

HW4 Debrief, Project 2, AI and labor

15-113 – Monday 2/16

Talk with your neighbors, you know the drill

- Move around and talk to some different people today!
- Take turns briefly showing what you created and how
- Discuss which models you used, **any difficult parts**, anything that seemed particularly effective
- Take notes! We'll ask people to report back briefly

We'll (anonymously) share these results

- I hope the name's right this time (HW4)! <https://bit.ly/4aji4zD>



What unexpected challenges did you find?

Share something you're proud of
or impressed by

Onward: Project 2!

This is very open-ended (though it overlaps the last few projects)

- Purpose: You should build something you're proud of *and* can *comfortably* talk about in interviews!
- This is a two-week project. Be detail-oriented and make it nice! (And don't forget the midpoint check-in)
- Due before spring break. After spring break, we'll do small-group presentations like for Project 1 (but with a little more detail)

Get creative!

You can take a prior project and make it really good if you like.

Or you can make something completely new! It's really up to you, and there's no grade (dis)advantage either way.

Just note that we'll be looking for 5-6 hours of new work for this project, whether that's building something new or improving on something you already have, and we'll be looking for indications that you've learned something and are using AI more effectively to achieve greater goals than you could at the beginning of the course.

HW5

Two articles to read (will take about an hour).
Take notes, and be ready to discuss!

- Due Wednesday 2/18 before class: Read and be ready to discuss “Cognitive Helmets for the AI Bicycle” here: <https://bit.ly/4tkll81>
- We’ll have another (TBD) due Monday 2/23 before class
 - This one will get us prepped for Illah, who’ll be giving our lecture on Thursday 2/25 (right before break)!

Discussion time!

Please participate! Voice your thoughts, ask questions! This is only going to be valuable if you're willing to add to it.

- It's ok (and probaby even good) to be uncertain about what you know or feel or think. Voice that too!
- I know you'll all respect each other. On top of that, let's add a house rule for anonymity: If you tell anyone about our discussions outside of class, please don't attach names to anyone's comments.
- On my end, I promise to represent your interests to the university, other faculty, etc.

First Topic:

One of the major promises of generative AI is that it will make us more productive. For example, we'll be able to generate code faster than we could write it ourselves.


- Do you feel this is already true?
- Do you feel that it is likely to be true in the near future?

For the moment, assume the standard of productivity is that the output “gets the job done,” and that we’re talking about low-creativity low-risk tasks. We’ll get to questions of safety, artistic merit, education etc later.

If modern *AI can* make us more productive...

- Who benefits from this?
- How might this change how much money you make, or how much free time you have?
- How might this change how you feel about the work that you do?

An aside: What's this?



Eli Whitney's cotton gin, 1793

- Function: Automates and *massively* accelerates cotton seed/fiber separation by 4900%
- Along with other industrialization, this created a bottleneck: harvesting the cotton
- Rather than easing the workload, a 'productivity tool' skyrocketed the demand for slave labor in the field, resulting in longer hours, worse conditions, and higher expectations
- Back to AI: Does this seem relevant? Or do you think this is too far in the past?

More AI and labor, learning, and sci-fi

15-113 – Wednesday 2/18

Picking up from last time:

- Excerpt from 60 Minutes on AI supervision and training (Nov 2024):
<https://www.youtube.com/watch?v=qZS50KXjAX0>
- CW: Corporate exploitation, poverty, descriptions of physical/sexual violence in the context of AI-generated imagery
 - General note: You are always welcome to step out of class if needed

What are your thoughts?

- Fill this out and then talk with each other <https://bit.ly/4ruKZHc>



Let's pivot to a different issue

- “Cognitive Helmets for the AI Bicycle”: <https://bit.ly/4tkll81>
- How many of you read at least some of the article?
- How many of you read it carefully?

Ask this to yourself:

- Would you have trouble recalling and confidently describing the key points in the article, even though you read it recently?

Here's what I think:

- Many of us (myself included) think we're reading, but we often end up just *looking*. We sometimes skim the words, and then actively think about a sub-selection of what we skim. We can read carefully by being very conscious and deliberate, but it's *hard*.
- ...do you agree or disagree with that?

Here are some generally-accepted notes on learning:

- We learn more by doing an action compared with watching someone or something do the same action.
- We learn *best* with repetition and proper spacing.
- Effective problem-solving is learned, not intrinsic.
- Problem-solving skills can transfer across domains/disciplines.
- Metacognition (thinking about how you think) helps you learn faster and become a better problem-solver

How might AI be hindering our learning?

- One might engage with problems more superficially
- One might resort to AI use when things get challenging
- One's learned problem-solving habits might rely on specific tech
- One might miss out on participatory problem-solving with others
- One might easily confuse *seeing* a solution with understanding (or even having written) the solution

Cat Hicks' article suggests:

Most broadly, that we think deliberately about problems and the outputs AI gives us. Pause before and after prompting.

- Try to force yourself to solve problems that are within or just beyond your skills
- Pause and try to predict the general outcome before asking AI, and then reflect on whether your expectations were met afterward
- Take notes, write documentation, formulate explanations
- Always critically evaluate the output of AI

Any other thoughts?

HW5: By Monday, read one of these!

***The Evitable Conflict* by Isaac Asimov:** Deals with algorithmic control of sociopolitical systems and unintended consequences of (mis)alignment

***Runaround* by Isaac Asimov:** Also relevant to prompting and alignment, but on a smaller scale

***R.U.R. (Rossum's Universal Robots)* by Karel Čapek:** Introduced the word “robot” to the world. Deals with various modern motivations for accelerating technology (money, power, pursuit of utopia, desire to create life) and the consequences of a fully-automated workforce. (Note: the romance subplot is very dated...)

- This is a play. If you'd prefer to watch instead of reading, that's ok.

I'll find PDF copies of all of these and post them online tonight.

For the remainder of today:

- First, spend some time ideating with your neighbors on what you might like to create for Project 2. Use the blackboards!
- Choose something you care about, something that feels achievable, and something that feels explainable.
- Then, use the rest of the time today to get started! The hardest part of most projects is just taking the first step, so take it before you leave the room today!

Note: For the required mid-semester checkin, you'll do these in OH like you did for Project 1. Come to OH any time starting tomorrow until Friday 2/27.