

CS182: Ethics, Public Policy, and Technological Change 2025-26
Assignment 2

Due: Thursday, February 12, 2026 no later than 11:59pm

Length: 1800-2200 words

For students enrolled in the Writing in the Major (WIM) version of the class, you will receive comments on this paper, and there will be a required revision.

The papers will be graded blindly. Submit the paper with **YOUR STANFORD ID** number on a cover page along with the **TITLE OF THE PAPER**, the **QUESTION YOU ARE ANSWERING**, and the **WORD COUNT** (not including footnotes/bibliography).

Your section leader will give you additional information regarding how to submit the assignment.

GENERAL INSTRUCTIONS

Choose one of the following paper topics. Your essay should state and defend a thesis. *You must use and cite specific examples from our readings to illustrate and support your argument.*

Although we encourage you to draw on lectures and section discussions to elucidate your argument, you must foreground your own insights and analysis of multiple readings and demonstrate your use of them beyond merely revisiting quotes or arguments that were discussed in lectures and sections. You are encouraged – but not required – to consult supplemental readings listed on the syllabus or that you identify in your own research.

Statement on use of AI

AI tools have an important place in educational and, more broadly, professional settings. But given our learning goals on the philosophy paper (see the following page), the use of AI to produce text will undermine rather than augment your learning. It can also constitute cheating. We have prompted the major AI tools (ChatGPT, Claude, Gemini) with the essay questions below, and based on the output, AI-written papers earn a low B, lacking nuance, creativity, and philosophical sophistication. So using such tools alone is a likely recipe for a middling grade.

To be clear, AI can be useful in exploring and refining ideas, for reviewing your own writing, and for quick formatting and re-formatting of your paper. If you use AI in any way, you are required to disclose your use

All students must submit (as an appendix to their paper) a document that discloses their use of AI tools. This appendix should include:

- List of the AI tools or systems used by the student (e.g., Claude, ChatGPT, Gemini, etc.)
- **Fully inclusive list** of the student-generated prompts (including all multi-turn interactions)

Stanford's plagiarism policy:

<https://communitystandards.stanford.edu/policies-and-guidance/bja-guidance-definitions-and-clarifications/what-plagiarism/plagiarism>

Do not plagiarize. Plagiarism frequently happens unintentionally, so if you are in doubt about your use of external sources, ask your instructor.

Some Guidelines on Writing a Philosophy Paper.

For any student who has never written a philosophy paper before, we strongly encourage you to read the paper writing guides on the course website and Canvas page.

Many students believe that writing a paper in philosophy is unlike any other kind of writing. This is an exaggeration. Writing a good philosophy paper is very much like writing a good paper in history, economics, literature, or biology. A good paper requires a particularly careful use of language and a particularly close examination of ideas and arguments. Because a good paper is a good paper, the guidelines sketched here apply with equal force in other courses as well.

It's important, however, to stress just how central good writing is when doing philosophy. As a general matter, your writing is a good measure of your ability to communicate ideas. But in philosophy, the quality of your writing is not just a measure of your ability to communicate; it is also a measure of your ability to think. If you cannot express your thoughts in writing in a clear, concise, and cohesive manner, then your thoughts themselves are not clear, concise, or cohesive. Philosophical writing exercises your thinking; learning to write better is inseparable from learning to think better. In philosophy, becoming a good writer is the same thing as become a good thinker.

For this essay on privacy, in addition to the many resources on the syllabus for this module, we recommend (but do not require) consulting some new books on privacy that have been published recently.

[Carissa Véliz, *Privacy Is Power: Why and How You Should Take Back Control of Your Data* \(Penguin Press, 2022\).](#)

Neil Richards, [*Why Privacy Matters* \(Oxford University Press, 2022.\)](#)

[Danielle Keats Citron, *The Fight For Privacy: Protecting Dignity, Identity, and Love in the Digital Age* \(WW Norton, 2022\).](#)

Lowry Pressly, [*A Right to Oblivion: Privacy and the Good Life* \(Harvard U Press, 2025\).](#)

Kashmir Hill, [*Your Face Belongs to Us* \(Penguin, 2023\).](#)

Daniel Solove, [*On Privacy and Technology* \(Oxford University Press, 2025\)](#)

In addition, there are many excellent online resources for doing further research in philosophy. We especially recommend the Stanford Encyclopedia of Philosophy, described as ("[the most interesting website on the internet](#)"). These entries in particular may be helpful:

Privacy: <https://plato.stanford.edu/entries/privacy/>

Privacy and Information Technology, <https://plato.stanford.edu/entries/it-privacy/>

Informed Consent: <https://plato.stanford.edu/entries/informed-consent/>

Personal Autonomy: <https://plato.stanford.edu/entries/personal-autonomy/>

Philosophy of Technology: <https://plato.stanford.edu/entries/technology/>

Artificial Intelligence: <https://plato.stanford.edu/entries/artificial-intelligence/>

History of Utilitarianism: <https://plato.stanford.edu/entries/utilitarianism-history/>

Utilitarianism/Consequentialism: <https://plato.stanford.edu/entries/consequentialism/>

Happiness: <https://plato.stanford.edu/entries/happiness/>

Well-Being: <https://plato.stanford.edu/entries/well-being/>

Rawls's Theory of Justice: <https://plato.stanford.edu/entries/rawls/>

Isaiah Berlin and Value Pluralism: <https://plato.stanford.edu/entries/berlin/>

Value Pluralism: <https://plato.stanford.edu/entries/value-pluralism>

Question 1: The Surveillance Advertising Economy

In a 2023 [Consumer Reports study](#) investigating the surveillance economy, over 700 volunteers downloaded their Facebook data using Facebook's "Download Your Information" tool and donated it to researchers for analysis. This tool reveals which companies have shared user data with Facebook for advertising purposes, including both "Custom Audiences" (lists of personal identifiers like email addresses and phone numbers used for ad targeting) and "Events" (records of user actions like website visits, purchases, or app usage).

The study ([full version here](#)) found that the average participant was identified in data shared by 2,230 different companies with Facebook, with some participants identified by over 7,000 companies. In total, researchers found slightly more than 186,000 different companies represented in the donated data. The study describes the surveillance economy as "cross-contextual," meaning it "uses information about individuals that's been collected in one context—such as a website visit, an action taken in an app, or a visit to a physical location—and applies it to another context to affect how you are advertised to, what prices you see, and how you are otherwise treated."

Consider these examples from the study:

- A small-town car dealership in San Benito, Texas (population 24,665) shared data on approximately 10% of study volunteers nationwide
- The Illinois Lottery shared data on nearly 70% of volunteers across the country
- 99% of participants were identified by at least one company with an unidentifiable name
- Data brokers like LiveRamp appeared in 96% of participant records

Helen Nissenbaum writes: "As social contexts, activities, roles, and rules migrate online, respective context-specific values, ends, and purposes serve as standards against which information-sharing practices can be evaluated as legitimate or problematic."

Elucidate Nissenbaum's argument concerning context-specific values, ends, and purposes as standards to evaluate information-sharing practices in the surveillance economy. Apply Nissenbaum's contextual integrity approach to two or three of the examples listed above. Your selection of examples should illuminate different privacy considerations within the framework of contextual integrity. Be sure to identify and consider strong arguments against Nissenbaum's own view (drawn from other readings on the syllabus) as applied to the examples you discuss.

Question 2: Tech and Immigration and Customs Enforcement (ICE)

News headlines over the past few months have focused world-wide attention on President Trump's efforts to identify, arrest, and deport undocumented immigrants in the United States. Much attention has focused on certain tactics used by ICE agents in the streets, e.g., wearing masks, heavily weaponized, surges into particular cities. Setting aside your views about these tactics, the identification and tracking of potential targets, and of protestors, has been supported by digital tools developed by technologists. Here are some representative articles that cover some of the main contracts between US ICE and tech companies.

New York Times, "[How ICE Already Knows Who Minneapolis Protestors Are](#)"

Washington Post, "[The Powerful Tools in ICE's Arsenal to Track Suspects – and Protestors](#)"

Note also [this blog post](#) by Palantir in January 2026

[Note: you can access these articles using Stanford Library if they are paywalled]

One of the [most important contracts](#) is with a company called Palantir Technologies Inc., [named after the all-seeing stones in J.R.R. Tolkien's *The Lord of the Rings*](#). Palantir specialized in data

integration and analysis and provides technologies that support various governments and their military, defense, and intelligence gathering efforts. Their primary civil offering is their [Foundry](#) platform, which aggregates information from separated data sources and provides government leaders a unified picture through an interactive “ontology” and counterfactual simulation via digital twins. Foundry underlies many Palantir deployments including, but not limited to (1) [HHS Protect](#), which guided the US COVID-19 response by [ingesting medical data from 5700 hospitals across the country](#), (2) [the United Nations World Food Programme](#), (3) the [Disaster Victim Database for earthquake response in Japan](#), and (4) and [ImmigrationOS](#), [Elite](#), and other tools support US Immigrations and Customers Enforcement (ICE).

Assess the privacy considerations involved in a technical system that constructs a picture of an individual through the consolidation of information possessed by all, or nearly all, organizations and institutions they have interacted with across their lives. Your analysis should include (at a minimum) the following:

- Discussion of at least three of the following authors (Nissenbaum and contextual integrity, Foucault on surveillance as power, Veliz on privacy and power, Solove).
- Discussion of the practical implications of your analysis. Should such data integration and analysis systems not be built? If so, what about all the potentially beneficial deployments? Should the deployment of such systems be restricted to particular uses? What principle(s) could identify permissible and impermissible uses? How does privacy figure into an answer, if at all?

Question 3. Privacy, the AI-Augmented Self, and the End of Forgetting

It is likely that 2026 will see the release of always-on AI wearables. Apple is [said to be](#) working on this, [as are Sam Altman and Jony Ive](#) at OpenAI. These devices use always-listening microphones and potentially cameras to record ambient interactions, uploading the data to create a searchable "Personal Memory" of every conversation and event that the user experiences. Imagine that advanced models could offer "Daily Insights," using sentiment analysis to monitor shifts in the wearer's relationships and emotional well-being based on the tone of their social interactions.

To highlight one potential privacy issue, consider a scenario that may already be present. Imagine that you discover one of your closest friends has been using an AI agent to help manage their communications, including sorting through emails, photos, social media feeds, and text messages. This means that sensitive personal information you've shared with them - secrets, struggles - has been processed by this AI system and potentially stored on corporate servers (potentially in a country where privacy law is lax). Your friend didn't ask for your consent, and there is no regulation requiring them to do so.

Write an essay identifying and analyzing the privacy concerns of these body-worn AI systems. In your response:

- **Analyze the Bystander Problem:** How might these devices affect the privacy interests of non-users (bystanders, friends as well as strangers) who have not provided informed consent to be part of another person's digital archive?
- **Contextual Integrity:** Using Nissenbaum's framework, evaluate whether the migration of private, unrecorded conversations into a "digitally indexed and searchable" format violates the norms of social contexts like a private dinner or a confidential workplace meeting.

- **Power and Surveillance:** Discuss whether wearing such a device creates a "power asymmetry" in social interactions. What might it mean, positive or negative, for relationships if AI pins make possible the ability to "rewind" and "fact-check" a conversation using AI-generated transcripts?
- Based on your analysis, what are the implications for a tech company designing an AI-pin, for the norms of human interaction in a world of AI pins, or for new laws and regulation needed to protect privacy (as you define and defend it).

Your essay should incorporate discussion of at least three authors from the syllabus (e.g., Nissenbaum, Solove, Véliz, Foucault, or Lowry Pressly's new book on [A Right to Oblivion](#)).