

Generative Artificial Intelligence Use at Randolph College

Scope of This Policy and Definitions of Terms

Definition of Generative Artificial Intelligence

Generative Artificial Intelligence (GAI) refers to a type of technology that uses algorithms and large amounts of data to create new content such as text, images, music, and code based on patterns it has learned. Unlike traditional software, which follows very specific, predefined rules, GAI can produce original output that resembles human-created content.

GAI Applications

GAI is at work in many applications outside familiar AI assistants like ChatGPT, Microsoft Copilot, and Google Gemini. GAI can be found in productivity suites (Microsoft Office, Google Workspace), design tools (Adobe products, Canva), and even CRMs and social media platforms. This is often, but not necessarily, made clear to the user. When in doubt, it is best to assume that GAI has been incorporated into an application.

Randolph College has not yet endorsed any specific GAI products as being compliant with sensitive data handling requirements or legal regulations. If and when such an endorsement comes, this policy will be amended.

Audience

This policy applies to the internal work of members of the Randolph College community — staff, faculty, and students, referred to here as the “user” — outside the academic setting. Note: A policy on the use of GAI in academics is under development. When complete, a link to that policy will be inserted here.

Departmental Adaptation

This policy applies to all units, departments, and offices of Randolph College. However, individual units, departments, or offices may establish additional restrictions on the use of GAI, based on the nature of the work performed, the sensitivity of the data handled, or other operational concerns.

Users should work with their direct supervisors to understand the specific GAI guidelines relevant to their roles, and supervisors should be ready to clarify whether additional restrictions or special considerations apply to the use of GAI in their areas.

Basic Expectations in the Use of GAI

Applicability of Existing Controls

GAI should never be used to create content that violates:

- Federal, state, or local laws or regulations;
- Regional, national, or programmatic accreditation standards;
- Guidelines or standards published by professional organizations applicable to Randolph College; or
- Randolph College institutional policies.

Workplace Expectations Regarding GAI Proficiency

The use of GAI is supported by Randolph College as part of the College’s ongoing effort to enhance efficiency and innovation. Although users are encouraged to integrate GAI tools responsibly into their job functions, there is no institutional expectation for GAI proficiency. However, supervisors may expect GAI proficiency in certain roles and even require it of current or future staff.

Transparency

In general, the use of GAI should be disclosed when the nature of the product is such that the reader might expect to be told, or might be concerned, that GAI was used to create it. Therefore:

- It is *not* expected that the user disclose GAI use for:
 - Routine, incidental tasks (including, but not limited to, unofficial emails, informal research, background reading, and repetitive data processing);
 - Minor, repetitive tasks for which GAI is used purely to save time; or
 - Grammatical, syntactical, or formatting assistance.
- It *is* expected that the user disclose the GAI when producing content that:
 - Will be germane to planning or decision-making;
 - Has ethical or legal implications;
 - Contains potentially sensitive or confidential information (see “Data Security” below);
 - Comprises a very large proportion of the final product; or
 - Is conspicuously represented as being the user’s own work.

Whether or not GAI use is disclosed, the user remains responsible for the issues described in the section “Working Responsibly with GAI Outputs” below. When in doubt, the user should err on the side of disclosure, and, if deemed necessary, GAI use should be disclosed to the degree of detail appropriate to the situation.

Sample Disclosure Statement

Disclaimer: This [document, report] includes significant [content, language] generated using Generative Artificial Intelligence (GAI). The author has reviewed and verified the accuracy of all information contained herein and is solely responsible for its content.

Questions regarding GAI disclosure should be directed to the user’s supervisor, who has the authority to adjust transparency expectations to ensure that GAI use is appropriate for the specific work being completed.

GAI and Data Security

Note: A comprehensive data security policy is currently under development by the IT Policy Committee. When that policy is completed, much of the material in this section will be superseded and therefore removed.

The Nature of an LLM

GAI is a “large language model” (LLM) that is “trained” using preexisting data. Every file, data set, or prompt entered or uploaded by the user potentially adds to LLM data set and therefore is no longer private.

Personally Identifiable Information

Data security policy depends on an understanding of Personally Identifiable Information (PII). PII refers to any data point, or combination of data points, that can be linked to a specific person and uniquely identify that person.

This identification can be direct or indirect. Direct identifiers include name, photograph, email address, phone number, Social Security number, student ID number, and other data elements. Indirect identifiers include date of birth, place of birth, address, ZIP code, IP address, demographic details, employment details, matriculation and/or graduation details, the names of family members, mother’s maiden name, and other data elements. Direct identifiers identify an individual independently, but a combination of indirect identifiers might identify an individual as well.

Information related to Randolph College employees, employment history, and employment performance should also be considered as PII. This includes financial records, performance reviews, medical details, and salary and benefit information, in addition to other standard PII elements.

PII is to be treated as confidential in a general sense, but of specific concern is adherence to the following: the Family Educational Rights and Privacy Act (FERPA); the Health Insurance Portability and Accountability Act (HIPAA); and the Gramm-Leach-Bliley Act (GLBA).

PII and SaaS Applications

Randolph College provides, and/or provides access to, various software applications that have been reviewed and approved by the Office of Information Technology (IT) for use in processing certain PII. Also in use by College staff, faculty, and students are various applications accessed directly on the Software as a Service (SaaS) model.

This is to be expected, but it is important to understand that not all SaaS applications have been sanctioned by IT. Employees should not use unsanctioned SaaS applications to process, store, or transmit PII when using GAI. (Here, “unsanctioned SaaS applications” denotes any cloud-based software that runs without the need for local installation and that has not been reviewed or approved by IT.) All Randolph College employees are responsible for exercising good judgment when selecting and using SaaS applications. The user should contact IT for assistance if unsure about an SaaS application.

Internal College Information

In addition to PII about individual people, the user should be vigilant regarding the use of Randolph College institutional information. The user should not input, upload, or otherwise disclose any non-public College information, unless explicitly authorized to do so by the supervisor. This includes, but is not limited to:

- Budgetary or financial data;
- System configurations or IT infrastructure details;
- Internal planning documents;
- Donor information;
- Any confidential internal communications

Questions regarding this matter, and with requests for exceptions, can be made in writing (email acceptable) to the appropriate member of the President’s Leadership Team.

Working Responsibly with GAI Outputs

Limitations of GAI

There are several inherent limitations of GAI that should always be borne in mind by the user:

- GAI-generated material can be inaccurate, misleading, and even fabricated. (The latter situations are sometimes called “hallucinations.”)
- In so far as GAI is a “large language model” (LLM) that has been trained using preexisting, human-created material, its product may contain copyrighted material, and its creative output may incorporate the creative work of someone else.
- As a mechanism that is non-human by definition, GAI lacks an ethical dimension. It is incapable of ensuring that its product, or the use of its product, embodies a respect for humanity, is fair and unbiased, is non-discriminatory, honors personal privacy, is not misleading, and/or embodies an inclusive viewpoint. GAI must be supported by the user’s own accountability in these areas.

Accuracy

As stated above, GAI outputs may contain errors, inaccuracies, or incomplete information that require human oversight and correction. Users are responsible for ensuring that all content generated by GAI systems is accurate.

Copy-Editing

The user should thoroughly review and manually edit all GAI-generated content before using or sharing it.

Plagiarism and Copyright

Typically, GAI does not give credit to, or acquire permission from, the sources on which it has been trained and which it incorporates into its output. Although its output is often described as “new” or “original,” GAI may create content very similar, or practically identical, to unnamed original works. This is true of any written, visual, musical, or other human creation.

Nevertheless, incorporating GAI output that itself is “using or closely imitating the language and thoughts of [an] author [or artist] without authorization, and the representation of that author’s work as [its] own ... by not crediting the original author” does meet the definition of plagiarism and, if the material is copyrighted, constitutes copyright infringement (“Plagiarism”).

The user should therefore review all GAI output for the presence of citable sources and copyrighted intellectual property. It can be difficult to identify the sources whose words, artistic creations, ideas, and data are part of a GAI output, so the thoroughness of such a review may depend on how the output will be used: on the formality of the product, on the ultimate audience for the product, and on whether that audience is internal or external to the College. The user should consult closely with the supervisor on these things, and, if the presence of a source is detected that requires attribution under the circumstances, a citation should be inserted in the appropriate style.

Additionally, inputting or uploading copyrighted materials to an LLM, even as a prompt, without permission of the author or creator may also constitute copyright violation.

Disclaimer: Research for this policy was conducted, and ideas for content were generated, using ChatGPT. The committee that created the document (see cover sheet for a list of committee members) decided on all final content, most of which was generated independently, and the final text was written by the committee. The committee is solely responsible for its content.

Work Cited

“Plagiarism.” Dictionary.com, Dictionary.com, 2024, www.dictionay.com/browse/plagiarism.