

**Report of the Committee on Use of Generative  
Artificial Intelligence in University of Louisville  
Academics**

March 15, 2024

# Contents

<b>Committee Members</b>	<b>4</b>
<b>Introduction, Charge, and Scope of Work</b>	<b>5</b>
<b>1 Ethical Considerations of Generative AI</b>	<b>6</b>
1.1 Guiding principles for policy recommendations for syllabi . . . . .	6
1.1.1 The use of GenAI is not permitted in this course for any assignment . . . . .	6
1.1.2 GenAI may be used with prior instructor permission and appropriate attribution and citation . . . . .	7
1.1.3 The use of GenAI will be allowed without any restrictions . . . . .	7
1.2 Recommended revisions to the Academic Dishonesty Policy in the Student Code . . . . .	7
<b>2 Generative AI in Teaching and Learning</b>	<b>9</b>
2.1 New challenges and opportunities . . . . .	9
2.1.1 Approaches to adapting course design for GenAI technologies . . . . .	9
2.2 Recommendations related to teaching and learning and GenAI . . . . .	10
<b>3 Training and Support for Generative AI</b>	<b>11</b>
3.1 From a student’s perspective . . . . .	11
3.2 Unit establishment options . . . . .	11
3.2.1 A center dedicated to AI research, education, and innovation . . . . .	12
3.2.2 Embed an AI center into existing teaching and learning structures . . . . .	12
3.2.3 Integration into an existing training and support unit . . . . .	13
<b>4 Generative AI in Research</b>	<b>14</b>
4.1 Large language models: a technical primer . . . . .	14
4.2 Chatbots . . . . .	14
4.3 AI detectors . . . . .	14
4.4 Literature review . . . . .	15
4.5 Programming and coding . . . . .	15
4.6 Other use cases . . . . .	15
4.7 Areas for consideration and caution . . . . .	15
4.7.1 Prompt engineering . . . . .	15
4.7.2 Privacy and security . . . . .	16
4.7.3 Fact-checking . . . . .	16
4.7.4 The “black box” problem . . . . .	17
4.7.5 Algorithmic bias . . . . .	17
4.7.6 Intellectual property . . . . .	17
4.7.7 Multifaceted costs . . . . .	17
4.7.8 AI as a tool, not a source . . . . .	17
4.7.9 Other limitations of AI for research . . . . .	18
4.8 Recommendations to the Provost related to GenAI in research . . . . .	18
<b>5 Recommendations from the Steering Committee to the Provost</b>	<b>19</b>
5.1 Recommendations for the Provost to consider now . . . . .	19
5.2 Recommendations for future standing committee(s) . . . . .	19
<b>References</b>	<b>21</b>
<b>A Resources</b>	<b>23</b>
A.1 Teaching and Learning . . . . .	23
A.1.1 Delphi Center web resources related to GenAI . . . . .	23
A.1.2 Web resources from other UofL units . . . . .	23
A.1.3 For further exploration: a selection of teaching and learning-related GenAI links . . . . .	23
A.2 Research . . . . .	23
A.2.1 Chatbots . . . . .	23

A.2.2	Programming . . . . .	23
A.2.3	Non-text media . . . . .	23
A.2.4	Literature review . . . . .	25
A.2.5	Research reporting guidelines . . . . .	25
A.2.6	Journals and preprint servers . . . . .	25
A.2.7	The AI landscape . . . . .	25

<b>B</b>	<b>Charge Letter</b>	<b>26</b>
----------	----------------------	-----------

## Committee Members

### COMMITTEE CO-CHAIRS

Beth Boehm, Jose Fernandez

### SUBCOMMITTEE ON ETHICS

Beth Boehm	<i>Graduate School, Office of the Provost</i>
Jose Fernandez	<i>College of Business, Faculty Senate</i>
Stacie McCutcheon	<i>Privacy Office</i>
Danny McDonald	<i>Delphi Center for Teaching and Learning, College of Arts and Sciences</i>
Susan Tanner	<i>Louis D. Brandeis School of Law</i>

### SUBCOMMITTEE ON TEACHING AND LEARNING

Jennifer Anderson	<i>Delphi Center for Teaching and Learning</i>
Jonathan Kopechek	<i>J.B. Speed School of Engineering</i>
Mary Sherafati	<i>College of Arts and Sciences</i>
Kelvin Thompson	<i>Delphi Center for Teaching and Learning, Office of the Provost</i>

### SUBCOMMITTEE ON TRAINING AND SUPPORT

Madison Colebank	<i>J.B. Speed School of Engineering</i>
Matt Elder	<i>College of Business</i>
Timothy Johnson	<i>College of Arts and Sciences, Writing Center</i>
David King	<i>University Libraries</i>

### SUBCOMMITTEE ON RESEARCH

Michael Brill	<i>Office of Research and Innovation</i>
Sohail Contractor	<i>School of Medicine</i>
Adam Gaweda	<i>School of Medicine</i>
Alex Glynn	<i>University Libraries</i>
Katherine Linzy	<i>Trager Institute</i>
Cynthia Logsdon	<i>Office of Research and Innovation</i>
William Metcalf	<i>Office of Research and Innovation</i>
Olfa Nasraoui	<i>J.B. Speed School of Engineering</i>

### COPY EDITING AND TYPESETTING

Alex Glynn, David King, Katherine Linzy, Olfa Nasraoui

# Introduction, Charge, and Scope of Work

## Generative artificial intelligence

Generative artificial intelligence (GenAI) refers to computational systems able to create new content, such as text or images, based on their training data. GenAI has recently risen to unprecedented prominence, particularly in the form of large language models (LLMs) such as OpenAI's ChatGPT, Microsoft's Copilot, and Google's Gemini: tools that can write and revise text on command, offering new ways for students to learn but also raising questions about academic integrity. GenAI tools can also create images such as artwork, graphs, and figures using prompts, including real and fabricated data. While ChatGPT, Copilot, and Gemini perhaps the best-known examples of GenAI, there are many more LLM tools available, and the technological landscape is rapidly changing.<sup>1</sup>

## The Committee

This committee was formed by the Provost at the request of the Faculty Senate, and Drs. Beth Boehm (Dean of the Graduate School and Professor of English) and Jose Fernandez (Professor of Economics in the College of Business and a faculty senator) were appointed as co-chairs, along with eight other faculty and staff.<sup>2</sup> The charge was “to comprehensively assess the potential applications and impact of ChatGPT and generative artificial intelligence in the domains of academics and undergraduate/graduate student research at the University of Louisville” and “to explore how AI can enrich learning experiences and empower our students and faculty to be at the forefront of technological advancements while maintaining the highest standards of ethical conduct and academic integrity.”

The scope of work asked for four actions:<sup>3</sup>

1. Identify relevant use cases.
2. Assess how generative AI can enhance teaching and learning.
3. Investigate the ethical implications of AI in academics and undergraduate and graduate student research.
4. Identify the training needs of faculty, researchers, and students to enable the effective and responsible use of AI.

The committee formed four subcommittees based on this scope of work:

1. Ethical Considerations of AI
2. Generative AI in Teaching and Learning
3. Training and Support for AI
4. Generative AI in Research

Each committee submitted its own report, and while we have endeavored to refrain from excessive duplication, there is obviously some overlap; the entire committee presents and endorses all four reports and acknowledges that there is much work left to be done in the area of artificial intelligence at the University of Louisville (including faculty research, business operations, admissions and enrollment management, and more) which was beyond the scope of our committee's charge. Legal and compliance topics such as intellectual property rights, privacy concerns, data security, contractual obligations with AI vendors, and adherence to governmental regulations were not within the scope of this committee.

---

<sup>1</sup>Additional tools are described in this report.

<sup>2</sup>Additional committee members were recruited during the development of this report.

<sup>3</sup>The complete charge letter is available in Appendix B.

# 1 Ethical Considerations of Generative AI

The charge to the subcommittee focused on “ethical considerations” was to investigate the ethical implications of AI in academics and undergraduate/graduate student research, with a focus on maintaining academic integrity, avoiding plagiarism, and ensuring that AI-generated content is appropriately attributed and used. Further, we were asked to provide “recommendations for ethical guidelines and policies for the responsible use of AI technologies in academic and undergraduate/graduate research student activities. This includes a review of the university’s current academic dishonesty policy regarding cheating, fabrication, and plagiarism to incorporate work completed by entities that are not human.”

UofL faculty must be aware of the availability of GenAI tools as well as their potential benefits and drawbacks. If we want our students to enter the workforce prepared to use technologies ethically, we must teach them how to use them responsibly. Students must learn to double- and triple-check the outputs of AI to guard against “hallucinations” or made-up references, misleading images, and outright falsehoods, as well as to understand the possible bias in the algorithms, privacy and data security concerns, and issues of “authorship.” While many faculty members may be most concerned now about “cheating” and plagiarism, we should turn our focus to teaching these tools as heuristics, as organizational and feedback tools. We know that they are frequently being used in the workplace, and teaching the ethical use and responsible acknowledgement of the use of these tools is paramount.

Of course, along with the growth of these LLM tools, an industry of AI detection software has developed. Such software is controversial, with the best having between 68 and 84% accuracy rates and all resulting in frequent false positives.<sup>4</sup> As LLMs become better at imitating human language, it is likely that AI detection software will become less effective. While we do not wish to ban the use of AI detection software, we do want to urge caution and suggest that it is better for faculty to teach students the ethical use of AI rather than counting on detection software to “catch” them cheating.

The subcommittee strongly believes that the University has an obligation to teach students and faculty about the ethical use of AI, and thus we recommend that faculty be required to have a policy on AI on their syllabi. While we have provided some examples, we are hesitant to prescribe any one policy for the entire university and instead ask faculty to work within their programs, departments, and colleges to ensure that their syllabi statements are explicit and

transparent about how students may or may not use GenAI in their particular courses.

## 1.1 Guiding principles for policy recommendations for syllabi

1. **All faculty members should include a statement about their policy on the use of GenAI in their syllabi.** There are many sample syllabus statements available for teachers to use, and we believe faculty should have wide latitude in choosing/composing statements that suit their discipline, the level of the course being taught, and student and faculty knowledge and facility with AI. Units, departments, and programs are encouraged to work together to develop policies that best suit their disciplines and the courses they teach.
2. **When multiple sections of a single course are taught each semester, we propose that course directors and instructors collaborate on developing policies to govern all sections** (just as grading scales currently must be the same across all sections of the same course).
3. After review of syllabus statements from various institutions, we found that many are prefaced with a general statement about academic integrity and the use of AI, such as **“Academic misconduct is present in academic work wherever AI assistance has been used when unauthorized, or when authorized, has not been disclosed as required”**.<sup>5</sup>

Whatever decisions course directors and individual faculty make about the use of AI in their courses, the policies should be stated explicitly and transparently in the syllabus. What follows are the three basic types that most institutions have developed, with the subcommittee’s comments.

### 1.1.1 The use of GenAI is not permitted in this course for any assignment

While this approach may be the simplest, it is hard to enforce and is not very practical since each essay produced from the same prompt will be different. Banning the use of AI also will not allow students to learn ethical uses of AI.

<sup>4</sup>Driessen 2024.

<sup>5</sup>AI Working Group n.d.

### 1.1.2 GenAI may be used with prior instructor permission and appropriate attribution and citation

This policy attempts to preserve academic integrity while allowing ethical use of AI. It encourages students to learn about AI and the ethical use of such tools, it allows instructors to specify which assignments might be completed with the use of AI, and it encourages both to learn appropriate citation of AI. At minimum, faculty should teach students how they are to acknowledge the use of AI and should reference the citation style guidebooks appropriate to their particular disciplines. Three common citation guides—APA, MLA, and Chicago—all agree that LLMs are *not* authors and therefore cannot be cited as human authors but nevertheless need to be acknowledged in papers. Instructors and students should consult the resources developed by the APA, the MLA, and the Chicago Manual of Style, or their style manual of choice, regarding how to reference AI-generated text.<sup>6</sup>

Additionally, instructors should encourage students to state how they have used AI to complete the assignment, how they have reviewed the output produced by AI, and whether they have paraphrased the language from the AI output or used it verbatim. Some instructors require students to turn in the AI output(s) with their final assignment since, unlike traditional sources, a citation to AI output will not lead the reader to the exact material used by the student.

### 1.1.3 The use of GenAI will be allowed without any restrictions

This approach depends on instructors integrating AI into the learning objectives and assignments of the course; the instructor will teach how AI is used in

the discipline and how to navigate its potential unethical use. An open policy may make assessment of students' understanding, skills, and critical thinking more difficult, and instructors will still need to provide guidance as to how to acknowledge the use of AI in written assignments.

Just as we recommend that faculty provide clear policies for the use of AI in their classes, we recommend that students seek clarification from faculty before using AI in particular courses. Because expectations for using GenAI will vary across courses and across assignments, students must read the expectations for each course carefully. As a general rule, students should disclose to instructors whether they are using GenAI platforms and in what manner they are using them in coursework. If the course policy on AI is not clearly stated in the assignment instructions and/or in the syllabus, students must communicate with their professor(s) to clarify before using GenAI in their coursework.

The use of GenAI without faculty permission may be considered a violation of the UofL Student Code of Conduct. Suspected violations of this nature will be reported to the Dean of Students.

Students should assume that in the absence of a policy on the syllabus, the use of GenAI tools to complete an assignment or exam is prohibited. Unauthorized use of AI shall be treated similarly to unauthorized assistance and/or plagiarism and be subject to Dean's Discipline.

## 1.2 Recommended revisions to the Academic Dishonesty Policy in the Student Code

In Box 1, we have included in bold our recommended changes to the Student Code.

### Box 1: Revised Academic Dishonesty Policy in the Student Code

The University of Louisville pursues excellence in its work to educate and serve its community with integrity. Academic dishonesty is prohibited at the University of Louisville because it diminishes the quality of scholarship, prohibits independent thought that is essential to intellectual growth and development, makes accurate evaluation of student progress impossible, and defrauds those in society who must ultimately depend upon the knowledge and integrity of the institution and its students and faculty.

Academic dishonesty includes, but is not limited to, the following:

#### A. Cheating:

- (a) Using or attempting to use books, notes, study aids, calculators, **generative AI tools**, or any other documents, devices, or information in any academic exercise without prior authorization by the instructor.
- (b) Copying or attempting to copy from another person's paper, report, laboratory work, computer program, or other work material in any academic exercise.

<sup>6</sup>McAdoo 2023; *How Do I Cite Generative AI in MLA Style?* 2023; *Citation, Documentation of Sources* 2024.

- (c) Procuring or using tests or examinations, or any other information regarding the content of a test or examination, before the scheduled exercise without prior authorization by the instructor.
  - (d) Unauthorized communication during any academic exercise. Except when otherwise explicitly stated by the instructor, examination questions shall become public after they have been given.
  - (e) Discussing the contents of tests or examinations with students who have not yet taken the tests or examinations if the instructor has forbidden such discussion.
  - (f) Sending a substitute to take one's examination, test, or quiz, or to perform one's field or laboratory work; acting as a substitute for another student at any examination, test, or quiz, or at a field or laboratory work assignment.
  - (g) Conducting research or preparing work for another student, or allowing others, **including non-human AI tools**, to conduct one's research or prepare one's work, without prior authorization by the instructor.
- B. Fabrication: Inventing or making up data, research results, information, or procedures, including a record or any portion thereof regarding internship, clinical, or practicum experience.
- C. Falsification: Altering or falsifying information, such as:
- (a) Changing grade reports or other academic records.
  - (b) Altering the record of experimental procedures, data, or results.
  - (c) Altering the record of or reporting false information about internship, clinical, or practicum experiences.
  - (d) Forging someone's signature or identification on an academic record.
  - (e) Altering a returned examination paper in order to claim that the examination was graded erroneously.
  - (f) Falsely citing a source of information.
- D. Multiple Submission: The submission of substantial portions of the same academic work, including oral reports, for credit more than once without prior authorization by the instructors involved.
- E. Plagiarism: Representing the words or ideas of someone else as one's own in any academic exercise, such as:
- (a) Submitting as one's own a paper written by another person or by a commercial "ghost writing" service, **or AI-generated text**.
  - (b) Exactly reproducing someone else's words without identifying the words with quotation marks or by appropriate indentation, or without properly citing the quotation in a footnote or reference.
  - (c) Paraphrasing or summarizing or **using AI to paraphrase or summarize** someone else's work without acknowledging the source with a footnote or reference.
  - (d) Using facts, data, graphs, charts, or other information without acknowledging the source with a footnote or reference; **using AI generated graphs, images, charts or other information without acknowledging that AI was used to generate them**. Borrowed facts or information obtained in one's research or reading must be acknowledged unless they are "common knowledge." Clear examples of "common knowledge" include the names of leaders of prominent nations, basic scientific laws, and the meaning of fundamental concepts and principles in a discipline. The specific audience for which a paper is written may determine what can be viewed as "common knowledge:" for example, the facts commonly known by a group of chemists will differ radically from those known by a more general audience. Students should check with their instructors regarding what can be viewed as "common knowledge" within a specific field or assignment, but often the student will have to make the final judgment. When in doubt, footnotes or references should be used.
- F. Complicity in Academic Dishonesty: Helping or attempting to commit an academically dishonest act. The academic units may have additional guidelines regarding academic dishonesty. It is the student's responsibility to check with their instructors and academic units to obtain those guidelines.



## 2 Generative AI in Teaching and Learning

### 2.1 New challenges and opportunities

The Teaching and Learning Subcommittee took up the Provost's charge to "assess how ChatGPT and AI-powered tools can be integrated into the teaching and learning process to improve student outcomes" and to "explore how generative AI can be harnessed to provide personalized learning paths, promote critical thinking, and address individual learning needs."

The development of new GenAI technologies has major implications for teaching and learning activities in higher education, which create new challenges but also provide new opportunities to increase the quality of learning. A survey by Tyton Partners in September 2023 of over 1,000 higher education faculty and over 1,600 postsecondary students found that 49% of students use GenAI tools, while only 22% of faculty have adopted these tools.<sup>7</sup> These results suggest that many faculty are still in the early stages of the adoption curve. As instructors consider the impact of GenAI on their teaching practices, many questions are raised:

- What is the potential impact of GenAI technologies in my courses?
- Should we encourage the use of GenAI technologies in our courses?
- Do we need to discourage/prohibit the use of GenAI technologies?
- Will use of GenAI technologies hinder our students from learning important knowledge and skills?
- Will use of GenAI technologies aid students in learning important knowledge and skills?
- Can we leverage GenAI technologies to aid our teaching and learning activities?

The answers to these questions may be different for each instructor and may vary further by individual course and program, and the consensus on these issues may continue to evolve in the coming months and years.

Many resources are already being deployed at the University of Louisville to assist course instructors with redesigning their teaching and learning activities in response to the widespread adoption of GenAI technologies. The Delphi Center for Teaching and Learning at the University of Louisville has vetted and shared a collection of useful resources to aid

teaching and learning.<sup>8</sup> In addition, the Delphi Center is committed to providing personalized guidance for course instructors who are adapting their courses for GenAI technologies. Many instructors initially reacted to the adoption of GenAI technologies with fear and trepidation due to the potential adverse impacts on their learning activities, but they can be reassured to know that resources and guidance are available to help them navigate the new landscape in teaching and learning. With these tools in hand to adapt our courses to GenAI technologies, we can confidently continue to provide high-quality teaching and learning activities that help our students thrive and become equipped to benefit our society well into the future.

#### 2.1.1 Approaches to adapting course design for GenAI technologies

**Review learning outcomes in response to GenAI technologies and reconceptualize as appropriate** Just as calculators and spreadsheets eliminated routine computation tasks, GenAI calls for a reassessment of existing learning outcomes and associated assignments. In some cases, GenAI tools may help students reach higher-order learning outcomes than was possible in the past. A few illustrative use cases follow.

**Evaluation** Pursuing this highest level of Bloom's *Taxonomy of Educational Objectives* (Cognitive Domain)<sup>9</sup> is more readily in the grasp of students by directing them to use AI to generate assignment-relevant output and then focusing their efforts on reviewing, assessing, explaining, or correcting this output based upon correctness, legitimacy, soundness, lack of bias, or any other domain-specific criterion relevant to their learning.<sup>10</sup>

**Personalized tutoring and practice** Providing one sample prompt for students to use with a GenAI tool can transform the platform into a personalized tutor that assesses, reviews, challenges, probes, and scaffolds the learning of a given topic.<sup>11</sup> For more discrete tasks, students and their "tutor" can practice over and over again to ensure mastery individually before moving forward in the course of study.

<sup>7</sup>Shaw et al. 2023.

<sup>8</sup>See Appendix A.1.

<sup>9</sup>Bloom 1956.

<sup>10</sup>See Yee et al. 2023, for several concrete evaluation-level assignment examples.

<sup>11</sup>See E. Mollick and L. Mollick 2023, for one such prompt.

**Language translation** GenAI’s ability to manage language translation exercises encourages educators to highlight refined linguistic abilities, such as understanding cultural idioms and employing creative language.

**Literature reviews** Instead of investing significant time in compiling and summarizing literature, students can utilize GenAI for initial searches and summaries, enabling them to delve into the material more profoundly by critiquing methodologies and pinpointing research gaps.

When designing learning activities that aim to develop student expertise around course learning outcomes, it is critical to understand how assignments integrate with GenAI. Instructors should consider running their assignments through GenAI technologies to better assess which parts of the assignment can be automated or answered with GenAI easily.

**Consider utilizing GenAI technologies to assist with teaching and learning activities** GenAI technologies can assist instructors with time-consuming tasks as a starting point for teaching and learning activities, including the following:

**Drafting a syllabus** See Weiss 2023.

**Compiling content and resources about a course topic** See guidance on “Search” from *Ideas for Experimenting with Generative AI* 2023 or “The Generative Textbook” in Alby 2024.

**Generating ideas for in-class learning activities given a specified topic, learning outcome, and allotted time** See “Lesson Planner” in Rice and Kaminski 2024.

**Formalizing the instructor’s feedback to student course work** See “Help Providing General Feedback on Student Writing” in Alby 2024.

**Responding to student questions about the course or upcoming assignments** See “Individualized Student Assistance” in Rice and Kaminski 2024.

## 2.2 Recommendations related to teaching and learning and GenAI

1. The Delphi Center for Teaching and Learning should continue to offer discrete GenAI sessions and incorporate considerations of GenAI into existing programming for UofL instructors.
2. The Delphi Center website should host multiple sample syllabus statements in alignment with the Subcommittee on Ethics’ identification of multiple instructor stances on GenAI in need of communication to enrolled students.
3. The University’s existing curricular review processes (e.g., academic unit committees, central Cardinal Core coordination, etc.) should incorporate a mindful review of learning outcomes in light of GenAI (e.g., outcomes that are easily duplicated by GenAI and/or higher-level outcomes that might be scaffolded by students’ effective use of GenAI).

### 3 Training and Support for Generative AI

The Subcommittee on Training and Support understands that the introduction of AI into the academic space is a uniquely disruptive event, one that will have a profound impact on how education is delivered to students and how research is conducted. With that understanding, the subcommittee further understands that this is also a transformational event that the University of Louisville must address and embrace; otherwise, the institution runs the risk of alienating its core constituents and reducing its effectiveness in delivering on its core principles. Therefore, the subcommittee recommends that the University establish a premier support and training infrastructure that is designed to deliver relevant help and resources to all university members, including students, faculty, staff, and researchers.

#### 3.1 From a student's perspective

One of the realities of the AI discussion is knowing that students, at all levels, are already integrating the use of GenAI into their coursework. These students, regardless of where they are in their educational progress, are already using AI systems, such as ChatGPT, to assist them in gaining knowledge and completing their applied exercises.

Many students believe that systems such as ChatGPT help to improve their understanding in a more time-efficient manner than traditional study methods alone. Furthermore, such systems can bridge the gap when a student is left behind in a course and is too timid to come to the professor. The student can also ask minor questions of the AI system and receive answers and information immediately, which is important when they are trying to understand concepts in the moment rather than waiting for a response several days later. The student can also query the systems in different ways—different from the way they can query their professors—and ask follow-up questions of the AI, helping to narrow down the answers to a more concise solution that meets their specific needs.

It is generally believed that systems such as ChatGPT will be used extensively in society and industry. This is especially true of the science, technology, engineering, and math (STEM) fields and the biological and physical sciences. The total penetration of AI into the world is not completely certain, but people want information quickly and correctly, so if today's student—tomorrow's engineer—knows how to utilize a GenAI system, their job security will likely be greater than that of someone who does not, simply because the person using the GenAI can program

faster and adjust more quickly than someone who does not use such a tool.

The importance of knowing what tools are used by those we are educating is paramount, and there are resources available to all educators today to bring into their curriculum to begin introducing AI concepts.

#### 3.2 Unit establishment options







Here, we consider the question of how UofL might create a center or institute dedicated to AI by looking at existing models from other universities. At the present moment, UofL does not have this kind of center, and part of the charge of the Training and Support Subcommittee was to consider the potential either for a dedicated center or for a center to grow out of existing institutional units. At present, there are a number of developing AI initiatives spread across units and departments at UofL:

- The School of Medicine is currently working on the [Center for AI in Radiological Sciences](#).
- The Speed School:
  - Certificate program in [artificial intelligence in medicine](#).
  - Artificial intelligence track in the [BA in computer science](#).
- The Law School has a professor developing a GenAI toolkit to aid legal instruction.<sup>12</sup>
- The Delphi Center:
  - [Resources for Teaching in the ChatGPT Era](#).
  - [Teaching and Learning with AI](#).

The new unit would model itself on other similar units found in comparable higher education and research institutions in the United States. This unit would provide education on the current state of AI, how to integrate the use of AI in delivering education, and how university members can employ AI in their daily work: research, education, planning, patient care, etc. The subcommittee has researched example centers from other institutions as well as how AI content could be integrated into instruction and student work. Furthermore, the subcommittee has researched various techniques that instructors could implement in their coursework to either bolster or curb the use of AI.

The following suggested unit designs could be implemented by the university, beginning with the most ambitious. Figure 1 illustrates additional example implementations undertaken by other higher education institutions.

<sup>12</sup>Scoggins 2023.

Many universities have focused first on Generative AI research and/or usage in the classroom		
Non-exhaustive		
Initiative	Description	
 USC	Center for Generative AI and Society focusing on AI impact on culture and society	Launched research center that will seed research, convene experts, expand the university's national leadership with a focus on the ethical use and innovation of generative AI
 AUBURN UNIVERSITY	AI@AU: Artificial Intelligence Initiative creating university-wide advantages including AI computational resources, certificate programs, and Think Tank	Investment to coordinate/support AI-related educational and research activities, install of AI computational resource to support AI-related education and research, create a distinguished lecture series, launch AI certificate programs, and create a faculty-led AI think tank
 Stanford University	Stanford Institute for Human-Centered AI awarding seed grants for innovative uses of generative AI in education	Awarded 16 teams \$625K in seed grants to explore innovative uses of generative AI in education to develop AI tools to improve teaching and learning and to better understand critical issues in the learning context
 Mizzou University of Missouri	AI@Mizzou and SEC AI Consortium schools promoting AI exploration through public workshops	Held a workshop open to the general public about generative AI models like ChatGPT with the goal of encouraging MU to create a campus-wide strategy to embrace the new technology
 UCLA MIT  CORNELL TECH	Partnership with Break Through Tech AI creating training and career opportunities for underrepresented students	National program to teach AI to underserved college student populations. UCLA, Cornell Tech, and MIT in partnership with Offers a free 18-month AI program for skills-based training, portfolio-development coaching and career mentoring

Source: dot.LA, USC, Auburn, Stanford, UCLA

McKinsey & Company 14

Figure 1: Example implementations of GenAI in higher education and research.

### 3.2.1 A center dedicated to AI research, education, and innovation

The most ambitious option would be to create a center dedicated to engaging in cutting-edge research that also, in turn, has an educational component dedicated to preparing people to use AI and to be adaptable to future developments. A number of colleges and universities with similar, disparate AI hubs have invested significant resources into AI “centers” that coordinate the training of faculty, staff, and students in the use of AI and the application of existing AI capabilities to community-based problems. Examples of similar units at other institutions include, but are not limited to the following.

#### University of Indiana Luddy Center for AI

- Brings together faculty (>50) who are using AI in their research.
- Has a dedicated building—the \$35 million, 58,000 square foot Luddy Artificial Intelligence Center (opened 2021).
- Funds were provided by a “mega”-donor.

#### Artificial Intelligence Center at the University of South Carolina

- Provides a dedicated graduate program.

#### University of South Florida Institute for Artificial Intelligence

#### University of Stony Brook Institute for AI-Driven Discovery and Innovation

- State-of-the-art facilities.
- The institute was initially funded by a large National Science Foundation (NSF) grant.

As these examples indicate, this approach is resource-intensive and has required either large donor contributions or government grants. Another model for funding a freestanding AI center, then, would be to aim for government/corporate/economic crossover; the Trager Institute has used this model at UofL. Louisville is a regional hub for several industries: shipping (UPS), food services (Texas Roadhouse, YUM! Brands, Brown Forman), automotive manufacturing (Ford), and healthcare (UofL Health, Norton Healthcare, Baptist Health). Other existing institutes, such as the Additive Manufacturing Institute of Science & Technology (AMIST), rely on government grants and funds to operate.

### 3.2.2 Embed an AI center into existing teaching and learning structures

Alternatively, a number of schools have opted to embed training and support concerning AI within existing learning centers. These sites are interested in re-

acting and adjusting to the rapidly changing environment but often leave research and innovation to individual units:

- [Harvard's Derek Bok Center for Teaching and Learning](#)
- [The University of Florida's AI2 Initiative](#)
- [Columbia University's Center for Teaching and Learning](#)
- [Emory University's Center for AI Learning](#)

These units primarily provide programming and resources to faculty that can be used on a voluntary basis. Part of the challenge for these centers is buy-in. Producing and compiling resources and training is one thing; getting faculty to use them is another.

### **3.2.3 Integration into an existing training and support unit**

A final option would be to integrate AI training and support into an existing training and support unit (e.g., the Delphi Center). This option would imply a dramatic shift in resources and focus for the affected unit, as the unit would require training and resources itself to bolster its capabilities and bandwidth. This would necessitate the additional allocation of funds and potential staffing but would likely be less resource-intensive than establishing a completely new unit or center.

Regardless of the form of support the University chooses to implement, there are many opportunities that educators, researchers, and students can take today to integrate the use of GenAI into their daily work.



## 4 Generative AI in Research

The Subcommittee on Research was focused on highlighting potential applications of GenAI in academics and undergraduate/graduate student research, as well as issues of potential concern demanding further consideration. The subcommittee’s composition, with significant involvement in the health sciences, enabled them also to consider issues specific to the use of AI in healthcare and health sciences research. While fully developed guidelines for the use of GenAI in research are beyond the scope of this Steering Committee, researchers should promptly be made aware of the areas of caution raised here, and the discussion below should inform future guideline development.

### 4.1 Large language models: a technical primer

Large pre-trained Language Models (pLLMs)<sup>13</sup> are machine learning models at the core of modern conversational GenAI tools. Language models (LMs) are trained to predict the next word given its context—typically the previous words—or to estimate the probability distribution of the words given their context in information retrieval.<sup>14</sup> An autoregressive LM will add this predicted word to the previous context to construct an updated context, then generate the next word in the sequence, and so on. By repeatedly predicting the next word and adding it to the next context, an LM can thus generate text.

Thanks to recent research on attention models<sup>15</sup>, modern large language models (LLMs) are now trained on massive quantities of text data that are publicly accessible from the World Wide Web, such as Wikipedia articles, Reddit posts, and electronic books. LLMs that are used to drive conversational AI, such as ChatGPT, are further trained using reinforcement learning with human feedback (RLHF) to generate text that is verified by humans for plausibility.<sup>16</sup> Their generic text generation abilities allow LLMs to be co-opted to perform a variety of other tasks, ranging from question answering to summarization and information extraction.

Importantly, when interacting with LLMs, one must remember that these tools do not contain an internal abstract representation of reality (world model). All information generated by LLMs is based on very complex statistical relationships between sequences of words learned from massive amounts of

data, encoded in massive numbers of parameters. In other words, LLMs lack one key element of human intelligence: the ability to form abstract concepts from perceptions.

Which task an LLM will perform depends on the prompt text it is given as input. How well an LLM will perform a particular task depends on how well the prompt is designed in a process known as prompt engineering (see §4.7.1).<sup>17</sup> Because most LLMs are trained for a generic task on generic data, they may not produce accurate results for certain specialized tasks. One way to improve their performance is by fine-tuning them, such as by retraining the model on a specialized data set or for a specialized task. This process, however, can be computationally costly.

### 4.2 Chatbots

Most widely-used GenAI systems take the form of chatbots, able to take instructions and answer questions in a conversational format. This enables the user to communicate with the bot in plain English, ask follow-up questions, and request changes to the bot’s output. Additionally:

- Some systems allow the user to communicate with the bot by speaking into a microphone or uploading files.
- Some are able to generate images or audio in addition to written text.
- Some are able to access the internet and other external data sources, while others are ignorant of information later than their knowledge cutoff.

A summary of the most prominent LLM chatbots currently in use is provided in Appendix A (Table 1).

### 4.3 AI detectors

Various tools have been designed to distinguish between content written by humans and AI-generated content. One analysis of ten such tools found that [Winston AI](#) (84% accuracy) and [Sapling](#) (68%) were the best paid and free AI detectors, respectively.<sup>18</sup> However, there is some controversy as to the real-world effectiveness of AI detectors.<sup>19</sup>

As LLMs improve in their ability to imitate human language, AI detectors may become less effective. We

<sup>13</sup>Devlin et al. 2019.

<sup>14</sup>Song and Croft 1999.

<sup>15</sup>Vaswani et al. 2017.

<sup>16</sup>Bai et al. 2022.

<sup>17</sup>Brown et al. 2020.

<sup>18</sup>Driessen 2024.

<sup>19</sup>Williams 2023.

must be also be wary of false positives. Since presenting an AI’s output as one’s own work may be considered tantamount to plagiarism, it is a serious accusation.

#### 4.4 Literature review

A number of tools have been specifically designed to facilitate literature review.<sup>20</sup> These tools endeavor to perform one or both of the following functions:

1. Find research relevant to supplied plain text (other research papers or plain text research questions) rather than a search string.
2. Synthesize or summarize key findings from search results.

The degree to which these objectives are achieved is as yet unknown.

#### 4.5 Programming and coding

Most LLMs capable of text processing can handle code in common languages (Python, JavaScript, etc.). LLM chatbots can write code snippets, functions, or even whole programs. They can also help debug and rewrite code; users can copy code and/or error messages into the chat interface and explain the problem in plain English.

Code writing assistant bots, such as GitHub Copilot, are available for integration with widely used integrated development environments (IDEs), such as Visual Studio Code.<sup>21</sup> These assistant bots can write code in response to plain text prompts or suggest “completions” based on code written, comparable to an advanced form of autocorrect or IntelliSense.

There are serious limitations to be aware of before using LLMs to help write code. LLMs make mistakes, and since small changes in syntax can radically alter the behavior of a program, it is imperative that AI-generated code be scrutinized by a human programmer before it is run. Limitations include:

- LLMs sometimes confuse languages that have similar syntax or share libraries, such as Python and R.
- Less commonly used languages, such as Perl or Lua, may not be as well-supported as Python or JavaScript. Some code assistants have a list of supported languages in their documentation.
- LLMs that have a knowledge cutoff, such as Claude, cannot assist with code that depends on updates or modules released after that cutoff.

<sup>20</sup>See Appendix A.2.4 for examples.

<sup>21</sup>See Appendix A.2.2 for examples.

<sup>22</sup>See Appendix A.2.3 for examples.

- It is not safe to give LLMs access to code containing confidential data, such as API keys.

#### 4.6 Other use cases

The most prominent GenAI systems are primarily concerned with textual inputs and outputs, but other applications have been successful.<sup>22</sup> GenAI can be used to generate outputs beyond text, ranging from the arts (visual, textual, or musical form, etc) to visualizations that are generated from datasets (e.g., trends) or textual descriptions (such as mathematical functions or reverse captions), and scientific diagrams (e.g. drawings of anatomy, cells, chemical elements, or geometrical concepts, etc.), even to proteins and chemical compounds.

The implications, benefits, and risks vary widely depending on the discipline and the case. For instance, AI-generated chemical compounds should be disclosed as such, along with the model source, in addition to their properties and environmental impact.

Because “fake” multimedia can be easily generated, researchers have the additional onus of interpreting such products in their research, with the implications varying widely according to discipline: from journalism to history and from physics to biology, etc.

With sophisticated prompt engineering, one can obtain “attempted” solutions to quantitative problems (mathematical proofs, logic, physics, business, etc.) often with a rationale that can range from being correct to totally illogical, although the text expressing it often sounds very convincing.

#### 4.7 Areas for consideration and caution

##### 4.7.1 Prompt engineering

Prompt engineering (PE) is crucial for obtaining the desired output for special tasks from GenAI tools, whether they generate text or other modalities, such as images. In addition, judicious prompt engineering can reduce the number of calls to a GenAI tool, which in turn can reduce computational and financial costs. Depending on the complexity of the task, the most common PE methods include zero-shot, few-shot, and chain of thought (CoT). Zero-shot methods simply provide the model with a verbal description of the task and the desired result. Few-shot methods provide several examples of input and output to demonstrate the task, followed by the input for which the answer is desired. CoT prompting describes the

reasoning steps that lead to the desired answer for complex tasks.

#### 4.7.2 Privacy and security

Researchers must not enter sensitive information into AI models. Even models that claim to protect input data may not adequately do so. Entering protected information may constitute a violation of state and federal laws that protect patient, student, or other personal information (e.g. HIPAA, FERPA, KRS 61.931), whether or not the data are stored on the developers’ servers.

We conclude that chatbots cannot comply with the Health Insurance Portability and Accountability Act (HIPAA) in any meaningful way despite industry assurances.

—Marks and Haupt (2023)

Inputting even deidentified data into a chatbot may give the developers enough information to make inferences about a patient’s health or a clinician’s prescribing practices. Such information is extremely valuable and can be sold to advertisers and data brokers. Experts can also use AI to reidentify data.<sup>23</sup> Although users may not intend to input sensitive information, prompting is a two-way street—chatbots prompt the user with follow-up questions. The Federal Trade Commission has expressed concerns about the ability of chatbots to gain undeserved trust from users.<sup>24</sup> Healthcare providers may be tempted to divulge more than they originally intend.

This consideration has implications on the ethical conduct of research and IRB review. For instance, it should influence what type of model is used in research—a closed source model requires entering information into an interface or online service, which divulges the data. In contrast, an open-source model can be downloaded and used in a secure mode on a private platform, hence providing better privacy, albeit at the cost of the increased computational resources necessary to store and execute a large model. Privacy implications may therefore constrain a researcher to use only open-source models that can be entirely downloaded and used within the most secure confines and settings.

“Enterprise” versions of various chatbots are also available for organizations to use internally. The idea is to utilize the power of the chatbot on an organization’s internal resources and data without compromising security. Typically, developers agree not to collect and leverage enterprise user inputs for further

training of their models. In practice, enterprise versions may not be secure, and products meeting commercial data protection standards do not necessarily meet the demands of protected health information, for example.

When Office 365 users log into Microsoft Copilot, the interface claims that their “personal and company data are protected in this chat.” They are not protected sufficiently. Data are still sent to Microsoft’s servers, which may violate the data protection standards to which we are subject.

#### 4.7.3 Fact-checking

Care should be taken when using the outputs of GPT-4, particularly in contexts where reliability is important.

—OpenAI et al. (2024)

GenAI models are known to “hallucinate” non-existent references or confidently declare falsehoods. AI is not a substitute for systematic literature search, which is a necessary component of evidence-based practice. In fact, LLMs have been shown to have high hallucination rates, while their accuracy remains disappointingly low, especially in the academic domain.<sup>25</sup>

LLMs cannot “understand” user input; they can only identify linguistic patterns and imitate them.

By default, language models optimize the *next word prediction* objective, which is *only a proxy* for what we want these models to do.

—Ouyang et al. (2022, emphasis added)

ChatGPT is fundamentally a text transformer—not an information retrieval system.

—Walters and Wilder (2023)

Consequently, LLMs will sometimes output text that appears credible but has no factual basis. In particular, LLMs have a known tendency to cite non-existent sources in convincing bibliographical style. Even when citing real sources, LLMs may paraphrase them inaccurately.

If a prompt is ambiguously phrased, LLMs may (wrongly) guess user intent rather than asking clarifying questions or admitting that they do not “understand” what is being asked. Indeed, they can be “confidently wrong”.<sup>26</sup> This is, perhaps, because LLMs

<sup>23</sup>Marks 2021.

<sup>24</sup>Atleson 2023.

<sup>25</sup>Sun et al. 2023.

<sup>26</sup>Gravel, D’Amours-Gravel, and Osmanliu 2023.

<sup>27</sup>Walters and Wilder 2023.



are optimized to provide answers satisfying to human users, who are biased in favor of confident responses over doubtful or noncommittal ones.<sup>27</sup>

ChatGPT sometimes writes plausible-sounding but incorrect or nonsensical answers. Fixing this issue is challenging as [during reinforcement learning training], there's currently no source of truth.

—*Introducing ChatGPT* (2022)

#### 4.7.4 The “black box” problem

AI increasingly relies on paradigms, such as deep learning, in which the developers themselves do not understand how their models make decisions. When using particularly complex models, such as LLMs, in research, it becomes impossible to document and reproduce research methods. While humans learn from experience, a bot's methods are more comparable to intuition than to a set of instructions that can be followed or reproduced.<sup>28</sup>

For example, if we use an LLM to conduct a literature search, we cannot document the search strategy or reproduce the results. For most LLMs, the same prompt will produce a different output each time, and we cannot determine or appraise the criteria by which articles were included or excluded. This is antithetical to the principles and practice of systematic search, which is a cornerstone of evidence-based practice.

Using a model that is not open source adds an additional layer of opaqueness because the model parameters (such as the connection weights) are unknown to the user.

#### 4.7.5 Algorithmic bias

As machines, AI systems may give the false impression of being impartial and objective, but they are the products of the data inputs used to train them, which were created by human beings, and the choices of the developers, also human beings. Thus, AI is subject to many of the same biases and errors as humans are.

Concerns have been raised that existing datasets may underrepresent particular sociodemographic groups<sup>29</sup> and, if used as training data, may yield inequitable AI models. Indeed, racial and political biases have been observed in the outputs of ChatGPT.<sup>30</sup> It may be possible to counter this issue by use of carefully selected training data.<sup>31</sup>

<sup>28</sup>Bathae 2018.

<sup>29</sup>Arias-Garzón et al. 2023.

<sup>30</sup>Deshpande et al. 2023; Baum and Villasenor 2023; Rozado 2023.

<sup>31</sup>Wang, Chaudhari, and Davatzikos 2023.

<sup>32</sup>*Recommendations for the Conduct, Reporting, Editing, and Publication of Scholarly Work in Medical Journals* 2023.

<sup>33</sup>*Recommendations for the Conduct, Reporting, Editing, and Publication of Scholarly Work in Medical Journals* 2023; “Tools Such as ChatGPT Threaten Transparent Science; Here Are Our Ground Rules for Their Use” 2023; Flanagin et al. 2023; Hosseini, Rasmussen, and Resnik 2023; Thorp 2023.

#### 4.7.6 Intellectual property

Technology often advances more swiftly than policy and legal frameworks can respond. Intellectual property (IP) rights, (specifically, the legal rights to the research output) should be considered prior to engaging in AI development, especially when using research data owned by the University. Researchers at universities may engage in the creation of innovative AI software and algorithms. It is essential to note that copyright protection extends to software source code, and algorithms may be eligible for patent protection.

When using third-party software, researchers should ensure that the software includes explicit terms of use affirming the user's ownership of output. Researchers should only utilize AI software that transparently discloses the origins of its training sets; AI software may incorporate content that is owned by others, and, as a result, the outputs could infringe on their copyright. Every researcher should coordinate with UofL Office of Innovation & Commercialization during the planning stages of their research. This office collaborates with researchers to safeguard proprietary rights in commercially significant training sets, all the while maintaining the flexibility to make them accessible for nonprofit research purposes.

#### 4.7.7 Multifaceted costs

Different methods and decisions as to how to use a GenAI tool, such as prompt engineering and fine-tuning, will generally lead to different computational and budgetary costs, which can drain available resources. Furthermore, when accumulated across a large number of tasks and users, different ways to use these tools may lead to varying impacts on the environment.

#### 4.7.8 AI as a tool, not a source

Chatbots (such as ChatGPT) should not be listed as authors because they cannot be responsible for the accuracy, integrity, and originality of the work.

—ICMJE (2023)<sup>32</sup>

The general consensus among prominent scientific publishing organizations is that AI models cannot be credited as authors because they cannot be held accountable for their statements.<sup>33</sup> Human authors alone must accept the responsibility of author-

ship. Simultaneously, presenting the output of an AI model as one's own work is unethical and compromises the integrity of the research.

In summary, AI should be used to support and compliment human efforts, not as a replacement for them. Any use of AI in research should be reported transparently in publication. It may be also good practice, and in some cases required for scientific integrity reasons such as reproducibility, to include as much information as possible about the context and parameters of the use of GenAI, such as:

- A link to the GenAI tool and version.
- The prompt (or sequence of prompts) used.
- Date(s) of the prompt(s).
- Other parameters depending on the tool used.

#### **4.7.9 Other limitations of AI for research**

Outputs are often unnecessarily verbose and repetitive. Due to attempts at moderation of harmful con-

tent by the developers, LLMs will sometimes be excessively cautious. Conversely, because the developers' moderation practices cannot account for everything, some harmful outputs will still be possible.

### **4.8 Recommendations to the Provost related to GenAI in research**

1. Investigate enterprise AI as a potential means to protect research and other data.
2. Allocate resources for enhancing campus computational infrastructure (hardware systems, storage, and software) to facilitate the use of GenAI.
3. Create a web-based portal of links to existing tools with clear categorization based on open source status, size, computational and monetary costs, and pre-trained models in different domains or for different tasks.
4. Develop formal guidelines for the secure and appropriate use of AI in research.

## 5 Recommendations from the Steering Committee to the Provost

### 5.1 Recommendations for the Provost to consider now

1. The committee feels strongly that the Provost should appoint a standing committee to organize the University's various projects, unit committees, and business operations centered on AI and future technological disruptions; this committee should propose an AI governance structure.
2. We also recommend that the Provost charge each unit with creating a committee to consider the impact of AI on the unit's curriculum, academic policies, research, business policies, and so on. There should be a representative from each unit committee to report to the university-wide standing committee.
3. Require every faculty member to include a statement about their policy on the use of GenAI on each syllabus.
4. Send the changes to the Academic Dishonesty Policy in the Student Code of Conduct suggested by the Subcommittee on Ethics to the Dean of Students for implementation.
5. The Delphi Center should continue to offer discrete training sessions for instructors on the use of GenAI, and they should begin integrating AI use into existing programming for instructors.
6. The Delphi Center website should host sample syllabus statements in alignment with the Subcommittee on Ethics' identification of multiple instructor stances.
7. The curricular review processes at the University should incorporate a rigorous review of learning outcomes in light of GenAI.
8. Researchers should be trained on the ethical use of AI (including acknowledgement and citation) and on the potential data and research security issues surrounding GenAI. We recommend that the research office incorporate training on AI use in its research training modules. The Subcommittee on Research raises considerations and cautions in their report that should be addressed promptly.
9. The bold recommendation from the committee, which comes first from the Training and Support subcommittee, is that the University of Louisville found a center dedicated to research on AI, a center that could also house and coordinate efforts around AI and academics/education. This is an expensive recommendation, and some of the same outcomes could be met by the next recommendation.

10. Provide additional resources and staff to the Delphi Center so that it can embed an AI center into its current structure.

### 5.2 Recommendations for future standing committee(s)

These topics were outside this committee's charge but were considered by the committee to be highly important subject material for the recommended standing committee(s) on AI.

1. Develop an AI governance structure for the University of Louisville.
2. Develop a communication plan for informing the community about AI and about these committees; we recommend that the Provost's website host a centralized page on AI to list resources, collect questions, and discuss policies and best practices.
3. Unit committees to consider domain-specific applications and concerns regarding AI.
4. Explore the implications of AI integration into "big systems" (PeopleSoft, Office 365, Blackboard, etc.) on our work and security of data in those systems. Monitor how other products used by the University are employing AI, including potential data breach risks.
5. Explore the implications of AI for medical and healthcare research and clinical practice, including the security of patient data. Develop guidelines to foster secure and ethical use of GenAI in clinical contexts.
6. Assess whether ITS is sufficiently resourced to support faculty and student use of AI, particularly with regard to the additional pressure AI places on information security and compliance issues.
7. Likewise, the General Counsel's office and the Risk, Audit, & Compliance Office should work with the standing committee on policies and guidelines governing faculty, staff, and student use of AI, since many GenAI products have free versions that do not require going through procurement procedures.
8. Explore enterprise AI, which may facilitate greater data security.
9. Explore AI for business operations: how AI can be used to streamline processes in budgeting, finance, admissions and enrollment, student success, HR, ITS, and so on; how AI is already doing so; how AI can free up faculty and staff time; how it can save the university money.

10. Explore the implications of GenAI on the University workforce and potential changes to their activities and responsibilities. Create training for faculty, students, and staff with a view to greater adaptability.

## References

- AI Working Group (n.d.). *A Report on AI and the UAB Community*. Birmingham (AL): University of Alabama at Birmingham. 16 pp. URL: [https://www.uab.edu/provost/images/images/initiatives/Report\\_-\\_AI\\_and\\_the\\_UAB\\_Community.pdf](https://www.uab.edu/provost/images/images/initiatives/Report_-_AI_and_the_UAB_Community.pdf).
- Alby, Cynthia (2024). *AI Prompts for Teaching: A Spellbook – EmTech*. SUNY Exploring Emerging Technologies for Lifelong Learning and Success. URL: <https://emtech.suny.edu/search/ai-prompts-for-teaching-a-spellbook/> (visited on 03/12/2024).
- Arias-Garzón, Daniel et al. (Mar. 1, 2023). “Biases Associated with Database Structure for COVID-19 Detection in X-ray Images”. In: *Scientific Reports* 13.1, p. 3477. ISSN: 2045-2322. DOI: 10.1038/s41598-023-30174-1.
- Atleson, Michael (May 1, 2023). *The Luring Test: AI and the Engineering of Consumer Trust*. Federal Trade Commission. URL: <https://www.ftc.gov/business-guidance/blog/2023/05/luring-test-ai-engineering-consumer-trust> (visited on 03/11/2024).
- Bai, Yuntao et al. (Apr. 12, 2022). *Training a Helpful and Harmless Assistant with Reinforcement Learning from Human Feedback*. arXiv: 2204.05862 [cs]. URL: <http://arxiv.org/abs/2204.05862> (visited on 03/14/2024). preprint.
- Bathae, Yavar (Spr. 2018). “The Artificial Intelligence Black Box and the Failure of Intent and Causation”. In: *Harvard Journal of Law & Technology* 31.2, pp. 889–938. URL: <https://jolt.law.harvard.edu/assets/articlesPDFs/v31/The-Artificial-Intelligence-Black-Box-and-the-Failure-of-Intent-and-Causation-Yavar-Bathae.pdf>.
- Baum, Jeremy and John Villasenor (May 8, 2023). *The Politics of AI: ChatGPT and Political Bias*. Brookings. URL: <https://www.brookings.edu/articles/the-politics-of-ai-chatgpt-and-political-bias/> (visited on 03/11/2024).
- Bloom, Benjamin S., ed. (1956). *Taxonomy of Educational Objectives: The Classification of Education Goals*. In collab. with Max D. Engelhart et al. Vol. 1. 2 vols. Ann Arbor (MI): David McKay Company. 216 pp.
- Brown, Tom et al. (2020). “Language Models Are Few-Shot Learners”. In: *Advances in Neural Information Processing Systems*. Ed. by H. Larochelle et al. Vol. 33. Curran Associates, Inc., pp. 1877–1901. URL: [https://proceedings.neurips.cc/paper\\_files/paper/2020/file/1457c0d6bfc4967418bfb8ac142f64a-Paper.pdf](https://proceedings.neurips.cc/paper_files/paper/2020/file/1457c0d6bfc4967418bfb8ac142f64a-Paper.pdf).
- Citation, Documentation of Sources* (2024). The Chicago Manual of Style Online. URL: <https://www.chicagomanualofstyle.org/qanda/data/faq/topics/Documentation/faq0422.html> (visited on 03/12/2024).
- Deshpande, Ameet et al. (Apr. 11, 2023). *Toxicity in ChatGPT: Analyzing Persona-assigned Language Models*. arXiv: 2304.05335 [cs]. URL: <http://arxiv.org/abs/2304.05335> (visited on 03/11/2024). Preprint.
- Devlin, Jacob et al. (May 24, 2019). *BERT: Pre-training of Deep Bidirectional Transformers for Language Understanding*. arXiv: 1810.04805 [cs]. URL: <http://arxiv.org/abs/1810.04805> (visited on 03/14/2024). Preprint.
- Driessen, Koen (Feb. 26, 2024). *Best AI Detector — Free & Premium Tools Tested*. Scribbr. URL: <https://www.scribbr.com/ai-tools/best-ai-detector/> (visited on 03/11/2024).
- Flanagin, Annette et al. (Feb. 28, 2023). “Nonhuman ”Authors” and Implications for the Integrity of Scientific Publication and Medical Knowledge”. In: *JAMA* 329.8, pp. 637–639. ISSN: 1538-3598. DOI: 10.1001/jama.2023.1344.
- Gravel, Jocelyn, Madeleine D’Amours-Gravel, and Esli Osmanliu (Sept. 2023). “Learning to Fake It: Limited Responses and Fabricated References Provided by ChatGPT for Medical Questions”. In: *Mayo Clinic Proceedings: Digital Health* 1.3, pp. 226–234. ISSN: 29497612. DOI: 10.1016/j.mcpdig.2023.05.004.
- Hosseini, Mohammad, Lisa M. Rasmussen, and David B. Resnik (Jan. 25, 2023). “Using AI to Write Scholarly Publications”. In: *Accountability in Research*, pp. 1–9. ISSN: 1545-5815. DOI: 10.1080/08989621.2023.2168535.
- How Do I Cite Generative AI in MLA Style?* (Mar. 17, 2023). MLA Style Center. URL: <https://style.mla.org/citing-generative-ai/> (visited on 03/12/2024).
- Ideas for Experimenting with Generative AI* (Nov. 15, 2023). *Ideas for Experimenting with Generative AI: Use Cases and Things to Keep in Mind*. Harvard University Information Technology. URL: <https://huit.harvard.edu/news/ai-use-cases> (visited on 03/11/2024).
- Introducing ChatGPT* (Nov. 30, 2022). OpenAI. URL: <https://openai.com/blog/chatgpt> (visited on 03/11/2024).
- Marks, Mason (Apr. 1, 2021). “Emergent Medical Data: Health Information Inferred by Artificial Intelligence”. In: *UC Irvine Law Review* 11.4, pp. 995–1066. ISSN: 2327-4514. URL: <https://scholarship.law.uci.edu/ucilr/vol11/iss4/7>.
- Marks, Mason and Claudia E. Haupt (July 25, 2023). “AI Chatbots, Health Privacy, and Challenges to HIPAA Compliance”. In: *JAMA* 330.4, pp. 309–310. ISSN: 1538-3598. DOI: 10.1001/jama.2023.9458.

- McAdoo, Timothy (Apr. 7, 2023). *How to Cite ChatGPT*. <https://apastyle.apa.org>. URL: <https://apastyle.apa.org/blog/how-to-cite-chatgpt> (visited on 03/12/2024).
- Mollick, Ethan and Lilach Mollick (Sept. 25, 2023). *AI as Personal Tutor*. Harvard Business Publishing Education. URL: <https://hbsp.harvard.edu/inspiring-minds/ai-as-personal-tutor> (visited on 03/11/2024).
- OpenAI et al. (Mar. 4, 2024). *GPT-4 Technical Report*. arXiv: 2303.08774 [cs]. URL: <http://arxiv.org/abs/2303.08774> (visited on 03/11/2024). Preprint.
- Ouyang, Long et al. (Mar. 4, 2022). *Training Language Models to Follow Instructions with Human Feedback*. arXiv: 2203.02155 [cs]. URL: <http://arxiv.org/abs/2203.02155> (visited on 03/11/2024). Preprint.
- Recommendations for the Conduct, Reporting, Editing, and Publication of Scholarly Work in Medical Journals* (May 2023). International Committee of Medical Journal Editors. URL: <https://icmje.org/icmje-recommendations.pdf> (visited on 10/16/2023).
- Rice, Charles and Peter Kaminski (Mar. 11, 2024). *Microsoft / Prompts-for-Edu*. Redmond (WA): Microsoft. URL: <https://github.com/microsoft/prompts-for-edu> (visited on 03/11/2024).
- Rozado, David (Mar. 2, 2023). “The Political Biases of ChatGPT”. In: *Social Sciences* 12.3, p. 148. ISSN: 2076-0760. DOI: 10.3390/socsci12030148.
- Scoggins, Jill (Nov. 14, 2023). *UofL Law Professor Developing Generative AI Toolkit to Aid Legal Writing Instruction*. UofLNews. URL: <https://www.uoflnews.com/section/science-and-tech/uofl-law-professor-developing-generative-ai-toolkit-to-aid-legal-writing-instruction/> (visited on 03/14/2024).
- Shaw, C. et al. (Aut. 2023). *GenAI in Higher Education: Fall 2023 Update Time for Class Study*. New York: Tyton Partners. 17 pp. URL: <https://tytonpartners.com/time-for-class-2023/GenAI-Update>.
- Song, Fei and W. Bruce Croft (Nov. 1999). “A General Language Model for Information Retrieval”. In: *Proceedings of the Eighth International Conference on Information and Knowledge Management*. CIKM99: Conference on Information and Knowledge Management. Kansas City Missouri USA: ACM, pp. 316–321. ISBN: 978-1-58113-146-8. DOI: 10.1145/319950.320022. URL: <https://dl.acm.org/doi/10.1145/319950.320022> (visited on 03/14/2024).
- Sun, Kai et al. (Aug. 20, 2023). *Head-to-Tail: How Knowledgeable Are Large Language Models (LLM)? A.K.A. Will LLMs Replace Knowledge Graphs?* arXiv: 2308.10168 [cs]. URL: <http://arxiv.org/abs/2308.10168> (visited on 03/14/2024). Preprint.
- Thorp, H. Holden (Jan. 27, 2023). “ChatGPT Is Fun, but Not an Author”. In: *Science (New York, N.Y.)* 379.6630, p. 313. ISSN: 1095-9203. DOI: 10.1126/science.adg7879.
- “Tools Such as ChatGPT Threaten Transparent Science; Here Are Our Ground Rules for Their Use” (Jan. 2023). In: *Nature* 613.7945, p. 612. ISSN: 1476-4687. DOI: 10.1038/d41586-023-00191-1.
- Vaswani, Ashish et al. (Dec. 4, 2017). “Attention Is All You Need”. In: *Proceedings of the 31st International Conference on Neural Information Processing Systems*. NIPS’17. Red Hook, NY, USA: Curran Associates Inc., pp. 6000–6010. ISBN: 978-1-5108-6096-4.
- Walters, William H. and Esther Isabelle Wilder (Sept. 7, 2023). “Fabrication and Errors in the Bibliographic Citations Generated by ChatGPT”. In: *Scientific Reports* 13.1, p. 14045. ISSN: 2045-2322. DOI: 10.1038/s41598-023-41032-5.
- Wang, Rongguang, Pratik Chaudhari, and Christos Davatzikos (Feb. 7, 2023). “Bias in Machine Learning Models Can Be Significantly Mitigated by Careful Training: Evidence from Neuroimaging Studies”. In: *Proceedings of the National Academy of Sciences of the United States of America* 120.6, e2211613120. ISSN: 1091-6490. DOI: 10.1073/pnas.2211613120.
- Weiss, Mitchell (July 27, 2023). *If Your Syllabus Needs a Refresh, Generative AI Can Help — Harvard Business Publishing Education*. Harvard Business Publishing Education. URL: <https://hbsp.harvard.edu/inspiring-minds/if-your-syllabus-needs-a-refresh-generative-ai-can-help> (visited on 03/11/2024).
- Williams, Rhiannon (July 7, 2023). *AI-text Detection Tools Are Really Easy to Fool*. MIT Technology Review. URL: <https://www.technologyreview.com/2023/07/07/1075982/ai-text-detection-tools-are-really-easy-to-fool/> (visited on 03/11/2024).
- Yee, Kevin et al. (2023). *ChatGPT Assignments to Use in Your Classroom Today*. UCF Created OER Works 8. Orlando: University of Central Florida. 125 pp. URL: <https://stars.library.ucf.edu/oer/8>.



## A Resources

### A.1 Teaching and Learning

#### A.1.1 Delphi Center web resources related to GenAI

**Resources for teaching in the ChatGPT era** Public compilation of slower-to-change resources. <https://louisville.edu/delphi/resources/chatgpt>

**Generative AI/ChatGPT workshops** Public listing of upcoming workshops related to teaching with generative AI. <https://louisville.edu/delphi/programs/featured/generative-ai-chatgpt-workshops>

**Teaching and learning with AI** UofL-restricted SharePoint site featuring ever-evolving, crowdsourced resources. <https://cardmaillouisville.sharepoint.com/sites/TeachingandLearningwithAI>

#### A.1.2 Web resources from other UofL units

**Generative AI LibGuide** Subject guide from the Kornhauser Health Sciences Library. <https://library.louisville.edu/kornhauser/generative-ai>

#### A.1.3 For further exploration: a selection of teaching and learning-related GenAI links

- Instructor-assignable (or modifiable) open-licensed course (from the University of Sydney): “AI in Education”.<sup>34</sup> <https://canvas.sydney.edu.au/courses/51655>
- Keynote address video (Dr. Chris Dede at Stanford University’s *AI X Education* conference. August, 2023): “If AI is the Answer, What is the Question: Thinking about Learning and Vice Versa”. <https://www.youtube.com/watch?v=1yZbUwcVR8k>
- Guiding principles and associated resources for deciding if/when to use GenAI (from Oregon State University’s Ecampus Unit): “Guidance for online course development and the use of artificial intelligence tools”. <https://ecampus.oregonstate.edu/faculty/artificial-intelligence-tools/>
- Resource page for teaching and learning (from the University of Central Florida): “Artificial Intelligence”.<sup>35</sup> <https://fctl.ucf.edu/technology/artificial-intelligence/>

## A.2 Research

### A.2.1 Chatbots

See Table 1 (p. 24).

### A.2.2 Programming

**Copilot** (GitHub/OpenAI) An LLM-powered code assistant. <https://github.com/features/copilot>

**Codeium** (Codeium) A free (but closed source) alternative to Copilot. <https://codeium.com>

### A.2.3 Non-text media

**DALL-E** (OpenAI) Generates images in a variety of styles (cartoon, photorealistic, etc.) in response to text prompts. <https://openai.com/dall-e-3>

**Whisper** (OpenAI) An automatic speech recognition (ASR) system, able to parse spoken language. For example, it can be used in conjunction with Kdenlive (a video editor) to autogenerate subtitles for videos. <https://openai.com/research/whisper>

---

<sup>34</sup>For students and built by students.

<sup>35</sup>A wealth of practical insights here, but of particular note relevant to UofL is the tip to protect university data using Microsoft Copilot.

Table 1: Most prominent LLM-based chatbots.

	ChatGPT		Claude		Copilot		Coral		Gemini		Perplexity	
Developers	OpenAI	Anthropic	Microsoft	Cohere	Google	Perplexity						
Free version	✓	✓	✓	✓	✓	✓						✓
Model	GPT-3.5	Claude 3 Sonnet	GPT-4	Command	Gemini Pro	GPT-3.5						
Paid version	✓	✓			✓	✓						✓
Behind paywall	GPT-4, DALL-E 3	Claude 3 Opus			Ultra 1.0	GPT-4(V)/Claude 3						
Price/month	\$20	\$20			\$20	\$20						\$20
Web search	🔒		✓	✓	✓	✓						✓
Text file input		✓			✓							
Image input	✓		✓			✓						✓
Image generation	🔒		✓			✓						✓
Voice input	🔒		✓			✓						✓
Voice output	🔒											✓
iOS/Android	✓		✓		✓	✓						✓

🔒 Paid version only.



#### A.2.4 Literature review

**Consensus** allows the user to ask a research question and find relevant papers, extracts “findings” from research, and synthesizes findings with a consensus percentage breakdown. <https://consensus.app>

**Elicit** also allows the user to ask a research question and find relevant papers, as well as providing one-sentence abstract summaries. Users can select relevant papers and search for more like them, extract details into tabular form, and synthesize themes and concepts. Elicit claims to be 90% accurate; it is unclear how accuracy is defined in this context. <https://elicit.com/>

**Keenious** finds relevant research by analyzing input text (typed, uploaded, or linked to). Users can highlight specific portions of the text to focus on. <https://keenious.com>

**LitSuggest** is hosted by the National Library of Medicine. The user provides a list of PubMed IDs (PMIDs) for relevant articles to train a model. The system outputs a list of similar articles in PubMed. <https://www.ncbi.nlm.nih.gov/research/litsuggest/>

**Scopus AI** allows the user to find research papers using plain-text queries. The system synthesizes findings of numerous research articles into “digestible summaries” and uses “advanced engineering,” which “limits the risk of hallucinations” (see §4.7.3). <https://www.elsevier.com/products/scopus/scopus-ai>

#### A.2.5 Research reporting guidelines

**CONSORT-AI** extends the CONSORT reporting guideline for clinical trials involving AI. <https://www.clinical-trials.ai/consort>

**SPIRIT-AI** extends the SPIRIT reporting guideline for clinical trial protocols involving AI. <https://www.clinical-trials.ai/spirit>

#### A.2.6 Journals and preprint servers

**JMIR AI** Subsidiary of the *Journal of Medical Internet Research*. <https://ai.jmir.org/>

**NEJM AI** Subsidiary of the *New England Journal of Medicine*. <https://ai.nejm.org/>

**arXiv** Preprint server, featuring research by OpenAI and others. <https://arxiv.org/>

#### A.2.7 The AI landscape

**AI Index Annual Report** Stanford University Institute for Human-Centered AI (HAI). Annual report on the state of AI globally. <https://aiindex.stanford.edu/>

**Awful AI** Curated list of harmful AI applications. <https://github.com/daviddao/awful-ai>.

**Epoch** Researchers investigating and forecasting the development of advanced AI. Includes datasets and visualization tools. <https://epochai.org/>

**Theres an AI for That** Aggregator of AI tools. <https://theresanaiforthat.com/>

## **B Charge Letter**

### **Committee on Use of Generative Artificial Intelligence in UofL Academics**

I am pleased to appoint each of you as members of the Committee on Use of Generative Artificial Intelligence in UofL Academics. Your expertise and diverse perspectives make you valuable assets to this committee, and I am confident that your recommendations will play a pivotal role in shaping the integration of AI technologies at our university.

#### **Purpose**

The purpose of this committee is to comprehensively assess the potential applications and impact of ChatGPT and generative artificial intelligence in the domains of academics and undergraduate/graduate student research at the University of Louisville. Our goal is to explore how AI can enrich learning experiences and empower our students and faculty to be at the forefront of technological advancements while maintaining the highest standards of ethical conduct and academic integrity.

#### **Scope of Work**

- **Identify Relevant Use Cases:** The committee shall identify and evaluate potential use cases for ChatGPT and generative AI in academic settings and student research projects. This could include, but is not limited to, content generation, data analysis, simulation, language translation, and data-driven research applications.
- **Enhance Teaching and Learning:** Assess how ChatGPT and AI-powered tools can be integrated into the teaching and learning process to improve student outcomes. Explore how generative AI can be harnessed to provide personalized learning paths, promote critical thinking, and address individual learning needs.
- **Ethical Considerations:** Thoroughly investigate the ethical implications of AI adoption in academics and undergraduate/graduate student research, with a focus on maintaining academic integrity, avoiding plagiarism, and ensuring that AI-generated content is appropriately attributed and used.
- **Training and Support:** Identify the training needs of faculty, researchers, and students to enable them to effectively and responsibly use AI technologies. Propose training programs and workshops to equip them with the necessary skills and knowledge for responsible AI adoption.

#### **Deliverables**

Based on the scope of work, the committee is expected to deliver the following:

- A comprehensive report highlighting potential applications of ChatGPT and generative AI in academics and undergraduate/graduate student research.
- Recommendations for ethical guidelines and policies for the responsible use of AI technologies in academic and undergraduate/graduate student research activities. This includes a review of the university's current academic dishonesty policy regarding cheating, fabrication, and plagiarism to incorporate work completed by entities that are not human.
- Recommend best practices, Training and support initiatives to prepare students and faculty for AI adoption.
- A final presentation to the university leadership, summarizing the committee's findings and recommendations.

#### **Timeline**

- The committee is expected to complete its work and submit its final recommendations by February 28, 2024.

Thank you in advance for your dedication and contribution to University of Louisville's academic excellence and undergraduate/graduate student research innovation.

