

John A. Burns School of Medicine (JABSOM)
Guidelines on Generative AI for Learners
May 2025

Introduction

Generative Artificial Intelligence (AI) represents a significant technological advancement, offering transformative possibilities in medical education, research, and administration. However, its use also raises critical considerations regarding academic integrity, data privacy, ethical standards, and professional conduct. This policy provides guidelines specific to the John A. Burns School of Medicine (JABSOM) for the appropriate and responsible use of generative AI by JABSOM learners.

This policy is designed to complement broader University of Hawai'i (UH) guidance on AI, and future policies, with additional specificity tailored to the unique context of JABSOM's MD Program and GME curricula and the critical importance of protecting patient and learner privacy under HIPAA and FERPA, respectively. This policy is subject to change to ensure alignment with UH, JABSOM, Hawai'i Residency Programs (HRP), and health system or clinical learning site policies.

Scope of Policy

This policy covers all forms of standalone and integrated generative AI, including but not limited to:

- **Large Language Models (LLMs)** such as Gemini or ChatGPT, used for generating text and future direct derivatives (i.e., Large World Models) as they evolve
- **Text-to-Image and Text-to-Video AI** tools, such as DALL-E, used for creating visual content.
- **Artificially Generated Audio**, including voice synthesis and deepfake technologies, such as Suno AI.
- **AI-driven Virtual Patients** and simulation tools, used in medical training.
- **Text-to-Speech and Speech-to-Text AI**, used in creating and transcribing educational and clinical content.
- **AI-driven Personalized Learning Systems**, used to customize educational experiences.
- **Retrieval-Augmented Generation (RAG)** an AI framework that combines traditional databases and search with the capabilities of LLMs.
- **AI that is integrated into enterprise software** such as Microsoft, Adobe, or Google products

General Guidelines

1. **Default Policy and Instructor Flexibility**
 - a. A "learner" is defined in this policy as a student educated within JABSOM (including but not limited to medical students and 'Imi Ho'ōla students), as well as those in graduate medical education programs for which UH JABSOM is the

John A. Burns School of Medicine (JABSOM)

Guidelines on Generative AI for Learners

May 2025

sponsoring institution, i.e., residents and fellows of UH residency and fellowship programs.

- b. A uniform approach to generative AI is the default across all JABSOM courses and modules. This policy applies unless instructors provide specific alternative guidelines. Appendix A elaborates on some default policies on generative AI use by learners.
 - c. Instructors retain the flexibility to customize generative AI use in their courses, provided that these expectations are clearly communicated to students at the outset and do not conflict with other aspects of the JABSOM or AI policy, or other state and federal policies, such as the use of sensitive information, etc. Any modifications should be explicitly stated in the course syllabi or communicated directly to students.
- 2. Generative AI as Educational Strategy (Androgogy)**
- a. Faculty are encouraged to explore the pedagogical potential of generative AI by integrating it into their teaching practices. However, they must ensure that students know these tools' limitations and are guided on how to use them responsibly.
 - b. Suggested educational strategies include assignments that critically evaluate AI-generated content, scenarios that require students to compare AI outputs with human-created materials, and discussions on the ethical implications of AI in medicine.
- 3. Professional Conduct and Oversight**
- a. Students and trainees are all subject to this policy and are expected to uphold the professional standards of JABSOM in their use of generative AI. This includes maintaining the confidentiality of sensitive information, avoiding any form of academic dishonesty, and contributing to a culture of integrity and respect.
 - b. Consent is always needed before an AI model incorporates another person's personal information, data, visual likeness, photographs, videos, voice, creative content, social media content, and any other outputs associated with that individual.
 - c. It is a breach of professionalism for any learner who, without explicit consent, falsifies representations of another person, using AI or otherwise. This may include falsified writing, images, videos, audio, and/or personal information and data of any kind.
 - d. In the above, a "person" may include any student, colleague, faculty, staff, patient, or other real human being.
 - e. Any use of AI that promotes harm, violence, discrimination, or graphic or explicit content is prohibited in the context of UH campuses and spaces, on UH-affiliated technological devices, or software resources provided by UH to its learners.
- 4. Recommendation for AI Training**
- a. While not mandatory, it is strongly recommended that faculty and students engage with available resources and training opportunities to deepen their understanding of generative AI. JABSOM will provide advisory materials and workshops to support the ethical and effective use of AI in medical education and practice.

John A. Burns School of Medicine (JABSOM)
Guidelines on Generative AI for Learners
May 2025

Academic Integrity

Maintaining academic integrity is paramount at JABSOM. As generative AI tools become increasingly integrated into educational practices, all learners, faculty, and staff must understand and adhere to the highest standards of honesty and transparency. In general, medical students may utilize generative AI as a tool, resource, or consultant and not as a substitute for their own knowledge acquisition, analysis, and self-reflection.

1. Use of AI in Assignments:

- a. Students are permitted to use generative AI tools for academic assignments only as described in Appendix A, or otherwise if explicitly allowed by the instructor. In such cases, the use of AI must be fully disclosed, and students must critically assess the output generated by these tools.
- b. The student must ensure that AI-generated content used in their work is accurate, relevant, and appropriately cited. AI should be used as a supplement to, rather than a replacement for, the student's own work.
- c. Instructors are encouraged to clarify in their syllabi and assignments what constitutes acceptable use of AI, including any specific limitations or expectations.

2. Disclosure of AI Use:

- a. When AI tools are used to generate content for an assignment, students must include a clear disclosure statement indicating how and to what extent AI was utilized. This disclosure should be placed at the beginning or end of the assignment, as directed by the instructor. See also Appendix A for additional guidance on disclosure in various contexts.
- b. Example Disclosure Statement:
 - i. *"In completing this assignment, I used ChatGPT to generate an initial draft of the literature review section. I reviewed and edited the AI-generated content to ensure accuracy and relevance. All AI-generated text has been critically evaluated, and the sources cited were independently verified."*
- c. Failure to disclose the use of AI in assignments may be considered a violation of academic integrity and could result in disciplinary action, as per JABSOM and UH policies.

3. Faculty Responsibilities:

- a. Faculty members are expected to model academic integrity by clearly stating their own use of AI tools in creating course materials, assignments, or assessments. This includes disclosing the use of AI for generating lecture content, exam questions, or any other educational resources. See [Appendix A](#) for additional examples and details.

John A. Burns School of Medicine (JABSOM)

Guidelines on Generative AI for Learners

May 2025

- i. Faculty may use generative AI to collate and summarize qualitative comments on assessments and evaluations, such as in summarizing multiple supervisor comments across a clerkship. When this occurs, the use of AI should be FERPA compliant (e.g., through de-identification and/or the use of UH-licensed and approved LLMs—See *Data Privacy and Security* below) and disclosed. The disclosure of AI use for summative comments would also carry over to a need for disclosure in subsequent documentation utilizing these summations, such as the MSPE.
 - b. Instructors should also guide students in critically engaging with AI-generated content, emphasizing the importance of accuracy, critical thinking, and ethical considerations.
- 4. Academic Misconduct:**
- a. Any attempt to use AI tools to deceive or mislead, such as generating work (e.g., assignments, reports, research) that is falsely presented as entirely original or failing to verify AI-generated content, will be treated as academic and professional misconduct.
 - b. At the course director's discretion, cases of suspected academic dishonesty involving AI by students will be brought to the Evaluation Remediation and Review Committee (ERRC). Depending on the severity of the violation, consequences may range from receiving a failing grade on the assignment to more severe disciplinary actions.
 - c. Misuse by GME learners will be reviewed by the program's director, faculty, and clinical competency committee as a breach of professionalism. Depending on the severity of the misuse, concerns with AI utilization will be reported to the Office of the Designated Institutional Officer as necessary for additional guidance.

5. Creation of New Intellectual Property

- a. A full dive into evolving copyright laws pertaining to creative works generated by artificial intelligence is beyond the scope of this policy; however, the following are key considerations:
 - i. Products of generative AI, such as generated images, music, videos, and other multimedia, as well as written works, are generally not able to be copyrighted as they were not considered to have been created by a human being.
 - ii. It is incumbent upon individuals at JABSOM creating media, products, inventions, ideas, and other potentially copyrightable outputs to be aware that AI-generated components of their works may not be copyrighted.
- b. Additionally, JABSOM advises against and discourages the “laundering” of previously copyrighted material through its re-generation via artificial intelligence, as this may lead to unethical or illegal acts of claiming rights or access.

John A. Burns School of Medicine (JABSOM)

Guidelines on Generative AI for Learners

May 2025

6. Monitoring for undisclosed generative AI Use in Education

a. AI Detection Tools

- i. Submissions of learners' work at JABSOM may be subject to appraisal with AI detection tools. While **JABSOM does not encourage using such tools** at this time, they are not forbidden. If faculty choose to use such tools in assessing a learner's work, this is only permissible when such tools and their use uphold all relevant data privacy and security standards, professionalism, and due diligence as otherwise delineated in this document and other JABSOM or UH policies.
- ii. Like all AI software used, any third-party tools for AI detection must be approved by UH.
- iii. Faculty are responsible for removing identifiable information (PII) in accordance with FERPA before running through such software. This includes indirectly identifiable information that may appear within the content of the submitted assignment.

b. Use of Activity Logging

- i. One potentially helpful mitigating strategy for demonstrating authentic, non-AI-generated task completion by learners is to leverage software that tracks the activity or progression of a document or use of an application. For example, word processing applications like Google Documents may offer such logging, so a learner performing a written assignment could demonstrate its manual completion without generative AI.
- ii. It is *not* assumed as a default policy that learners must routinely use such logging. However, individual faculty, educators, and course leaders reserve the right to request this function for specific assignments. When requesting this, it must be clearly stated in advance, with adequate guidance for the learners to easily and effectively utilize these tools. In cases where logging was *not* specifically requested before the assignment, it is unreasonable to require the learner to procure such logs retroactively. For example, if there is a concern for inappropriate use of generative AI, and no specific logging requirement was stated before the assignment. Decisions about whether inappropriate use occurred cannot be predicated on the existence or absence of such a log.
 1. Note that this policy element may require revision as technology evolves, such as logging becoming so ubiquitous and automatic that it becomes more reasonable to request without prior warning. Pending such revisions, logging functionality will be considered opt-in.

John A. Burns School of Medicine (JABSOM)
Guidelines on Generative AI for Learners
May 2025

Data Privacy and Security

The protection of sensitive information, including patient health information (PHI) and student records, is of utmost importance at JABSOM. Generative AI tools must be used responsibly to ensure compliance with HIPAA, FERPA, and other relevant data protection laws.

1. Prohibited Use of Sensitive Information:

- a. Generative AI tools must not be used to process, store, or transmit any sensitive information covered by HIPAA, FERPA, or other applicable data protection laws. This includes, but is not limited to, patient health information, student records, and other forms of personally identifiable information (PII).
- b. Sensitive data such as that covered by HIPAA, FERPA, or other applicable data protection laws, or protected by copyright, may not be used to train or fine-tune an AI model without all appropriate and applicable forms of consent.
- c. Exceptions may include the use of AI platforms that have been deemed compliant, such as through a business agreement with the university. Additional research agreements for using a platform may be required when used in research contexts. JABSOM learners seeking to pursue compliant AI models in this way must pursue this through relevant course directors or supervising faculty, initially. They, in turn, may contact the JABSOM Associate Dean for Academic Affairs for inquiries about such business agreements and related health system-approved platforms.
- d. Example of Inappropriate Use:
 - i. A student uses a generative AI tool to transcribe notes from a patient interview, including the patient's name, date of birth, and specific medical conditions, and inputs this data into an open AI platform like ChatGPT. This constitutes a breach of HIPAA, as the patient's PHI is exposed to a non-secure, third-party system.

2. De-identification and Secure Handling:

- a. If there is a legitimate educational or research need to use AI tools with data derived from sensitive information, the data must first be deidentified to remove any personally identifiable information. De-identification should be conducted according to the highest standards, ensuring no individual can be readily identified from the data.
- b. If used for research purposes, Institutional Review Board (IRB) approval must be obtained before use. The IRB or a privacy board must approve any waivers or alterations of authorization. Additionally, consent and authorization forms related to participants in research should note the potential use of their data with AI technologies, as applicable.
- c. It is the responsibility of the user to verify the accuracy of any AI-derived analysis, conclusions, or other outputs.

John A. Burns School of Medicine (JABSOM)

Guidelines on Generative AI for Learners

May 2025

d. Examples of Appropriate Use:

- i. Quality Improvement: A student is working on a quality improvement project with a faculty mentor and wants to use an AI tool to analyze trends in patient care at the faculty member's clinic. The goal is to optimize their personal clinical practice, and it is not for dissemination as research. The team also runs this project through the IRB and confirms that it is deemed Not Human Subjects Research (NHSR), given its quality improvement nature. Before inputting the data, they remove all identifiers such as names, dates of birth, and any specific details that could trace back to an individual patient. The deidentified data is then processed using an AI tool hosted on a secure, UH-approved platform designed to handle such data. This approach maintains compliance with HIPAA while leveraging AI for such activities.

e. Examples of Inappropriate Use

- i. Risk of indirect identification: A student conducting a mentored medical education research project for MD5 interviewed five students from underrepresented backgrounds about their experiences of inequity in education throughout their lives. The student obtained IRB approval to interview the students and perform transcription and thematic content analysis but did not have the idea to use AI to assist until now. The student researcher transcribes the recordings and redacts them anywhere the person's name is mentioned. Without additional review for other risks of identification, the essays are fed into chatGPT (without any specific settings like Temporary Chat) so that it can assist with thematic analysis. When finished, the faculty deletes that conversation from the list of prior conversations in chatGPT.
 1. This example demonstrates inadequate precautions, given the potential presence of indirectly identifiable, sensitive aspects of the interview participants' stories. These narrative essays may pose a high risk for students to disclose direct or indirect identifiers, such as through comments about their background, personal views, or experiences or through transcribed speech patterns alone. In this scenario, multiple additional measures may be needed to safely use LLMs to support the project, potentially including some combination of:
 - a. [Obtaining consent from students](#)
 - b. Manual review for indirectly identifiable content requiring redaction
 - c. Identification of a lower-risk LLM, such as an institution-specific, approved LLM, which will not save or train on the inputted student data.

3. Secure Platform Usage:

- a. All members of JABSOM must ensure that sensitive or potentially identifiable data is never input into open or unsecured AI platforms. Approved, secure

John A. Burns School of Medicine (JABSOM) Guidelines on Generative AI for Learners

May 2025

platforms should be used exclusively when dealing with these forms of high-risk data.

- b. Students must be trained to recognize and use these secure platforms appropriately, avoiding the temptation to use freely available AI tools that do not meet JABSOM's security standards.
- 4. Emphasis on Proactive Consent Form Implementation**
- a. In research contexts, consent forms may need to disclose the use of AI with participant data, and the IRB should provide guidance therein.
 - b. In educational contexts, we recommend implementing consent forms, such as for learners either upon matriculation into the school and/or at the start of specific courses or assignments that are known to particularly depend on inputting learner-related data into artificial intelligence.
- 5. Consequences of Misuse:**
- a. Students' misuse of generative AI tools in ways that compromise sensitive information will be subject to review by the relevant course director, mentor, or other supervisor as applicable. It may be referred to the ERRC on a case-by-case basis. Depending on the severity of the breach, consequences may include retraining, loss of access to AI tools, or more severe disciplinary actions in cases of gross negligence or repeated violations.
 - b. Misuse of this nature by GME learners will be reviewed by the program's director, faculty, and clinical competency committee as a breach of professionalism. Depending on the severity of the misuse, concerns with AI utilization will be reported to the Office of the Designated Institutional Official as necessary for additional guidance.

Sustainability and Environmental Concerns

The growing use of artificial intelligence, including generative AI has been observed to significantly increase the use of natural resources, including land, water, and energy, with multiple potential downstream effects, including significant increases in carbon dioxide emissions. Particularly given JABSOM's aspiration to uphold the University of Hawaii's identity as a Native Hawaiian place of learning, we recommend that JABSOM's present and future leaders and community continue to embody values of aloha 'aina and malama honua by also considering the environmental impacts of our AI use. While the present policy does not offer specific recommendations to guide the quantity or nature of generative AI utilization by its end users based on environmental impact, JABSOM or its parent university may someday be positioned to do so. In the meantime, we encourage JABSOM and the University of Hawaii to instead focus on advocacy for sustainable practices in data management locally, nationally, and globally.

Conclusion

John A. Burns School of Medicine (JABSOM) Guidelines on Generative AI for Learners

May 2025

Generative AI is a powerful tool that, when used responsibly, can enhance medical education, research, and administration at JABSOM. This policy is intended to guide the JABSOM community in navigating the opportunities and challenges presented by these technologies, ensuring that their use aligns with our values of integrity, professionalism, and respect for privacy.

By following these guidelines, JABSOM will continue to lead in medical education while safeguarding the ethical standards essential to our mission.

Disclosure of AI Use in Policy Authorship

OpenAI's ChatGPT was used in the information gathering, initial brainstorming, and early first draft outline to create this policy. Subsequent revisions and additions were made manually by the JABSOM Education Technology Committee without further contribution from artificial intelligence. All portions of this document, including those generated by AI have been reviewed by the team for accuracy.

Relevant Institutional Policies

Adhere to the rules that apply to all other aspects of your responsibilities as a JABSOM student, including professionalism and integrity, confidentiality, privacy, and security. The University of Hawaii is a public institution, and JABSOM students are expected to uphold their commitment to safeguarding and preserving state resources. Relevant institutional policies include, but are not limited to, the following listed policies and any subsequent revisions. Students are responsible for being knowledgeable about current institutional policies.

Current Institutional Policies

Professionalism and Integrity

- A. JABSOM Policy – Guidelines for Appropriate Appearance and Attire
- B. JABSOM Policy – Medical Student Mistreatment Guidelines and Procedures
- C. JABSOM Social Networking and Policy Guidelines
- D. UH Executive Policy EP7.205 – Systemwide Student Disciplinary Sanctions
- E. UH Executive Policy EP7.208 – Systemwide Student Conduct Code
- F. UH Executive Policy EP12.214 – Conflicts of Interest and Commitment

Confidentiality, Privacy, and Security

- A. JABSOM GME Policy - HIPAA
- B. JABSOM Policy – HIPAA
- C. JABSOM Policy – Public Computer Use
- D. UH Executive Policy EP2.210 – Use and Management of Information Security Resources
- E. UH Executive Policy EP2.214 – Institutional Data Security Classification Categories and Information Security Guidelines
- F. UH Executive Policy EP2.215 – Institutional Data Governance
- G. UH Executive Policy EP2.217 – HIPAA Policy

Applicable State Laws

- A. Hawai'i Revised Statutes Chapter 84-10 (Code of Ethics)
- B. Hawai'i Revised Statutes Chapter 708-890 (Computer Crime)

John A. Burns School of Medicine (JABSOM)
Guidelines on Generative AI for Learners
May 2025

Appendix A: Default Permissible Use of Generative AI

In the absence of specific, alternative instructions from an instructor, mentor, or supervisor, the following default rules apply to the use of generative AI tools by students and trainees at JABSOM:

1. **Grammar, Syntax, and Idea Generation:**
 - o **Permitted:**
 1. Students may use generative AI tools to assist with grammar and syntax corrections in their writing. AI tools can also be used to generate ideas, such as brainstorming topics for research projects or determining effective teaching formats.
 2. AI may be used to organize existing information into outlines or other educational formats.
 - o **Not Permitted:** Generative AI is not allowed to produce differential diagnoses, treatment plans, or other clinical reasoning outputs. These tasks require independent critical thinking and must be performed by the student.
2. **Content Generation:**
 - o **Not Permitted:** AI tools should not be used to generate substantive content for assignments, such as writing sections of essays or reports. Any use of AI allowed under item 1 must be disclosed appropriately to ensure transparency.
3. **Critical Analysis and Verification:**
 - o **Required:** Any AI-generated content included in an assignment must be critically analyzed and verified by the student to ensure accuracy and relevance. This applies to all allowed uses of AI under this policy.
4. **Data Privacy:**
 - o **Strictly Prohibited:** Students must not use generative AI tools to process, store, or transmit any sensitive or identifiable information. Any such use must first be discussed with the appropriate faculty, mentor, or supervisor before proceeding.
5. **Sharing of JABSOM Curricular Materials and Intellectual Property**
 - o **Permitted:** Learners may upload their own responses to assignments into third-party LLMs or AI databases.
 - o **Permitted with Limitations:** If rubrics, grading criteria, or similar materials are available to a learner, such as in the context of an assignment or upcoming assessment, such scoring criteria materials may be uploaded to third-party LLMs to prepare for assessments or seek feedback on a learner's work.
 - o **Strictly Prohibited:** Sharing of instructional materials, such as syllabi, problem-based learning (PBL) case materials, simulation materials, and lecture slides, is not permitted unless with the written permission of all educators or creators involved in the content creation.
6. **Use in Final Submissions:**
 - o **Permitted with Limitations:** AI-generated non-textual content, such as images, art, video, music, or multimedia elements, may be included in final submissions

John A. Burns School of Medicine (JABSOM)
Guidelines on Generative AI for Learners

May 2025

and presentations, provided that the AI-generated material is properly cited. However, the use of AI to generate textual content in final submissions is not permitted.

7. Group Work and Collaboration:

- o **Permitted:** Generative AI tools may be used in group assignments or collaborative projects, following the same guidelines as individual work.

8. Presentation and Multimedia Projects:

- o **Permitted:** Students may use generative AI to create visual aids, such as slides or videos, for presentations. All AI-generated materials must be properly disclosed and cited.

9. Exams and Quizzes:

- o **Strictly Prohibited:** The use of generative AI tools is not allowed for take-home exams, quizzes, or any assessments intended to measure individual knowledge and understanding.

10. Simulations and Virtual Patients:

- o **Not Permitted:** The use of generative AI in simulations, virtual patient interactions, or any clinical education scenarios is not allowed unless explicitly permitted by the instructor.

11. Faculty Use for Instruction:

- o **Permitted with Disclosure:** Faculty may use generative AI tools to create teaching materials, including lecture notes and exam questions. However, any use of AI for creating multimedia content (e.g., images, audio, videos) or gathering information must be disclosed to students, and AI-generated content must be properly cited.