

Artificial Intelligence (AI) Use in Research and Scholarly Projects

Guidance for Medical Education Sponsored Research

Kaiser Permanente Northern California (KPNC)

Version: October 11, 2026

From Tracy Lieu and Juleon Rabbani, with input from many others¹

Purpose: To recommend best practices for using AI in research and scholarly projects in order to support effective learning and academic integrity, and to mitigate risk to patient confidentiality.

Audience: Learners (residents, fellows, medical students), clinicians, researchers, and staff with medical education-sponsored research or scholarly projects. Medical education-sponsored projects are those supported by graduate medical education (GME), undergraduate medical education (UME), or continuing medical education (CME).

Focus: This guidance addresses AI used to support research-related tasks, not AI evaluated as a clinical intervention in a research or scholarly project. Research or evaluation of AI use as a clinical intervention requires organizational review by KP.

Key Points

- **DO use AI as an assistant, not a creator or decision maker**
 - Employ AI to generate search terms, organize information, and refine writing
 - Maintain human responsibility for key steps, including selecting study design, forming analysis plans, interpreting results, drawing conclusions, and writing first drafts of presentations and manuscripts
- **DO verify all content generated by AI**
 - Use dedicated search engines such as PubMed rather than relying solely on AI summaries; read original references and confirm the accuracy of citations
 - Independently validate statements generated by AI
- **DON'T put confidential patient information into AI tools**
 - Never enter patients' protected health information (PHI) into AI tools, even KP-approved tools such as CoPilot, for research or scholarly projects.
 - Don't enter confidential or sensitive information, even if it doesn't include PHI, into AI tools not approved by KP
- **DO support accountability in AI use**
 - Document how AI was used throughout each project
 - Disclose AI use in manuscript and scholarly products

Best Practices

¹ This draft reflects input from KPNC Division of Research (DOR) Biostatistical Consulting Unit Analytic Programming Team managers Lue-Yen Tucker, Amanda Jain, Zahra Samiezade-Yazd, and Julia Wei; DOR subject matter experts Cathy Chou (regulatory and compliance), Christine Moreno (IT user support), and Jingrong Yang (research analytics); Eve Melton, Director of the KP Libraries; Ted O'Connell, Director of Medical Education, and Theresa Azevedo-Rousso, DIO and Regional Senior Director, UME and GME. See Appendix C for Methods.

Others who have reviewed and given helpful input include: GME DMEs and RPMs, CRC chair Betty Suh-Burgmann, TPMG Chief Research Officer Doug Corley, DARE director Jacek Skarbinski, and DOR director Rachel Ramoni

	Uses of AI that are:		
Task	Recommended or acceptable	Not recommended or prohibited	Rationale and evidence
Literature searching	Identify potentially relevant references - Note: General-purpose generative AI tools such as Microsoft Copilot may provide Web-only searching. They do not necessarily search scholarly databases such as PubMed directly and cannot access scholarly materials that are behind paywalls.^a Generate a summary of an article or set of articles to help the user gain context Identify or generate potential search terms or keywords to use in PubMed or other scholarly databases	Generate a reference list without verifying each source against the original publication Reading only summaries without reading the original publications Generate a systematic review or meta-analysis without expertise in these methods Prohibited: Using KP-licensed scholarly articles in AI tools that are being trained on user provided content (this includes most general-purpose chatbots, such as ChatGPT)	Reviewing original articles provides essential context and supports deeper, more accurate interpretation. It also ensures that you are not using a non-existent hallucinated reference. Fabricated citations can be detected by algorithms. The rate has been increasing since 2023 and journals are being encouraged to adopt automated reference verification [Topaz Lancet 2026] Most licensing agreements do not allow for uploading copyright-protected works to third-party platforms, including Generative AI tools.
Project design	Suggest potential subtopics within a broader area of study Suggest ways to organize information Suggest options for study designs	Draft the initial project outline or analysis plan without the study leaders first thinking through the objectives Select a study design without real-time discussion with collaborators Select a study design without consideration of available resources or expertise	Developing a project outline or analysis plan and selecting a study design require subject-matter expertise, knowledge of available resources, conceptual judgment, and collaboration. AI cannot replace these inputs and may short-circuit the learning that medical education sponsored projects are intended to foster

	Uses of AI that are:		
Task	Recommended or acceptable	Not recommended or prohibited	Rationale and evidence
Identifying what regulatory approvals are needed, e.g. Institutional Review Board (IRB)		Using AI as the main source of advice on whether a study needs review by the IRB, Research Determination Official team, or what arrangements are required for data sharing	Input from local regulatory experts is needed to make these determinations.
Data collection	Data analysts may use KP-approved AI tools or methods to collect or extract data from medical records or images under protocols approved by the IRB, with tools approved by KP and DOR IT.^b Evaluate the validity of data extracted via AI or other algorithms via chart review	Prohibited: Entering PHI, including visit notes, into general-use AI tools such as Copilot, even those that are KP-approved for other purposes Using AI tools that are not KP-approved Using AI tools to extract data without clinician validation of accuracy	See details in Appendix A.
Data analysis	Use KP-approved AI tools to: Review and debug human-generated programming code Generate template programming code (starter code) with human validation Conduct simple analyses of deidentified datasets, such as creating frequency tables or other univariate statistics that can be cross-checked by the user	Prohibited: Enter personally identifiable information, protected health information (PHI), confidential data or any information not safe to share publicly into an unapproved or public AI tool Prohibited: Enter PHI into Copilot (see Appendix A) Alter clinical images with AI without full disclosure Conduct complex or cutting-edge analyses	Entering identifiable or confidential data into an AI tool may breach privacy and confidentiality requirements Human expertise should guide the final interpretation of results because AI-generated interpretations may be inaccurate or incomplete Nature and some other journals do not allow use of generative AI to alter (clinical or original) images, to maintain research integrity, consent, privacy, and IP protection

	Uses of AI that are:		
Task	Recommended or acceptable	Not recommended or prohibited	Rationale and evidence
		that require iterative decision making Develop final interpretations or conclusions from analysis results	
Manuscript writing	Suggest writing refinements after the researcher has prepared a first draft Suggest structures for, or refine, tables and figures Suggest ways of streamlining text Maintain documentation and plan to clearly disclose how AI tools were used in developing the study, analysis, and manuscript Ask AI for input, then implement the edits or other changes yourself If AI tools are used to generate tables or figures, double check all numbers and data representations	Draft full sections describing the findings, discussion, or conclusions Copy substantial AI-generated text without disclosure Use AI to write letters to the editor or opinion papers (prohibited by some journals) Cite AI-generated references without verifying them against the original articles Allow AI to make final decisions about what is included in a paper or other scholarly product	Using AI too early in the writing process may hinder creative and conceptual thinking, as described in this article: What 370,000 College Essays Tell Us About A.I.'s Effects on Creativity AI systems may generate errors. Their statements need to be checked by human domain experts for accuracy, bias, relevance, and reasoning [Hosseini Accountability in Research 2024] Copying substantial AI-generated text without disclosure may constitute plagiarism Many journals require disclosure of how AI tools were used in manuscript writing, including which tools and versions were used
Poster creation	The relevant notes and activities from the above section on “Manuscript writing” apply		

^a From Eve Melton, KP Libraries: There are AI research tools that do surface information from scholarly databases and/or peer-reviewed published articles, but so far none have been approved for use by Kaiser employees. Examples are: Research Rabbit, Elicit, Consensus AI, Semantic Scholar, OpenEvidence (pending approval).

^b This includes the use of programming code that implements AI methods, and specific tools approved for research, such as Linguamatics that conducts natural language processing with medical records.

Related recommendations

1. See the additional guidance in “Use of AI tools for research purposes” in Appendix A
2. As an author, you take responsibility for the accuracy of the content you submit.
3. Document the tasks for which AI was used in your project.
4. Plan to disclose how you used AI in your project when you submit your manuscript, even if you only used it for editing. See this [2026 JAMA editorial](#) and the [2024 version](#) for helpful specifics.
5. If you are peer reviewing a grant proposal, abstract, or manuscript, do not enter all of the content or any sections of it into an AI program. This may violate confidentiality.

References

Flanagin A, Perlis RH, Bibbins-Domingo K. Updated guidance for author use of AI in medical publication. *JAMA* 2026 Aug 10; doi:10.1001/jama.2024.3471.

Flanagin A, Pirracchio R, Khera R, Berkwits M, Hsuen Y, Bibbins-Domingo K. Reporting use of AI in research and scholarly publication—JAMA Network guidance. *JAMA*. 2024;331(13):1096-1098. doi:[10.1001/jama.2024.3471](https://doi.org/10.1001/jama.2024.3471)

National Institutes of Health. Artificial Intelligence in Research: Policy Considerations and Guidance. URL: [Artificial Intelligence - Office of Science Policy](#). Accessed July 27, 2026.

University of Michigan: Michigan Institute for Data & AI in Society. Generative AI for Research Guide. URL: [Generative AI for Research Guide – MIDAS](#). Accessed August 10, 2026.

University of Pretoria. Artificial Intelligence Guidelines: How can you use AI for research? URL: [How can you use AI for Research? - Artificial Intelligence Guidelines - Subject Guides at University of Pretoria](#). Accessed August 10, 2026.

Yale Standards for Using AI Tools and Chatbots in Health Sciences. URL: [Yale Standards for Using AI Tools and Chatbots in Health Sciences | AI at Yale](#). Accessed August 10, 2026.

Yale: Using AI Safely and Responsibly in Health Care and Research. [Using AI Safely and Responsibly in Health Care and Research | AI at Yale](#). Accessed August 10, 2026.

Appendix A

Use of AI tools for research purposes (from KP NCAL Division of Research IT, June 2, 2026):

When using AI platforms/tools in your research:

- Use AI platforms/tools that have been approved by Kaiser Permanente for KP work.
 - Currently, KP has approved only Copilot and Abridge for work use.
 - For a list of AI tools available at KP, see this webpage: [KP Artificial Intelligence \(AI\) / Generative AI \(Gen AI\)](#)
 - See also **Appendix B**
- Never use your personal AI accounts to conduct KP business.
- Do not enter PHI, PII, or other confidential/sensitive information into any AI platform/tool, even if the platform/tool has been approved by KP.
 - Information uploaded or entered into unapproved or public AI tools could be shared outside KP.
 - Current KP guidance specifies that PHI should *not* be used in any AI tool, even licensed premium or standard versions of Copilot at KP.

The above guidance is from a post from DOR in July 2026 titled [“Are you using AI in research at DOR appropriately?”](#) You need to log into the KP network to access it.

For additional details, see:

[CoPilot PHI Use Case Guidance : DOR-SUPPORT](#)

[Microsoft Copilot Usage Guidelines](#)

[GitHub Copilot at KP-DOR: Permitted Use & Compliance Guide : DOR-SUPPORT](#)

If there is a situation where PHI needs to be entered into a generative AI model (an LLM-based application), a user can make an application to KP’s NexusOS AI Council. However, this has not yet been attempted for a UME/GME/CME-sponsored project. The relatively small size of these projects may mean that this approach does not warrant the effort required for the approval and setup process.

[NexusOS | AI Enablement Platform](#)

Never enter the following into any public or unapproved AI tool:

- PHI, such as after-visit summaries or visit notes
- PII, such as member names or addresses
- Credit card or payment information
- Confidential, private, or restricted KP business information

For a detailed definition of PHI, see this [document](#). You will need to log into the KP network to access it.

Additional up-to-date guidance from KP Northern California’s Division of Research (DOR) IT on the use of AI in research is available at this intranet site:

<https://insidedor.kaiserpermanente.org/technology/ai-technology/> You need to log into the KP network to access it.

Appendix B

KP AI Tools – PHI and Non-PHI Use Matrix
from Christine Moreno of DOR IT, Sep. 2026

Current version: [KP AI Tools - PHI and Non PHI Use Matrix : DOR-SUPPORT](#) [You must be on the KP network to access this]

Note: Use the above link to access the most current version of this information. A recent version is provided below for convenience.

KP AI Tools: PHI and Non-PHI Use Matrix

Purpose: Quick reference for KP workforce members choosing an approved or enabled AI tool.

Important: This matrix summarizes current KP guidance. A tool being listed does not create blanket approval for every use case. Follow KP policy, minimum-necessary rules, local procedures, and required AI Council, Privacy, Compliance, Legal, Security, and Enterprise Architecture reviews.

Status legend

- ☐ **Yes** — permitted for the stated use when normal KP controls are followed.
- ☐ **No** — do not enter or process PHI.
- ☐ **Conditional** — only in a specifically approved, governed implementation; not blanket approval.
- ☐ **Not established** — the reviewed guidance did not explicitly authorize PHI use.

Tool matrix

Tool	Use without PHI	Use with PHI	Conditions and restrictions
Abridge	Yes	Yes — approved care-delivery use	Use for approved clinical workflows. Obtain consent from all participants before each encounter and maintain clinician/human oversight.
Epic AI / Epic Nebula features	Yes	Yes — only within approved Epic/KPHC workflows	Features must pass KP AI governance, architecture, Compliance, Legal/regulatory, responsible-AI, and KPHC feature reviews. Users must verify AI-generated content.
Microsoft 365 Copilot (licensed)	Yes	No	Don't enter PHI, PII, sensitive confidential KP information, or

Tool	Use without PHI	Use with PHI	Conditions and restrictions
			patient/member-facing content into prompts. Regulated or high-impact use cases require a formal AI request; don't assume this changes the PHI prohibition.
Microsoft 365 Copilot Chat / Copilot Enterprise	Yes	No	Enterprise Data Protection is enabled, but PHI remains prohibited. Don't use for patient/member-facing interactions.
Microsoft Copilot Studio	Limited / conditional	Not established	Evaluation-scoped or limited availability in the reviewed guidance. No general production or PHI approval. Requires enterprise governance, controls, and monitoring before deployment.
Avaamo	Yes — approved AI utility platform	Not established	The reviewed sources describe approved conversational workflows but don't provide blanket PHI authorization. Obtain use-case-specific governance and Compliance approval before processing PHI.
Azure AI Services, including Azure OpenAI in KP Azure	Yes — in approved KP environments	Conditional	PHI may be used only in a specifically approved KP implementation. Required steps include AI Council approval, Enterprise Architecture approval, NexusOS AI Platform onboarding, and Azure onboarding. Services must be accessed through KP Azure for required security and legal protections.

Tool	Use without PHI	Use with PHI	Conditions and restrictions
Azure Databricks in KP Azure	Yes — in approved KP environments	Conditional	Use PHI only in an approved, governed workload with applicable privacy, security, data, and AI controls. There is no blanket PHI approval for every Databricks use case.
KP Thrive AI	Yes — for approved KP workloads	Conditional	KP-managed on-premises platform intended to run AI close to KP data. PHI use must be tied to an approved workload and responsible-AI governance; listing the platform isn't blanket PHI approval.

Quick decision guide: [KP Artificial Intelligence \(AI\) / Generative AI \(Gen AI\)](#)

1. **If PHI is involved, don't use Microsoft 365 Copilot, Copilot Chat, or Copilot Enterprise.**
2. **For approved clinical workflows, use Abridge or approved Epic AI/Nebula features** and follow workflow-specific consent, oversight, and verification requirements.
3. **For custom solutions involving PHI, use only a specifically approved KP implementation** in KP Azure, Databricks, or Thrive AI after all required governance and onboarding steps.
4. **If the guidance says “conditional” or “not established,” stop and obtain approval before entering PHI.**
5. Review Strategic Technology Connect for the current technology lifecycle status and submit new AI solutions to the AI Council before adoption.

Sources reviewed

6. **KP Artificial Intelligence (AI) / Generative AI (Gen AI)** — Enterprise Architecture Portal technology guidance.
7. **Microsoft 365 Copilot** — KP Collaboration Solutions onboarding and use guidance.
8. **Microsoft 365 Copilot Chat** — KP Collaboration Solutions guidance.
9. **AI Guidance 06-10-26** — KP compliance guidance identifying Copilot and Abridge restrictions.
10. **Epic AI Overview @ KP** and **Generative AI in Epic General Overview** — Epic/KPHC governance and workflow guidance.
11. **Artificial Intelligence (AI) Strategy 2025** — KP platform strategy for Avaamo, KP Azure AI, Databricks, and Thrive AI.
12. **Azure OpenAI Service Engagement** — required onboarding and governance steps.

Because AI products and approvals change, confirm the latest guidance before starting a new or regulated use case.

Appendix C – Methods

To develop this guidance, we:

1. Reviewed or requested any existing published and publicly accessible guidance on the use of AI:
 - a. In research – published guidance from selected high-impact journals
 - b. In research - from KP internal departments including KPSOM and IT departments
 - i. KPSOM, SCAL (HI), NCAL, WA, NW, CO, GA, MAS
 - c. In medical education (for students, residents, and fellows) – published reference documents
2. Gathered examples from medical schools ranked within the top 15 of NIH funding recipients, plus selected other medical schools
3. Created a first draft
4. Invited input from subject matter experts within the KPNC Division of Research (including the Biostatistical Consulting Unit), plus Eve Melton of the KP Libraries
5. Created a second draft
6. Invited input from directors of medical education leaders, research project managers, and specific KPNC research leaders
7. Finalized recommendations for fall 2026
8. Plan to update in 6 to 12 months