UCI School of Medicine
Office of Education Compliance and Quality

# Twelve Tips to Guide the Use of Data with Generative AI for Accreditation and Continuous Quality Improvement

Authors:
Julie Youm, PhD
Zahra Dabzadeh

March 30, 2024

**Summary**

This resource explores the use of generative AI applications for accreditation and CQI activities in medical education. The tips are organized to address two phases of generative AI use. The first phase focuses on activities that are recommended prior to the use of generative AI including preparation of the data and data protection, privacy and ethical considerations. The second phase focuses on relevant applications of generative AI to support data management and analysis. While we acknowledge limitations of these tips due to the rapidly changing generative AI tools and the expertise surrounding its use, we find implications for sharing these recommendations early to encourage the broader adoption of generative AI in medical education through a deliberate and informed perspective.

**Introduction**

OpenAI's ChatGPT launched on November 30, 2022 and had a broad and notable impact on the mainstream interest to artificial intelligence (AI) in a manner previously unseen (Edwards 2023). In particular, the hype of generative AI has dominated the AI discussion and brought curiosity often associated with those of technology backgrounds to the average everyday user. Prior attention about AI most often focused on machine-learning, where a model is trained to make predictions based on a set of data. The current attention on generative AI is a result of the freely available tools including ChatGPT, Microsoft Copilot (Microsoft 2024) and Google Gemini (Google 2024).

In simple terms, a generative AI model differs from machine-learning in that it learns to create new data from the dataset rather than focusing on making predictions (Zewe 2023). The newly generated data can take on different forms, for example, text, images, audio, based on the data the generative AI model is trained upon. The ability to generate data in this way offers potential in many areas of health professions education. In this article, we specifically explore the use of generative AI to support accreditation and continuous quality improvement (CQI) efforts in medical education. Accreditation and CQI are activities that are inherently dependent on data and analysis.

There have been clear applications of the predictive capabilities of machine-learning for solving medical education problems like identifying students at academic risk based on a summative assessment data or understanding the impact of pathway programs on improving the diversity of a medical school applicant pool. The nexus of generative AI encourages us to consider how to move accreditation and CQI efforts forward with new and progressive approaches. However, the complexity and evolving nature of generative AI warrants a degree of mindfulness and caution for those who will engage with it. These twelve tips present an initial set of considerations for medical educators as they pursue the integration of generative AI in their own accreditation and CQI practices. These tips are divided so that the first six tips present guidance on preparing to use data with generative AI, and the second six tips focus on using and analyzing data with generative AI.

## Preparing to Use Data with Generative AI

**Tip 1: Review classification levels for your data**
Generative AI is predicated on data from which it can train. As members of a medical school accredited by the Liaison Committee on Medical Education (LCME) (LCME 2024) medical school or a large academic institution, awareness and compliance with institutional policies and practices that protect data are necessary. There are many types of data that are used in practice and the level of protections that are needed will differ accordingly. Formalizing the classification of data helps to appropriately align policies and practices across an organization. Data classifications are often required for federal regulations and state requirements as well (Educause 2015).

Classifications can include levels of data protection (concerns related to confidentiality and security), availability (the business impact if loss of availability) and integrity (accuracy, consistency and completeness). While researchers and clinicians focus on data security as a core aspect of their roles, it is imperative that all members of the institution fully understand the scope of data classifications as related to their daily work, especially as these are often discussed from an information technology perspective and may not seem relevant to administrative staff, teaching faculty and students. However, in October 2023, Educause, a nationally recognized association whose mission is to “to advance higher education through the use of information technology” (Educause 2024), conducted a poll of the Educause community (Burns, Robert and Muscanell 2023) and found that only a third of respondents reported institutional privacy policies (35%, N=162) included as part of campus data security awareness trainings. Further, less than a quarter of respondents (23%, N=162) reported the inclusion of institutional data governance policies in these trainings.

It is critical for faculty and staff who work with accreditation and CQI data to know relevant institutional, state and federal policies and practices before they use generative AI to support their efforts. For example, if evaluating student performance on the USMLE Step 1 examination as a programmatic outcome measure, the data should not contain student identifiers when reporting on this to remain compliant with the federal Family Educational Rights and Privacy Act of 1974 (FERPA). Identifiers that one may be commonly aware of include name and email address. There are also those that are not always common in a dataset such as the Association of American Medical Colleges (AAMC) identifier that is used by the National Board of Medical Examiners (NBME) to track medical student examination data and should have additional considerations when shared as input into generative AI tools.

**Tip 2: Connect with the people that have a role in data management for input and any required approval/authorization**

In addition to the classification of data, an institution may have a policy on data governance, the exercise of authority and control over the management of data (DAMA International 2017). Data governance policies also include identification of the people who participate in the management of data. Connecting with the people who have a role in the management of the data you need to access should be a foundational step in your use of the data.

The people who are involved in data management can take on different types of roles. The Data Management Association International (DAMA-I) outlines several of these roles (DAMA International 2017), which are differentiated based on the role's position in the organization and/or its function:

- Chief Data Stewards: Lead data governance efforts.
- Executive Data Stewards: Active in data governance efforts in senior management positions.
- Enterprise Data Stewards: Provide oversight of a data domain across business units.
- Business Data Stewards: Provide subject matter expertise for a data domain to define and control data.
- Data Owner: Serve as a business data steward with approval authority within their data domain.
- Technical Data Stewards: Operate with one of the data management knowledge areas like data integration, database administration, business intelligence, data quality, and metadata management.
- Coordinating Data Stewards: Lead and represent teams of business and technical data stewards in interactions with Executive Data Steward and in discussions across business units.

Each role in the data governance process needs to be involved in considering the implications of generative AI from their relevant perspectives. Simply having access to the data is not enough. For example, if medical school applicant data is needed for an accreditation team to evaluate the effectiveness of pre-matriculation pathway programs, they must connect with the admissions team if they are the data owners for the applicant data to determine which data will best answer the evaluation need and obtain permission to use this data. They may then need to work with the technical data steward of the admissions data to obtain appropriate access and continue working with the data owner or business data steward if there are questions about interpreting the data received.

**Tip 3: Evaluate generative AI tools for data privacy, risk and integrity**
The landscape of generative AI tools and their capabilities are changing rapidly. Once the decision is made to use generative AI in your work, the process to select an available tool should consider multiple factors including (University of Illinois Urbana-Champaign 2024):

- Enterprise licenses: If your institution has an enterprise license for the tool, then there are likely commercial data protections built-in both contractually and within the advanced features made available in the version that becomes available for your use.
- Compliance with privacy and usage policies: Generative AI is subject to any institutional, state and federal privacy laws such as FERPA and Health Insurance Portability and Accountability Act (HIPAA). If you want to use the tool with protected data, then ensure that the tool you select can provide appropriate levels of privacy and alignment with institutional and other privacy usage policies.
- Training data: The data that the generative AI tool was trained on may include data that violate privacy and copyright laws and/or was possibly collected in an unethical manner. Datasets may also introduce limitations or bias if its scope is not broad and comprehensive. If there is no initial transparency about how the training data was created for a generative AI tool, it is prudent to learn more about this before selecting and using it.
- Input data: The data that you input into the generative AI may in turn be used as training data. Further, user behaviors and analytics may be collected during your use of a generative AI tool and shared with third parties. Understanding how the data you submit as prompts to a generative AI tool will be used can help you determine the types of data that are appropriate to be used with the tool.
- Integrity of output: The output created by generative AI may not be accurate or true since there is a dependence on the strength and integrity of the models and training

data they are built upon. These inaccuracies are referred to as "hallucinations" (Sallam 2023). More significantly, generative AI models are designed to generate plausible content by predicting its next item, e.g., similar to the next word of an autocomplete, and not necessarily to verify the truth of what is produced (Mesko and Topol 2023). It is always necessary for human review and validation of any generated output before it is used as a source of truth elsewhere.

When using generative AI for accreditation and CQI, the same data privacy and usage considerations that were in place previously continue to apply. Additional considerations revolve around the integration of this new class of tools that may introduce risk and integrity concerns that may not seem obvious. For example, the conversational nature of generative AI chatbots may draw a deeper engagement with the tool when attempting to perform a data analysis than when using of a tool like Microsoft Excel to perform the same task. In other words, the process of prompt refinement may lead to the sharing of data that was not originally intended.

**Tip 4: Preparing your data to use with generative AI**

Taking the steps to prepare data is critical to having high quality, reliable and accurate data for CQI and analysis. Raw data is often plagued with errors, inconsistencies and missing values (Kwak and Kim 2017). Raw data can be further impacted by an imbalance in the dataset, e.g., overrepresentation of one class of data, that results in biased output from an analysis.

Data preparation has always had a significant role in data analysis and it continues to remain a necessary prerequisite for high quality data to be used with generative AI. The steps for data preparation should not be overlooked and include the concepts found in Table 1 (Alteryx 2024).

*Table 1. Data Preparation Concepts*

| Concept | Definition |
|---|---|
| **Cleaning** | Data cleansing, also known as data cleaning or scrubbing, identifies and fixes errors, duplicates, and irrelevant data from a raw dataset. (*Source: https://www.alteryx.com/glossary/data-cleansing*) |
| **Transforming** | Data transformation is the process of converting data from one format to another. The most common data transformations are converting raw data |

| | |
|---|---|
| | into a clean and usable form, converting data types, removing duplicate data, and enriching the data to benefit an organization. (*Source: https://www.alteryx.com/glossary/data-transformation*) |
| **Validating** | Data validation is the process of checking data that meets requirements by comparing it to a set of rules that have already been set up or defined. *(Source: https://www.questionpro.com/blog/data-validation)* |

Privacy and confidentiality can be further ensured by de-identifying data to keep it anonymized. There are a variety of techniques that can be used to de-identify data such as removing direct identifiers, e.g., name, email, removing or re-coding specific dates, e.g., birthdates, and removing institution-specific references (Johns Hopkins University 2024).

Another way to prepare data is to think about how to train the generative AI model you will use by providing the resources and context for your specific work. Plugins and application programming interfaces (APIs) for generative AI tools enable you to customize and refine the responses that are generated from your prompts. For example, if you want to use generative AI to help you with curriculum mapping, providing the AI information like the AAMC keywords list and associated definitions, the MedBiquitous standardized vocabularies for instructional methods, assessment methods and resource types (MedBiquitous 2024) and the educational program objectives for your medical school can help generate responses that support your institutional mapping goals.

**Tip 5: Understand the potential for bias embedded in AI tools and models**
Generative AI relies on large language models (LLMs) which are AI models trained on vast amounts of data to generate and process human language. Existing literature demonstrates evidence that LLMs can encode and extend the biases, e.g., racial and gender-related, that are embedded in its training data (Zack et al. 2024). These concerns of bias exist and may even be amplified when the training data is from our own institutions and influenced by the type, quantity and method of collecting data for use with AI.

Until comprehensive and transparent assessments are developed to help recognize the introduction of bias in generative AI tools, it is necessary for us to find ways to mitigate the impact of bias when using these tools in our own work. Some strategies include augmenting the training data, altering the training data and applying pre- and/or post-processing corrections to the data (Hastings 2024).

Though there is a significant amount of research focused on approaches to mitigating bias in AI, it is too early to cite specific examples of successful approaches. Moving forward with this in mind will mean reflecting on the specific nature of the AI task, seeking knowledge about the training data and LLM, and recognizing when the potential for bias outweighs the benefits of the AI tool. For example, when evaluating the outcomes of a faculty diversity hiring initiative, understanding how the data related to the desired diversity categories was captured, e.g., granularity of race/ethnicity categories, options for gender identification, may reveal that a standardized approach to analyzing the faculty hiring data is not possible without misrepresentation and bias in the results.

**Tip 6: Create effective prompts**

While the prior tips could broadly be applied to data analysis in general, this tip relates directly to the use of generative AI in accreditation and CQI work. Creating effective prompts is a new type of skill that should be obtained to effectively use generative AI tools. A prompt is the way in which you interact with generative AI to perform a task (Amazon Web Services 2024). The quality of the prompt determines the value of the output generated by the AI tool. Prompts can be as short as a single word or as complex as a series of well-detailed steps.

The ability to create effective prompts to achieve meaningful generated responses is an art and the term "prompt engineering" is used to describe the process of designing and refining prompts (Mesko 2023). "Prompt engineering is a relatively new field of research that refers to the practice of designing, refining, and implementing prompts or instructions that guide the output of large language models (LLMs) to help in various tasks." (Mesko 2023)

To be an effective user of generative AI, one must be a prompt engineer through practice, iteration and study of the nuances that context, clarity and completeness bring to a prompt (Mesko 2023). There are many emerging frameworks and guidelines to support prompt engineering. For example, Bertalan Mesko (2023) shares an "Introduction to Prompt Engineering in Healthcare" infographic with specific prompt recommendations such as describing a setting and providing a context, identifying the overall goal of a prompt first and asking the generative AI to take on a role with healthcare-based examples to improve generative AI interactions.

**Tip 7: Plan an effective and efficient data story**
Data storytelling transforms numbers into a persuasive narrative that drives action (Duarte 2019). Investing time in data story planning at the start of any data analysis project provides the analyst a roadmap that keeps them connected to the story through each step of the analytics process. No matter the framework used to plan a data story or whether the plan is developed inductively or deductively, the key questions supporting analysis objectives are generally identified based on the analyst's intuition and domain expertise and they are revisited during the analytics process as needed. Generative AI can be used for ideation to improve the quality of questions identified by the analyst. In simple terms, generative AI can expand and shape approaches, increase their knowledge in a topic area, and ultimately even be considered a sounding board in the planning step of the analytics process. For example, Element 8.7 on Comparability of Education and Assessment from the LCME Data Collection Instrument (DCI) (LCME 2024) can lead to collaboration between curriculum and accreditation teams to define comparability components and identify relevant assessment outcomes. These efforts can be organized considerably through data story planning by introducing the situation and how data measures align with core accreditation requirements detailed in the DCI element, exploring complications that may be contributing to the positive or negative direction of data measures, and presenting actionable recommendations that address the complications through analytical sub actions. Each of these components of the data story can be boosted through ideation with generative AI tools that expand how units approach this DCI element and navigate the analytics process.

**Tip 8**: **Boost statistical analysis in exploratory data analysis**
Fields like accreditation and assessment/evaluation bring together teams from different backgrounds with varying skill sets and levels of proficiency in data analysis and data science (Fayyad and Hamutcu 2020). The significance of statistics in data analysis and the appreciation for evidence-based results in academic environments suggests consideration of generative AI as a guided resource for medical educators who have higher proficiency in functional areas other than statistics. Generative AI's ability to build upon a discussion with the analyst, compared to a typical web-based search, makes it a valuable tool to effectively and efficiently choose the appropriate statistical tests to conduct a desired evaluation. As an example, student performance on internally developed or national exams are rich data sets that can play a significant role in identifying students who are at risk of academic difficulty as called for in LCME DCI Element 11.1 on Academic Advising and Academic Counseling (LCME 2024). Fully exploring the data collected through individual student

assessments can mean performing various statistical analyses to identify relationships between exam components, pre-matriculation data, and other measures that may predict student performance. Generative AI tools can be a starting point for selecting suitable methods to determine appropriate statistical tests based on primary data set characteristics.

**Tip 9**: **Consider generative AI tools for programming needs**

Generative AI can support various programming activities, from coding and debugging to algorithm design, understanding programming concepts, and guidance on best practices. In this application, analysts who need to perform complex analyses can utilize code snippets and explanations that generative AI tools provide and help to brainstorm solutions. Inspired by pair programming, a software development method in which two programmers collaborate on a single block of code to improve accuracy and efficiency (Codeacademy 2021), analysts can pair themselves with generative AI tools for suggestions on code completion or responses to specific questions. An example of this application can be described by taking a closer look at LCME DCI Element 1.1 on Strategic Planning and Continuous Quality Improvement (LCME 2024) that highlights how CQI efforts are directly related to review cycles, data sources, and data stakeholders. To support CQI efforts in a data-driven manner, analysts may be just as involved with data sensemaking and storytelling as they are with organizing data collected through medical school management systems. Using generative AI tools can help boost skills in programming and scripting languages even for those who may not have this prior experience and better organize data at scale no matter their technical background or the organization's data management approach.

**Tip 10**: **Conduct visual exploratory data analysis and create client-ready data visualizations**

Accreditation and CQI relies on many types of reports, tables and graphs to perform the monitoring required in determining compliance with quality standards. These include data-driven charts used in exploratory data analysis that visually analyze and make sense of summary statistics and relationships describing the data as well as client-ready visualizations that are intentionally crafted based on principles including prioritizing the information you want to share, establishing meaning through color and selecting the right geometries to represent data (Midway 2020). The data analysis features of popular generative AI tools can be used to create charts that explore the data or charts to refine as needed for inclusion in presentations to stakeholders. For example, the AAMC Medical Student Graduation Questionnaire (GQ) (AAMC 2024) which is frequently referenced throughout the LCME DCI (LCME 2024), can be explored through visual analysis and data visualizations powered by generative AI tools for program evaluation and CQI monitoring. This exploration can efficiently produce a first draft of line graph (Figure 1), heat map (Figure 2), or deviation graph visualizations, which are often used to display performance trends at the school and national level and further generate narrative summaries corresponding to these visualizations.

*Figure 1. Line graph produced using GPT-4o to visually analyze questions from the 2021 GQ All Schools Summary report for exploratory data analysis*

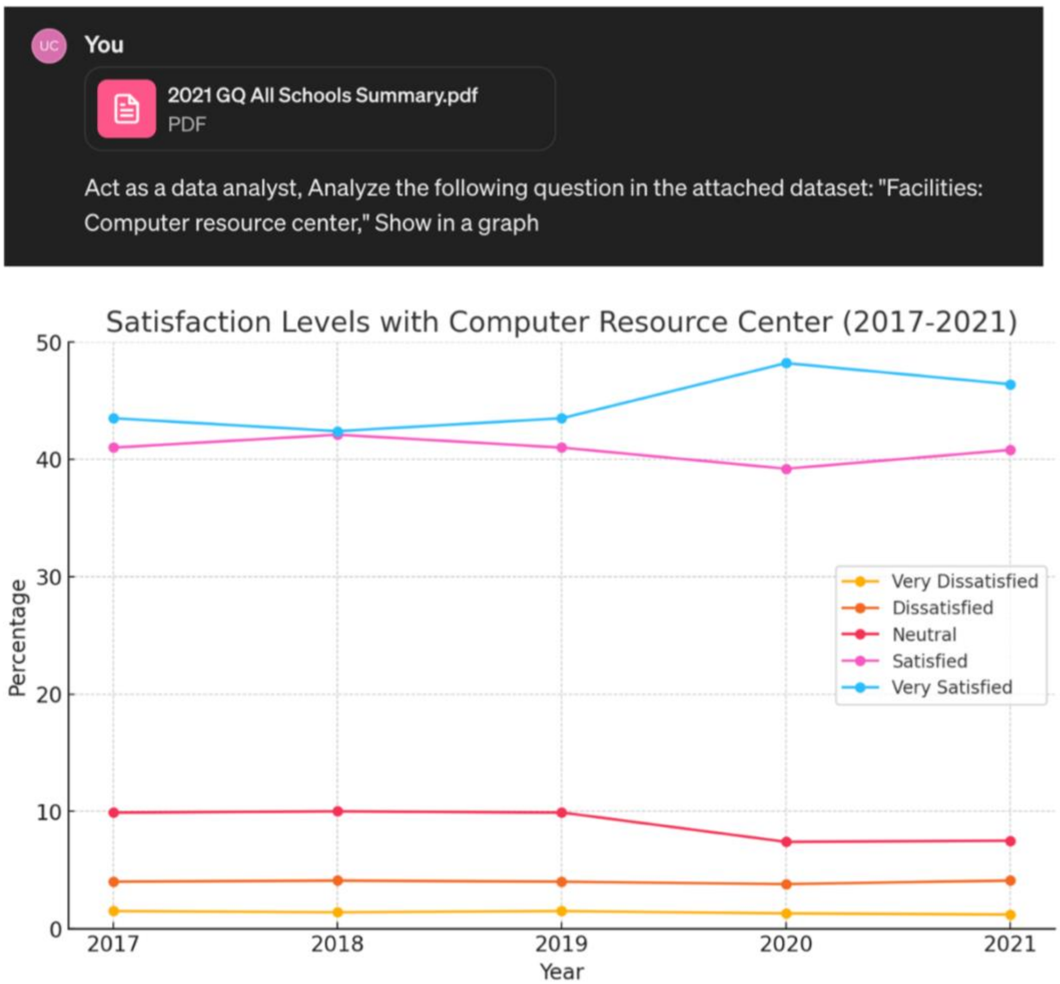

*Figure 2. Heatmap produced using GPT-4o to visually analyze questions from the 2021 GQ All Schools Summary report for exploratory data analysis*

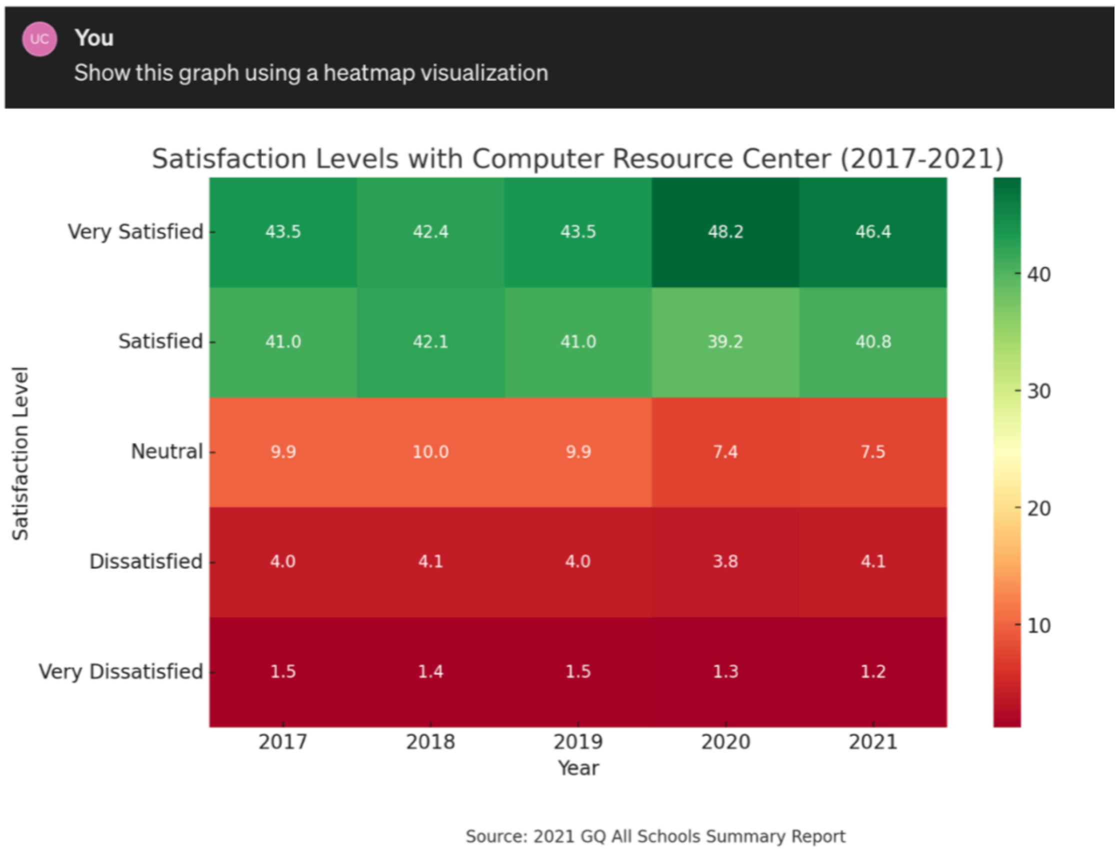


There are also different categories of AI tools that you can consider for your work beyond a generative AI chat bot. Productivity tools such as those for data visualization, slide presentation creation and project management are integrating generative AI capabilities into their feature updates. For example, Microsoft Powerpoint offers a built-in “Designer” feature that generates slide suggestions based on the content found on the slide and graphics tools, e.g., Venngage (Venngage, Inc.) and Canva (Canva), integrate AI features to produce customized visual content based on user prompts.

**Tip 11: Build trust in generative AI tools and the generated results**

As AI technologies, including generative AI, improve and more decisions and processes can be automated, algorithm aversion and lack of trust in these tools may arise more frequently. Awareness of valid concerns with the ethics of technology companies that rely on data collected from users as a core component of their business models is also important (Walsh 2019). Medical educators can navigate algorithm aversion by recognizing experiences with hallucinations produced by generative AI tools (Sallam 2023) and intentionally establishing safety controls and operating boundaries in the analytics process. These controls can be designed as simply as regularly checking in with domain experts as part of the analytics process and turning off chat history when using external generative AI tools, or in a more sophisticated manner to reproduce two-layer verification models. For example, if data analysis features of generative AI tools are used to produce reports for LCME DCI Element 8.8 on Monitoring Student Time (LCME 2024), analysts can review these reports with team members who are the subject matter experts with the policies and procedures related to duty hours, rotations, and the curriculum schedules to verify results.

**Tip 12: Consider ethics and transparency when using generative AI tools**
Including adequate disclosures and appropriate citations when using generative AI tools in data analysis projects, as you would with any other resource, reflects the integrity of project owners. Guidelines on how to cite and acknowledge content produced by generative AI tools are still in development and the acceptability of generative AI taking the role of author vs. tool for research and scholarship remains questionable (Sallam 2023). However, these disclosures are necessary to support project transparency and broaden stakeholder understanding of the use cases for generative AI tools. Other ethical considerations of generative AI tools may include the human, climate and environmental impacts related to sources powering AI research and development (Keller, Dononghoe, Perry 2024). While navigating global issues requires collective action and decision making, an example of this application at an individual level is appropriately acknowledging any content, data-related or otherwise, that is created using generative AI tools as recommendations around their use develop from institutions, accrediting bodies and the medical education community.

**Conclusion**
The desire to find use for generative AI is reaching ubiquity across many fields. This article explores the use of generative AI applications for accreditation and CQI activities in medical education. We organized our tips to address two phases of generative AI use. The first phase focuses on activities that are recommended prior to the use of generative AI including preparation of the data and data protection, privacy and ethical considerations. These activities are general recommendations that should be applied to any projects involving data and continue to remain important when engaging generative AI capabilities. We bring these to the audience of the broader medical education community as the interest

in this technology expands beyond those who had roles responsible for data management and analysis previously.

The second phase focuses on relevant applications of generative AI to support data management and analysis. The novel capabilities that generative AI offers expand the use cases for both experienced analysts and those experimenting on their own. While we acknowledge limitations of the tips for this second phase due to the rapidly changing generative AI tools and the expertise surrounding its use, we find implications for sharing these recommendations early to encourage the broader adoption of generative AI in medical education through a deliberate and informed perspective.

**References**

About Educause [Internet]. Educause. [Accessed on March 30, 2024]. Available from: https://www.educause.edu/about

Burns S, Robert J, Muscanell N. EDUCAUSE QuickPoll results: Growing needs and opportunities for security awareness training [Internet]. Educause. October 30, 2023. [Accessed on March 30, 2024]. Available from: https://er.educause.edu/articles/2023/10/educause-quickpoll-results-growing-needs-and-opportunities-for-security-awareness-training

DAMA International. DAMA-DMBOK: Data management body of knowledge. Vancouver, WA: Technics Publications, LLC; 2017.

Duarte N. DataStory: Explain data and inspire action through story. United States of America: Ideapress Publishing; 2019.

Glossary term: Data preparation [Internet]. Alteryx. [Accessed on March 30, 2024]. Available from: https://www.alteryx.com/glossary/data-preparation

Educational Methodology [Internet]. MedBiquitous. [Accessed on March 30, 2024]. Available from: https://www.medbiq.org/initiatives/educational-methodology

Edwards B. ChatGPT is one year old: Here's how it changed the tech world [Internet]. Ars Technica. November 30, 2023. [Accessed on March 30, 2024]. Available from: https://arstechnica.com/information-technology/2023/11/chatgpt-was-the-spark-that-lit-the-fire-under-generative-ai-one-year-ago-today/

Fayyad U, Hamutcu H. Toward foundations for data science and analytics: A knowledge framework for professional standards. *Harvard Data Science Review, 2(2).* doi:10.1162/99608f92.1a99e67a.

Graduation Questionnaire (GQ) [Internet]. Association of American Medical Colleges (AAMC). [Accessed on March 30, 2024]. Available from: https://www.aamc.org/data-reports/students-residents/report/graduation-questionnaire-gq

Keller J, Dononghoe M, Perry A. The US must balance climate justice challenges in the era of artificial intelligence [Internet]. The Brookings Institution. [Accessed on March 30, 2024].

Available from: https://www.brookings.edu/articles/the-us-must-balance-climate-justice-challenges-in-the-era-of-artificial-intelligence/

Kwak SK, Kim JH. Statistical data preparation: management of missing values and outliers. *Korean J Anesthesiol*. 2017;70(4):407-411. doi:10.4097/kjae.2017.70.4.407

Meskó B. Prompt engineering as an important emerging skill for medical professionals: Tutorial. *J Med Internet Res*. 2023;25:e50638. doi:10.2196/50638

Meskó B, Topol EJ. The imperative for regulatory oversight of large language models (or generative AI) in healthcare. *NPJ Digit Med*. 2023;6(1):120. doi:10.1038/s41746-023-00873-0

Midway SR. Principles of Effective Data Visualization. *Patterns (N Y)*. 2020;1(9):100141. doi:10.1016/j.patter.2020.100141

Sallam M. ChatGPT utility in healthcare education, research, and practice: Systematic review on the promising perspectives and valid concerns. *Healthcare (Basel)*. 2023;11(6):887. doi:10.3390/healthcare11060887

Standards, publications, & notification forms [Internet]. Liaison Committee on Medical Education (LCME). [Accessed on March 30, 2024]. Available from: https://lcme.org/publications/

Gemini [Internet]. Google. [Accessed on March 30, 2024]. Available from: https://gemini.google.com/

Hastings J. Preventing harm from non-conscious bias in medical generative AI. *Lancet Digit Health*. 2024;6(1):e2-e3. doi:10.1016/S2589-7500(23)00246-7

Copilot [Internet]. Microsoft. [Accessed on March 30, 2024]. Available from: https://copilot.microsoft.com/

Privacy considerations for generative AI [Internet]. University of Illinois Urbana-Champaign. [Accessed on March 30, 2024]. Available from: https://cybersecurity.illinois.edu/policies-governance/privacy-considerations-for-generative-ai/

Protecting human subject identifiers [Internet]. Johns Hopkins University. [Accessed on March 30, 2024]. Available from: https://guides.library.jhu.edu/protecting_identifiers/de-id_steps

Walsh M. The algorithmic leader: how to be smart when machines are smarter than you. Canada: Page Two Books; 2019.

What is pair programming? [Blog]. Codeacademy. September 24, 2021. Accessed on March 30, 2024]. Available from: https://www.codecademy.com/resources/blog/what-is-pair-programming

Zack T, Lehman E, Suzgun M, et al. Assessing the potential of GPT-4 to perpetuate racial and gender biases in health care: a model evaluation study. *Lancet Digit Health*. 2024;6(1):e12-e22. doi:10.1016/S2589-7500(23)00225-X

Zewe A. Explained: Generative AI: How do powerful generative AI systems like ChatGPT work, and what makes them different from other types of artificial intelligence? [Internet]. MIT News. November 9, 2023. [Accessed on March 30, 2024]. Available from: https://news.mit.edu/2023/explained-generative-ai-1109