be acceptable. In some cases, sponsors note definitively whether such charges to sponsored project accounts are allowed, but this is not always the case.

If a principal investigator becomes aware that her graduate student queried a generative AI tool (e.g., ChatGPT) with proprietary data obtained appropriately from a company when summarizing research team meeting notes, what should her next steps be? Who is responsible for notifying the company? Who is responsible for remedying the action if the company has grounds to sue for breach of the data use agreement?

The PI is responsible for what their students do in the course of their Cornell work, and is therefore responsible for ensuring that these individuals use GenAI resources appropriately. That said, mistakes are bound to happen, and they present great opportunities for education and training of both the faculty and of the students.

Further, in these situations when proprietary information is input into GenAI inappropriately, it is reasonable that the PI may feel compelled to directly report this issue to their technical contact at the company, but doing so may not align with Cornell's processes for resolution. Therefore, we should educate faculty about the appropriate way to resolve something like this, which Cornell resources are available to them, and what offices – such as OSP or Counsel's Office – are available to help.

What tools or approaches might Cornell researchers find useful for shared awareness of responsible GenAI use?

The use of GenAI comes with significant privacy and security concerns, and it may be important for the university to gain an understanding of the privacy policies of GenAI companies in order to determine whether they are safe to use. Also of concern are plug-ins to GenAI programs, which may come with their own privacy and security issues.

Although Cornell should provide guidance on risks and tools to avoid, it would also be very useful to provide researchers with information about what tools and resources they can or should use, as well as access to those tools, and confirmation that they've been vetted and found to be secure.

The university could also provide information about the use of GenAI through other means:

- Creation of a tool – “Asking for a Friend” – which could be used to answer questions researchers may have (ex.: “Can I use GenAI to edit my scope of work?”).
- Training should not only focus on risks and concerns, we should also provide training on how to get the most out of these types of tools, and how to use them better. “Hackathons as Training” – could make it fun for researchers to gain new skills, while also contributing to the safe and responsible use of GenAI.
- The IT@Cornell web page is a centralized location that can be used to post preliminary guidelines, general information about GenAI, and what researchers need to know about it.

In order to educate researchers on the use of GenAI, communication and outreach are key. We should educate researchers about the central offices that issue training, guidance, etc. that can help them, rather than leaving them to rely on potentially siloed offices in the units that may not provide consistent advice. If the university as a whole is looking to the same resources, and inquiries consistently come to the same/appropriate offices, approaches/advice/guidance given is more likely to be consistent university-wide.

Finally, much like training on the use of other things in research (animals, human participants, biological agents, etc.), education and training should be provided on how to use GenAI safely.

Appendix 1. Existing Community Publication Policies

We surveyed current policies regarding the use of GenAI in research from funders, journals, professional societies, and peers. We found that most of these examples were stated by journals, professional societies, and research funders, and centered around the research dissemination phase. These include the authorship and review of publications. As of fall 2023, we found relatively little policy about the "private" phases of research, such as ideation or data analysis in what Fig. 1 of our report describes as research stage A, the ideation and execution phase. In these policies, institutions tend to be cautious rather than eager to embrace the possibilities of AI.

Current policies on AI use are often cited in the context of publication, through journals, funding agencies, or professional societies that run peer-reviewed conferences. Many express an openness to the use of GenAI as a tool for writing and editing, especially in so far as it "levels the playing field" for researchers who are not native English speakers. But many also express serious concerns about generation of text beyond grammatical polishing of author-written text. Potential harms usually fall into two categories. First, AI can produce plausible-sounding information that is not, in fact, correct. The published record could become increasingly polluted by unfounded information that is extremely difficult to detect. Second, AI can produce verbatim copies of existing work, possibly causing unintentional plagiarism. More subtly, AI could introduce ideas and findings from actual published work but omit or provide inaccurate citations.

There is significant concern about responsibility. We can find no example of a journal that allows non-human authors, and several that explicitly ban such a practice, as it cannot meet the authorship criteria of accountability for the work. But there are also more subtle questions of duty. Given that the use of generative AI provides substantial risk of inappropriate output (either false or inadequately cited) and may require substantial work to fact-check, who should carry out that work, and who should be punished if it is not done adequately? There is unlikely to be a single policy throughout academia as there are many distinct cultures of collaboration and responsibility. Some fields make strong distinctions between a principal investigator (PI)/advisor's work and PhD student work. In this case a PI may have comparatively little responsibility to check an advisee's use of GenAI. In other fields PIs and students work collaboratively on multi-author publications, where a senior or last author may be expected to have a more supervisory (and therefore responsible) role.

Some agencies also raise the issue of sensitive data. The best current language models are beyond the capabilities of typical laptop hardware, so they are often available as a cloud-based service. While there have been differing statements about what OpenAI or Google might do with information uploaded to the systems, the bottom line is that using such tools exposes potentially sensitive information to third parties. Therefore, institutions explicitly ban entering confidential, sensitive, proprietary, and export-controlled information into publicly available GenAI tools. Similarly, grant agencies and publications prohibit using AI tools in the peer review process to avoid the confidentiality breach.

When AI tools are permitted, there is a consensus about documenting their use in research conception and execution for reporting transparency and reproducibility/replicability purposes. Most publications require disclosure of GenAI use in the Materials and Methods section of a submitted manuscript as well as in a disclosure to editors, except when AI is used as an editorial assistant for author-written text.

Living guidelines for generative AI published in Nature, <https://www.nature.com/articles/d41586-023-03266-1>:

“For Researchers, reviewers and editors of scientific journals

1. Because the veracity of generative AI-generated output cannot be guaranteed, and sources cannot be reliably traced and credited, we always need human actors to take on the final responsibility for scientific output. This means that we need human verification for at least the following steps in the research process:• Interpretation of data analysis;• Writing of manuscripts;• Evaluating manuscripts (journal editors);• Peer review;• Identifying research gaps;• Formulating research aims;• Developing hypotheses.

2. Researchers should always acknowledge and specify for which tasks they have used generative AI in (scientific) research publications or presentations.

3. Researchers should acknowledge which generative AI tools (including which versions) they used in their work.

4. To adhere to open-science principles, researchers should preregister the use of generative AI in scientific research (such as which prompts they will use) and make the input and output of generative AI tools available with the publication.

5. Researchers who have extensively used a generative AI tool in their work are recommended to replicate their findings with a different generative AI tool (if applicable).

6. Scientific journals should acknowledge their use of generative AI for peer review or selection purposes.

7. Scientific journals should ask reviewers to what extent they used generative AI for their review.”

Appendix 2. References Consulted or Cited

Conroy, G. (2023). How ChatGPT and other AI tools could disrupt scientific publishing, *Nature* 622, 234-236.

Current Cornell guidance on Gen AI for use in research, education, and administration (2023): <https://it.cornell.edu/ai>

Current Cornell guidance on Gen AI in teaching (Nov 2023): Available for download at https://teaching.cornell.edu/sites/default/files/2023-08/Cornell-GenerativeAIForEducation-Report_2.pdf and accessible online at <https://teaching.cornell.edu/generative-artificial-intelligence/cu-committee-report-generative-artificial-intelligence-education>.

Bockting, C. L., van Dis, E. A. M., van Rooij, R., Zuidema, W., & Bollen, J. (2023). Living guidelines for generative AI—why scientists must oversee its use. *Nature*, 622(7984), 693-696.

Davies, R. (1989). The creation of new knowledge by information retrieval and classification. *Journal of documentation*, 45(4), 273-301.

Dwivedi, Y. K., Kshetri, N., Hughes, L., Slade, E. L., Jeyaraj, A., Kar, A. K., ... & Wright, R. (2023). “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, 102642.

Dowling, M., & Lucey, B. (2023). ChatGPT for (finance) research: The Bananarama conjecture. *Finance Research Letters*, 53, 103662.

European Union Commission (2023). Press release, Dec 9, 2023: Commission welcomes political agreement on Artificial Intelligence Act.

https://ec.europa.eu/commission/presscorner/detail/%20en/ip_23_6473

Figure 1: Subfigure component sources were obtained from three sources: Microsoft Powerpoint; Freepix free license (patent certificate icon): <https://www.flaticon.com/free-icons/certificate> title="certificate icons. Certificate icons created by Freepix – Flaticon; and iStock (istockphoto.com) standard license with subscription (top hat icon).

Free GenAI to write your research paper with/for you: <https://jenni.ai>

Free GenAI to assist with literature review: <https://www.semanticscholar.org/> (among many available AI-powered tools for literature review, such as Iris.ai, Microsoft Academic, Scopus AI, Elicit, Scite, and Consensus).

Glossary of GenAI-related terms. Steven Rosenbush, Isabelle Bousquette and Belle Lin, “Learn these AI basics,” the Wall Street Journal <https://www.wsj.com/story/learn-these-ai-basics-39247aaf>

Institutional Review Board considerations: Some research centers consider these implications as part of the scholarly effort of the research practice, such as the Center for Precision Nutrition and Health: <https://www.cpnh.cornell.edu/bond-kids-1>.

King, R., & Zenil, H. (2023). Artificial Intelligence in Science: Artificial intelligence in scientific discovery: Challenges and opportunities. OECD Publishing, Paris, <https://doi.org/10.1787/a8d820bd-en>.

Lauer, M., Constant, S., & Wernimont, A. (2023). Using AI in peer review is a breach of confidentiality. *National Institutes of Health*, 23.

Lin, B. (2023). Meta and IBM Launch AI Alliance. Dec 5, 2023, *Wall Street Journal*.

Luccioni, A.S., Jernite, Y., & Strubell, E. (28 Nov 2023). Power Hungry Processing: Watts Driving the Cost of AI Deployment? <https://arxiv.org/pdf/2311.16863.pdf>

National Science Foundation (NSF) [Notice to research community: Use of generative artificial intelligence technology in the NSF merit review process](https://www.nsf.gov/news/notice-to-the-research-community-on-ai) (Dec 14, 2023): https://new.nsf.gov/news/notice-to-the-research-community-on-ai?utm_medium=email&utm_source=govdelivery

Nordling, L. (2023). How ChatGPT is transforming the postdoc experience. *Nature*, 622(7983), 655-657.

Smalheiser, N. R., Hahn-Powell, G., Hristovski, D., & Sebastian, Y. (2023). From knowledge discovery to knowledge creation: How can literature-based discovery accelerate progress in science?. In *Artificial Intelligence in Science: Challenges, Opportunities and the Future of Research*. OECD Publishing.

Swanson, D. R. (1986). Undiscovered public knowledge. *The Library Quarterly*, 56(2), 103-118.

Table 1: Content summarized in tabular form by N. Bazarova as Associate Vice Provost, Research & Innovation; and Z. Jacques as Director, Research Administration Information Services for Cornell's Ithaca and Cornell Tech campuses, with format suggested by B. Maddox, Chief Information Officer at Cornell's Ithaca campus.

Van Dis, E. A., Bollen, J., Zuidema, W., van Rooij, R., & Bockting, C. L. (2023). ChatGPT: five priorities for research. *Nature*, 614(7947), 224-226.

Van Noorden, R., & Perkel, J. M. (2023). AI and science: what 1,600 researchers think. *Nature*, 621(7980), 672-675.

Verma, P. and Ormeus, W. These lawyers used ChatGPT to save time. They got fired and fined. *Washington Post*, published 11/16/2023.

Waitman, L.R., Bailey, L.C., Becich, M.J., Chung-Bridges, K., Dusetzina, S.B., Espino, J.U., Hogan, R., Kaushal, R., McClay, J.C., Merritt, J.G., Rothman, R.L., Shenkman, E.A., Song, X., & Nauman, E. (2023). Avenues for strengthening PCORnet's capacity to advance patient-centered economic outcomes in patient-centered outcomes research (PCOR). *Medical Care* 61(12), S153-S160.

Zewe, A. Explained: What is Generative AI? MIT News, published 11/9/2023.

Appendix 3. Task Force Charge

The following charge was provided to the task force by Cornell's vice president for research & innovation, Krystyn J. Van Vliet, who worked with the task force comprising membership across Cornell's several campuses of research communities in New York state (Ithaca and Geneva, Cornell Tech, and Weill Cornell Medicine) to finalize the report prior to public release. The task force engaged a wider cross-section of the research community's departments and disciplines through discussions during the report development, and a Cornell-internal comment period on a draft version of the report in fall 2023 engaged faculty and staff from additional departments, colleges, interdisciplinary research centers, and offices.

Charge on Generative AI in Academic Research: Perspectives and Cultural Norms

Overview

Cornell's leadership recognizes the opportunity and challenge of generative artificial intelligence (GenAI) on academic research, as well as the communication and translation of research outcomes to research peers and broader society. The Vice President for Research & Innovation charges this task force to discuss and offer guidelines and practices for GenAI in the practice and dissemination of research. The outcome of this ad hoc task force provides clarity in establishing perspectives and cultural norms for Cornell researchers and research team leaders, as internal advice, and is not meant to be a set of binding rules.

Charge to Task Force

Generative artificial intelligence is a tool that is now widely available to the research community. Such capabilities can provide new efficiencies and insights in research, and can also introduce new quandaries for the responsible conduct of research. Faculty and senior research scientists (also called principal investigators of externally funded research) are leaders of research projects, and are thus ultimately responsible for setting and adhering to such norms – particularly when formal guidelines are nascent or disparate. Cornell now has the opportunity to discuss and establish these cultural and professional norms, consistent with our wider institutional values in responsible research across many fields.

This group of staff and faculty is charged to consider any guidelines and best practices on appropriate use and attribution of generative AI that should be shared with the Cornell research community of students, staff and faculty. This task force should identify the range of cultural norms consistent with Cornell values when using this class of tools for research. These recommendations should be communicated in a brief (<10 pages written), internal advisory report by Monday 6 November 2023.

The task force should not include extensions to Cornell education or admissions or hiring practices or institutional communications; those use cases are under consideration elsewhere.

Task Force Roster (listed alphabetically by family name)

- **Natalie Bazarova**
Department of Communication, College of Agriculture and Life Sciences
- **Michèle Belot**
Department of Economics, School of Industrial and Labor Relations

- **Olivier Elemento**
Department of Physiology and Biophysics, Weill Cornell Medicine
- **Thorsten Joachims**
Departments of Computer Science and Information Science, Cornell Bowers CIS
- **Alice Li**
Cornell Center for Technology Licensing, OVPRI
- **David Mimno**
Information Science, Cornell Bowers CIS
- **Bridget MacRae**
Office of Research Integrity Assurance, OVPRI
- **Alexander (Sasha) M. Rush**
Cornell Tech and Department of Computer Science, Cornell Bowers CIS
- **Lisa Placanica**
Cornell Center for Technology Licensing, OVPRI and Weill Cornell Medicine
- **Stephen Shu**
Dyson School, SC Johnson College of Business
- **Simeon Warner**
Cornell University Library
- **Fengqi You**
Smith School of Chemical and Biomolecular Engineering, College of Engineering