Policy on the Use of AI Text Generation

Preamble

This policy derives from a case study conducted on Tuesday January 24, 2023, by the 47 juniors and seniors in **Boston University Computing & Data Sciences** DS380, "Data, Ethics, and Society." The focus of the case study was the use of OpenAI's ChatGPT in university education. The case study involved analyzing ideas such as skill building and cheating, as well as the interests of stakeholders. The challenge posed within the case study was to devise a policy for use in their own class that takes full account of the aims of university education and the perspectives of relevant stakeholders, and can be shared for consideration in other classes. The case study was revisited each class session thereafter until a consensus was achieved. The result was the **Generative AI Assistance (GAIA) policy**.

The fundamental concerns of class members included wanting to learn how to use AI text generators skillfully while enhancing rather than damaging their own developing abilities as writers and thinkers, to honor stakeholder interests, and to ensure that grading is fair between people who do not use AI text generators and people who do use them, in a variety of different ways. This introduces new demands on those charged with grading student work. Thus, the GAIA policy stresses transparency, fairness, and obligations for both students and teachers.

The Generative AI Assistance (GAIA) Policy

We welcome AI language-generation tools (collectively, large language models, or LLMs) into the learning process, in a way that preserves fairness, optimizes student skill building, and honors relevant stakeholder perspectives. These include our own perspectives as students eager to learn and build careers, and the perspectives of our families who send us to university, our professors who are charged with teaching us vital skills, our university that has a responsibility to attest to our competency with diplomas, our future employers who invest in us because of our abilities and character, and our colleagues who lack privileged access to valuable resources. To that end, the GAIA policy adopts a few commonsense limitations on an otherwise embracing approach to LLMs.

1. As students, we should give credit to LLMs whenever they are used, even if only to generate ideas rather than usable text.
2. When writing papers or take-home assignments using LLMs, we should include in an appendix our entire exchange(s) with the LLM, highlighting the most relevant sections. We should write a narrative at the beginning of the appendix explaining precisely *how* we used the LLM (to generate ideas, turns of phrase, elements of text, long stretches of text, lines of argument, pieces of evidence, etc.). This narrative should also explain *why* we used the LLM (to save time, to surmount writer's block, to stimulate our thinking, to handle mounting stress, to clarify prose, to experiment for fun, etc.).

3. We should not use LLMs to help with in-class examinations, tests, or assignments, unless they are explicitly organized around an LLM (e.g. an in-class assignment might challenge students to elicit from a LLM the most profound statement on the future of AI).
4. If we choose to use an LLM, we should use it wisely and intelligently, aiming to deepen our grasp of the subject matter and to support our learning.
5. Our professor should seek to understand how LLMs work, including their strengths and weaknesses, as well as tools for detecting LLM-generated output, in order both to optimize their value for our learning, and to incorporate that awareness into grading procedures.
6. Our professor should:
   - Treat work submitted by students who declare no use of LLMs as the baseline for grading.
   - Use a lower baseline for students who declare use of LLMs, depending on how extensive the usage, while rewarding creativity, critical nuance, and the correction of factual inaccuracies or superficial interpretations in response to suggestions made by the LLM.
   - Employ AI-text-generated-detection tools to evaluate the degree to which machine-generated text is likely to be present in student work.
   - Impose a significant penalty for low-energy or unreflective reuse of wording generated by LLMs, to the point of assigning zero points for merely reproducing LLM output.

Furthermore:

- We acknowledge that some may prefer stronger restrictions on the use of LLMs, and therefore that the GAIA policy may make them feel uncomfortable.
- We acknowledge that some aspects of this policy are non-enforceable and depend on goodwill, a sense of fairness, and honorable character.
- We acknowledge that this policy will need to be revisited when subscription versions of LLMs become available, which may be affordable for some students but not for others.
- We acknowledge that this policy may need to be revised in light of other policies that would supersede it or due to novel technological developments.

Authors

Students: Ghalia Alshanbari, Adriana Alvarado, Lamya Alzahrani, Matthew Batacan, Olivia Bene, Gauri Bhandarwar, Jack Campbell, Timothy Chang, Natalia Clark, Jesse Conde, Osama Dabbousi, Matteo David, Aaron Fox, Jingyi Gao, Zachary Gou, Kanishka Gupta, Mengying He, Iris Ho, Clarissa Hoyt, Julia Hu, Junyi Huang, Dmitriy Kazlouski, Rebecca Lawrence, Yagev Levi, Victoria Luo, Weining Mai, Ella McNally, Misha Nasir, Yusra Nasir, Vincent Nguyen, Jeremy Ortiz, Rafael Perron, Oliver Samuels, Maria Shevchuk, Pratham Shroff, Jonathan Simen, Daniel Skahill, Mustafa Taybah, Sarai Thach, George Trammell, Andrew Tuckman, Carter Vande Moore, Ziyang Wang, Ryan Wong, Zining Ye, Sing Yu Yeung, Wilson Zhang

Teaching Assistants: Tejovan Parker, Zhengyang Shan. **Professor:** Wesley J. Wildman