TechXplore

Privacy in the AI era: How do we protect our personal information?

March 20 2024, by Katharine Miller

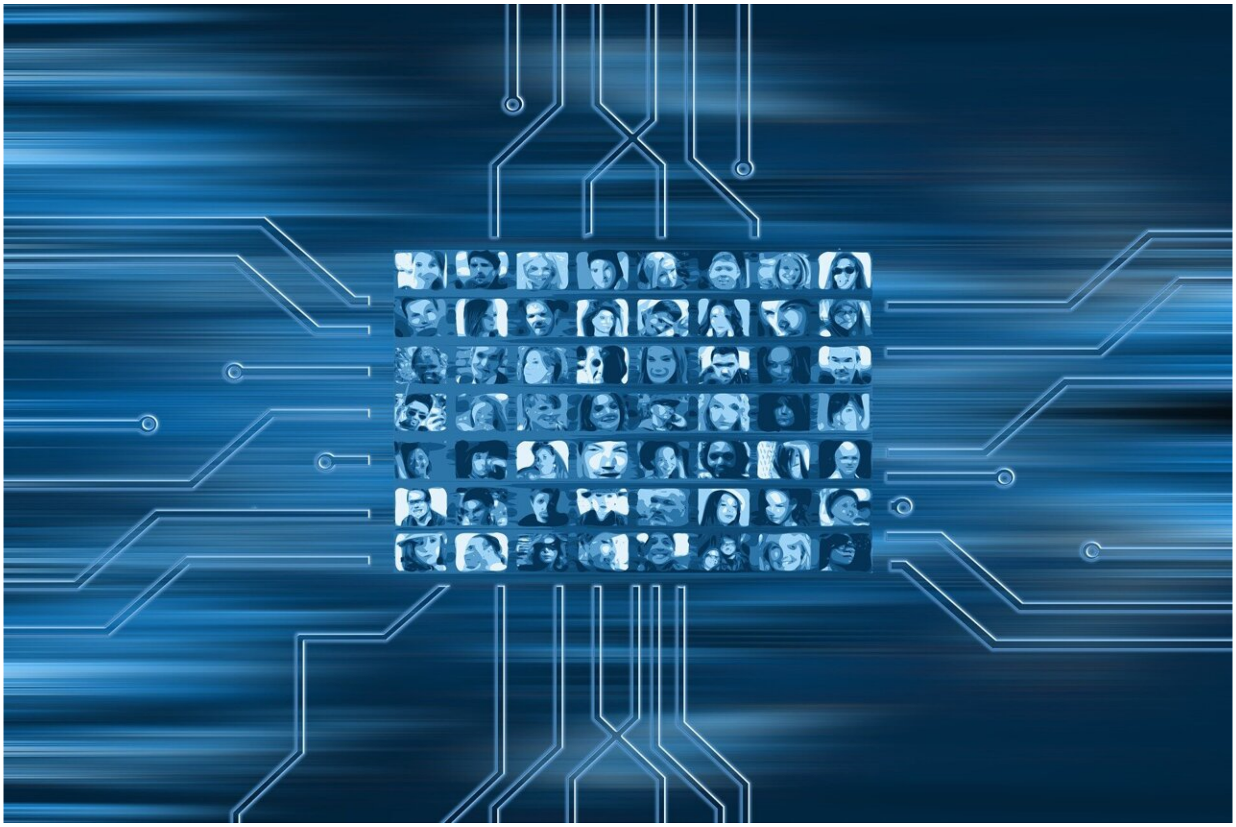

Credit: Pixabay/CC0 Public Domain

The AI boom, including the advent of large language models (LLMs) and their associated chatbots, poses new challenges for privacy. Is our personal information part of a model's training data? Are our prompts being shared with law enforcement? Will chatbots connect diverse threads from our online lives and output them to anyone?

To better understand these threats and to wrestle with potential solutions, Jennifer King, privacy and data policy fellow at the Stanford University Institute for Human-Centered Artificial Intelligence (Stanford HAI), and Caroline Meinhardt, Stanford HAI's policy research manager, published a white paper titled "Rethinking Privacy in the AI Era: Policy Provocations for a Data-Centric World." Here, King describes their main findings.

What kinds of risks do we face, as our data is being bought and sold and used by AI systems?

First, AI systems pose many of the same privacy risks we've been facing during the past decades of internet commercialization and mostly unrestrained data collection. The difference is the scale: AI systems are so data-hungry and intransparent that we have even less control over what information about us is collected, what it is used for, and how we might correct or remove such personal information. Today, it is basically impossible for people using online products or services to escape systematic digital surveillance across most facets of life—and AI may make matters even worse.

Second, there's the risk of others using our data and AI tools for anti-social purposes. For example, generative AI tools trained with data scraped from the internet may memorize personal information about people, as well as relational data about their family and friends. This data helps enable spear-phishing—the deliberate targeting of people for purposes of identity theft or fraud. Already, bad actors are using AI voice cloning to impersonate people and then extort them over good old-fashioned phones.

Third, we're seeing data such as a resume or photograph that we've shared or posted for one purpose being repurposed for training AI systems, often without our knowledge or consent and sometimes with direct civil rights implications.

Predictive systems are being used to help screen candidates and help employers decide whom to interview for open jobs. However, there have been instances where the AI used to help with selecting candidates has been biased. For example, Amazon famously built its own AI hiring screening tool only to discover that it was biased against female hires.

Another example involves the use of facial recognition to identify and apprehend people who have committed crimes. It's easy to think, "It's good to have a tool like facial recognition because it'll catch the bad guys." But instead, because of the bias inherent in the data used to train existing facial recognition algorithms, we're seeing numerous false arrests of black men. The algorithms simply misidentify them.

Have we become so numb to the idea that companies are taking all our data that it's now too late to do anything?

I'm an optimist. There's certainly a lot of data that's been collected about all of us, but that doesn't mean we can't still create a much stronger regulatory system that requires users to opt in to their data being collected or forces companies to delete data when it's being misused.

Currently, practically any place you go online, your movement across different websites is being tracked. And if you're using a mobile app and you have GPS enabled on your phone, your location data is being collected. This default is the result of the industry convincing the Federal Trade Commission about 20 years ago that if we switched from opt-out to opt-in data collection, we'd never have a commercial internet. At this point I think we've established the utility of the internet. I don't think companies need that excuse for collecting people's data.

In my view, when I'm browsing online, my data should not be collected unless or until I make some affirmative choice, like signing up for the service or creating an account. And even then, my data shouldn't be considered public unless I've agreed to share it.

Ten years ago, most people thought about data privacy in terms of online shopping. They thought, "I don't know if I care if these companies know what I buy and what I'm looking for, because sometimes it's helpful." But now we've seen companies shift to this ubiquitous data collection that trains AI systems, which can have major impact across society, especially our civil rights. I don't think it's too late to roll things back. These default rules and practices aren't etched in stone.

As a general approach to data privacy protection, why isn't it enough to pass data minimization and purpose limitation regulations that say companies can only gather the data they need for a limited purpose?

These types of rules are critical and necessary. They play a key role in the European privacy law [the GDPR] and in the California equivalent [the CPPA] and are an important part of the federally proposed privacy law [the ADPPA]. But I'm concerned about the way regulators end up operationalizing these rules.

For example, how does a regulator make the assessment that a company has collected too much information for the purpose for which it wants to use it? In some instances, it could be clear that a company completely overreached by collecting data it didn't need. But it's a more difficult question when companies (think Amazon or Google) can realistically say that they do a lot of different things, meaning they can justify collecting a lot of data. It's not an insurmountable problem with these rules, but it's a real issue.

Your white paper identifies several possible solutions to the data privacy problems posed by AI. First, you propose a shift from opt-out to opt-in data sharing, which could be made more seamless using software. How would that work?

I would argue that the default should be that our data is not collected unless we affirmatively ask for it to be collected. There have been a few movements and tech solutions in that direction.

One is Apple's App Tracking Transparency (Apple ATT), which Apple launched in 2021 to address concerns about how much user data was being collected by third-party apps. Now, when iPhone users download a new app, Apple's iOS system asks if they want to allow the app to track them across other apps and websites. Marketing industry reports estimate that 80% to 90% of people presented with that choice say no.

Another option is for web browsers to have a built-in opt-out signal, such as Global Privacy Control, that prevents the placement of cookies by third parties or the sale of individuals' data without the need to check a box. Currently, the California Privacy Protection Act (CPPA) provides that browsers may include this capability, but it has not been mandatory. And while some browsers (Firefox and Brave, for example) have a built-in op-out signal, the big browser companies (such as Microsoft Edge, Apple's Safari, and Google Chrome) do not. Interestingly though, a California legislator recently proposed a change to the CPPA that would require all browser makers to respect third-party opt-out signals. This is exactly what we need so that data is not collected by every actor possible and every place you go.

You also propose taking a supply chain approach to data privacy. What do you envision that would mean?

When I'm talking about the data supply chain, I'm talking about the ways that AI systems raise issues on the data input side and the data output side. On the input side I'm referring to the training data piece, which is where we worry about whether an individual's personal information is being scraped from the internet and included in a system's training data. In turn, the presence of our personal information in the training set potentially has an influence on the output side. For example, a generative AI system might have memorized my personally identifiable information and provide it as output. Or, a generative AI system could reveal something about me that is based on an inference from multiple data points that aren't otherwise known or connected and are unrelated to any personally identifiable information in the training dataset.

At present, we depend on the AI companies to remove personal information from their training data or to set guardrails that prevent personal information from coming out on the output side. And that's not really an acceptable situation, because we are dependent on them choosing to do the right thing.

Regulating AI requires paying specific attention to the entire supply chain for the data piece—not just to protect our privacy, but also to avoid bias and improve AI models. Unfortunately, some of the discussions that we've had about regulating AI in the United States haven't been dealing with the data at all. We've been focused on transparency requirements around the purpose of companies' algorithmic systems. Even the AI Act in Europe, which already has the GDPR as a privacy baseline, didn't take a broad look at the data ecosystem that feeds AI. It was only mentioned in the context of high-risk AI systems. So, this is an area where there is a lot of work to do if we're going to have any sense that our personal information is protected from inclusion in AI systems, including very large systems such as foundation models.

You note in your report that the focus on individual privacy rights is too limited and we need to consider collective solutions. What do you mean?

If we want to give people more control over their data in a context where huge amounts of data are being generated and collected, it's clear to me that doubling down on individual rights isn't sufficient.

In California where we have a data privacy law, most of us don't even know what rights we do have, let alone the time to figure out how to exercise them. And if we did want to exercise them, we'd have to make individual requests to every company we've interacted with to demand that they not sell our personal information—requests that we'd have to make every two years, given that these "do not sell" opt-outs are not permanent.

This all points toward the need for a collective solution so that the public has enough leverage to negotiate for their data rights at scale. To me, the concept of a data intermediary makes the most sense. It involves delegating the negotiating power over your data rights to a collective that does the work for you, which gives consumers more leverage.

We're already seeing data intermediaries take shape in some business-to-business contexts and they can take various forms, such as a data steward, trust, cooperative, collaborative, or commons. Implementing these in the consumer space would be more challenging, but I don't think it's impossible by any means.

More information: Rethinking Privacy in the AI Era: Policy Provocations for a Data-Centric World. hai.stanford.edu/white-paper-r … s-data-centric-world

Provided by Stanford University

Citation: Privacy in the AI era: How do we protect our personal information? (2024, March 20) retrieved 9 May 2024 from https://techxplore.com/news/2024-03-privacy-ai-era-personal.html