



JOHNS HOPKINS
WHITING SCHOOL
of ENGINEERING

AI Environmental Impacts / Course Projects

9/21/26

Reflection until 9:15am

- Energy Consumption of AI until **9:15**:
 - Luccioni, Sasha, Yacine Jernite, and Emma Strubell. "Power hungry processing: Watts driving the cost of AI deployment?." Proceedings of the 2024 ACM conference on fairness, accountability, and transparency. 2024.
 - Andy Masley. "Using ChatGPT is not bad for the environment - a cheat sheet", Blog post, 2025.

Notes on reflections:

- No need to summarize – we have done the readings too!
- Write ~1 paragraph
- Make clear to us that you have done the readings (and thought about them)
 - E.g. mention a detail or two from each reading

Class Time

- 9-9:15am: Reflection on Data Privacy
- 9:15-9:55am: Discussion on Energy Consumption
- 9:55-10:05: Discussion recap
- 10:05-10:25: Break / Reflection on Data Centers
- 10:25-11:00: Discussion on Data Centers
- 11:00-11:10: Discussion recap
- 11:10-11:30: Introduction to course projects

Course Project

- <http://ai-ethics-601-770.cs.jhu.edu/fa2023/project.html>
- 1. Literature Review [Due October 9]
- 2. Project Proposal (4-8 pages)
 - Written Report [Due October 30]
 - Short in-class presentation
- 3. Final Report [Due at the end of the semester]
 - Individual Q&A after you hand in your report

Literature Review

- Short paper (min 2 pages, max 6 pages) summarizing and synthesizing several papers (or works in general)
 - Should focus on a specific topic, not all of AI Ethics
- Groups of one should review at least 5 papers, groups of two should review 7 papers, and groups of three should review 9 (you are welcome to review more)
- Most people will choose the same topic for your lit review and final paper, but this is not required, and the lit review will be graded on its own
- More tips on the course website: <https://ai-ethics-601-770.cs.jhu.edu/fa2026/project.html>
- Please use the ACM template, with “nonacm” and “sigconf parameters (2-column format) <https://www.acm.org/publications/proceedings-template>. No need to include the teaser figure, CSS concepts, or keywords

Project Topics

- Can be anything related to the course
- We encourage you to talk to instructors about ideas (via email, Piazza, office hours, schedule additional appointment)
- Literature review focuses on identifying broad topic area – you have more time for specific project ideas!

Project Examples

- Conduct a technical assessment of a model or deployed system, such as probing for energy costs or unfairness
 - Luccioni, Sasha, Yacine Jernite, and Emma Strubell. "Power hungry processing: Watts driving the cost of AI deployment?." Proceedings of the 2024 ACM conference on fairness, accountability, and transparency. 2024.
 - Van Koevering, Katherine, Anjalie Field. "It's How You Ask: Gender-Associated Linguistic Bias in LLMs". COLM (2026)
 - Buolamwini, Joy, and Timnit Gebru. "Gender shades: Intersectional accuracy disparities in commercial gender classification." FAccT. PMLR, 2018

Project Examples

- Write a critical survey of a topic in AI
 - Language (Technology) is Power: A Critical Survey of “Bias” in NLP. Su Lin Blodgett, Solon Barocas, Hal Daumé III, and Hanna Wallach. Annual Meeting of the Association for Computational Linguistics (ACL)
 - Anjalie Field, Su Lin Blodgett, Zeerak Waseem, and Yulia Tsvetkov. "A Survey of Race, Racism, and Anti-Racism in NLP" ACL (2021).

Project Examples

- Conduct a technical study, focused on building models to address ethical concerns.
 - Methods for reducing toxicity in model outputs
 - Methods for reducing stereotyping in model outputs
 - Methods for privacy preservation and anonymization

Project Examples

- Write a survey structured around a policy or regulation question (what should policy makers understand about AI in writing regulations?)
 - Some starting points:
 - AI Bill of Rights: <https://www.whitehouse.gov/ostp/ai-bill-of-rights/>
 - <https://www.vox.com/future-perfect/2023/7/3/23779794/artificial-intelligence-regulation-ai-risk-congress-sam-altman-chatgpt-openai>
 - <https://techpolicy.press/>

Project Examples

- Conduct a manual or automated meta-analysis of research processes in papers or other avenues where research discussions take place (social media, news)
 - Birhane, Abeba, et al. "The values encoded in machine learning research." *Proceedings of the 2022 ACM Conference on Fairness, Accountability, and Transparency*. 2022
 - Leah Hope Ajmani, Stevie Chancellor, Bijal Mehta, Casey Fiesler, Michael Zimmer, and Munmun De Choudhury. *A Systematic Review of Ethics Disclosures in Predictive Mental Health Research*. In Proceedings of the 2023 ACM Conference on Fairness, Accountability, and Transparency. 2023
 - Media hype around AI
 - Prevalence of corporate research at different ML venues
 - Policy discussions around AI

Ideas for data sources

- Generally, data that has been used in prior research
 - Identify from reading research papers, which may have directly released the data or cited how they collected it (for example, [this paper](#) has a pointer to [this archive](#) of US Congressional Records). Authors are also often willing to share data if you reach out to them, even if they didn't post it publicly.
 - Shared tasks associated with workshops. For example, past versions of the workshop on [NLP for Internet Freedom](#) has had shared tasks on misinformation and censorship
 - Benchmark/shared task data is useful as analysis data, not just as model training data, for example [this paper](#) and [this paper](#) criticize standard NLP benchmarks
- There are some tools for data collection
 - [Semantic Scholar API](#) provides data on research publications (more details in [the paper](#)), which could be used for a variety of meta-analyses, like [this paper](#) on Big Tech influence in research
 - Many websites are possible to scrape (if you pay attention to terms of service and rate limits)
 - E.g. Wikipedia is a great source of data
- There are some pre-collected archives of data
 - Common Crawl <https://commoncrawl.org/> - large archive of web data
 - Many datasets on huggingface

Next class: Fairness, Bias, and Stereotypes

- Part 1: Overview of major research on this topic from the last decade, with a focus on classification
 - No required readings
- Part 2: Fairness in Generative AI systems
 - [Bianchi, Federico, et al. "Easily accessible text-to-image generation amplifies demographic stereotypes at large scale." FAccT 2023.](#)
 - [Katherine Van Koeveering, Anjalie Field. "It's How You Ask: Gender-Associated Linguistic Bias in LLMs". COLM \(2026\)](#)
 - **Group 3 needs a discussion leader**