Final Project

The goal of the final project is to give you a more in-depth look at a topic of your choice within Human-AI interaction. This can take the form of a blog-style essay, an interactive system, a user study, or an assignment on image generation with GANs. The options are described below.

Regardless of which project option you select, there are three main parts:

1. **Project Check-in:** This functions as a Project Proposal. See the Check-in Assignment on Glow for specific details.
2. **Project Pitch:** A short ~2 minute overview/motivation/pitch of your project idea and any progress, to be presented in the last Conference Section.
3. **Final Project Submission:** Submit all your hard work!

Option A: Persuasive HAII Article

Write an article on a Human-AI topic of your choice. This could be a deeper look at a topic we covered in class, or approach an entirely new topic. This article should be in the style of an article or blog post that a public citizen would read on the web.

Format

Depth & Length To give you a sense of depth and length your article should go into, here are two examples: one from The Atlantic and another from Buzzfeed on the same topic of fake news analysis. Note with web format, spacing, images, and larger font sizes give things more space. Roughly, though, these articles are between 1,600 - 3,000 words, which should be the general range your article falls in.

Submission Format You are free to post your article on any platform, like Medium or your own website. However, please submit your article as a PDF *not a link* to Glow, even if that just means printing your online article to a pdf. The reason is for giving feedback to you in PDF, and not running into web issues like Medium paywalls during grading.

Citations You *must must must* provide proper credit for any research, articles, or images you reference in your article. Since this is for educational purposes, you are free to use any images you find on the web… but put a nice little caption below the image to say who the image belongs to. See the Atlantic and Buzzfeed examples for some nice inline examples of citations, which can be as simple as web links inline. If you cite a research article, put a formal citation (any style like APA, Chicago, etc) at the bottom of your article as well as an inline citation like this: [1] or like this: (Linneaus 1758). Normal Williams College Honor Code applies.

Content

The tone and focus of your article can be any argument that you want the range of Human-AI Interaction topics. It can address a more general public or more technical audience. However,

although a lot of your article may be discussion and synthesis of these topics, to give your work more depth, **you must cite at least 4 peer-reviewed academic papers on your topic**. Many AI topics have a ton of super active research work going on right now, so this shouldn't be hard to find some good foundational or summary papers on your topic to get you started. I highly recommend Google Scholar and using the advanced search or "source:" search operators for specific venues, such as: FaccT, AIES, IUI, HCOMP, CHI, UIST, MD4SG, ICML, and AAAI, among others. (Note that papers on Arxiv that have not been reviewed at a conference or journal don't count. Ask Iris if you aren't sure.)

Meanwhile, cite however many other kinds of sources, like news articles or whatever, that you feel help build your story.

- You may decide to write a speculative fiction about how Human-AI interaction will look in the future. If you choose this avenue, the same requirements apply. (And yes, the fiction must have an argument.) For instance, you may include an "author's note" at the end of your fiction that connects ideas you've imagined to current events or technologies today, including the citations.

What is this argument thing we keep talking about? An argument is your position on a question that has multiple possible answers. Here are some example arguments:

1. cameras in public spaces must, by law, be prohibited from using face recognition technology.
2. In order to make true advances in human-facing AI applications, it is essential to test with users in emerging economies (e.g. China, India, Ethiopia). Users from developed countries (e.g. the US, Iceland) bring in too many expectations from their prior experience with technology
3. CAVs offer unique opportunities to make deep-learning understandable, that other explanation and visualization techniques do not.

Effort of your essays:

1. 20% is argument clear? Tip: Have a sentence in the first paragraph or two that says: "I argue that…"
2. 15% How have others argued in this situation?
3. 15% Why are they wrong? Cite sources (if possible) in support of them being wrong.
4. 20% What evidence do you have in support of argument?
5. 20% Is there any evidence that suggests you may be wrong? How do you respond to this evidence?
6. 10% is your article easy to read? Do you use appropriate graphics/visualizations to make your point? (You are allowed to use interactive visualizations for machine learning models if you want.)

Note that the four citation requirement is embedded in #2-#5.

Teams

No teams, sorry.

1. Linnaeus, C. (1758). "Felis Catus". Systema naturae per regna tria naturae: secundum classes, ordines, genera, species, cum characteribus, differentiis, synonymis, locis (in Latin). 1 (Tenth reformed ed.). Holmiae: Laurentii Salvii. p. 42.

Option B: Interactive AI System

Create something cool! Build an interactive system that a user can interact with in some form, and that is in some form powered by AI.

Scope

Keep in mind that the end of the semester gets busy, so scope something that will be achievable in 2-ish weeks. Better to have a polished demo that does one interesting thing with AI really well than try to do a ton of things and end up not getting any of them to work well. Make sure your project as a solid *core interaction* that you can achieve (a “minimum goal”, described in the check-in) as well as some ideas you’d like to implement with time but may be more risky (“reach goals”). In other words, try to separate out *must haves* from *would be nice to have* as you brainstorm your interactive system.

Teams

Feel free to work on this project with a developer team of a few of your classmates! Limit to team size is 3 people. I’m going to expect that the more people you add, the more ambitious or polished your final product will be.

Submission format

Code You are welcome to make a local app, a browser app, phone app, or whatever you deem appropriate for the interaction. Regardless of the format, you will need to submit a zip file or a Github repo containing all of your code resources and other resources to make your project run.

Video Walk-through The rationale here is: I can’t always get everyone’s code to replicate on my machine without bugs. To give your project the best possible opportunity for feedback in case I run into problems running it, you will also submit a demo video. The demo video can be a cheap screen capture with a voice over where you film basically what your interaction does and how the user interacts with it. Don’t bother with amazing high-end video quality unless this is something for your own benefit that you’d like to capture for your own professional photo. It just needs to be the quality that we can understand what’s going on and hear your commentary during the walk-through. You can upload this to Glow per your choice as an .mp3 or upload it to Youtube and submit the link through Glow.

Option C: User Research on HAII

Do user research. This option is only available to students who have taken an HCI course or previously done user studies before.

Scope

As you've hopefully figured out this semester, AI impacts people. If you choose this option you'll do some substantive user research with a group of people who are impacted by AI one way or another. This could be a group of people who use an AI-powered device/tool/software/webpage/platform/etc, but it could also be people who *don't* use the thing for an interesting reason (or a mix of people who do and don't use it).

As you (hopefully) know, user research is an open-ended process with many possible methods and approaches. The amount of up-front work that you do should be roughly equivalent to 4(ish) half-hour interviews (contextual or otherwise) if you're working alone or 7(ish) if working in a pair, but you don't necessarily need to do interviews; other user research methods may be more appropriate for the situation you're exploring.

After you've done the above work, you'll also need to analyze the data you've gathered. This could come through affinity diagramming, modeling, survey analysis, etc etc, depending on what research method you chose.

Teams

Feel free to work on this project with a one other classmate who also has human-computer interaction experience! Limit to team size is 2 people for this option. If you're working in pairs, you'll need to interview more people.

Submission format

Scholarly Write-up. In 3-4 pages (1" margins, 1.5 spacing, Times New Roman 12 pt font), answer the following:

1. What technology are you investigating, and who are the users (and/or non-users) you're focusing on?
2. For users (suggested questions, but you can focus on other interesting things that came up in your research too):
    1. How do users use the technology?
    2. Why do users use the technology?
    3. What are the users' mental models of how the technology works, if applicable?
    4. What did users do before they had access to this technology (i.e., what tool or process does this technology replace?)
    5. In what ways is this technology currently failing to serve users effectively (if any)?
3. For non-users (suggested questions, but you can focus on other interesting things that came up in your research too):
    1. Why don't these people use this technology?
    2. What is their mental model of how the technology works, if applicable?

3. What do they do instead of using this technology?

Option D: Image Generation

This is our would-be Assignment 5, which will now be a final project option in order to reduce stress at the end of the semester.

The goal of this assignment is to give you hands-on experience with various techniques for generating images using computer vision + machine learning (CVML) techniques… with dogs.

Scope

The assignment is detailed in the Github repository README, available here: https://github.com/UberHowley/haii-a5. This assignment will be done in Google Colab, so make sure you watch the Google Colab instructional video first.

Teams

No teams, sorry.

Submission format

Submission instructions are detailed in the README in the Github repository.

Final Project Submission

Submit to Glow! This includes your write-up and the other details described above.