Statement of Maciej Cegłowski, Founder, Pinboard
Before the U.S. Senate Committee on Banking, Housing, and Urban Development
On the topic of "Privacy Rights and Data Collection in a Digital Economy"

May 7, 2019

Thank you for the opportunity to address you today.

I am the founder and sole employee of Pinboard, a small for-profit archiving service founded in 2009 that competes in part on the basis of personal privacy. I have also been a frequent critic of Silicon Valley's reliance on business models requiring mass surveillance, speaking on the topic at conferences both in the United States and abroad.

As someone who earns his living through data collection, I am acutely aware of the power the tools we are building give us over our fellow citizens' private lives, and the danger they pose to our liberty. I am grateful to Chairman Crapo, ranking member Brown, and the committee for the opportunity to testify on this vital matter.

The internet economy in 2019 is dominated by five American tech companies: Apple, Microsoft, Google, Facebook, and Amazon. These are also the five most valuable corporations in the world, with a combined market capitalization exceeding four trillion dollars[1]. Between them, these companies control the market for online advertising, mobile and desktop operating systems, office software, document storage, search, cloud computing, and many other areas of the digital economy. They also own and operate a significant portion of the physical infrastructure of the internet, and act as its *de facto* regulating authority.

The concentration of power in the hands of these giant firms is the epilogue to a spectacular story of American innovation and dynamism. The technologies underpinning the internet were all developed here in the United States, and the many fortunes that they produced owe their thanks to fruitful cooperation between government, industry, and the research community. Working together, the public and private sectors created the conditions for a startup culture unlike any other in the world.

Today, however, that culture of dynamism is at risk. The surveillance business model has eroded user trust to such a point that it is impeding our ability to innovate.

In many ways, the five internet giants operate like sovereign states. Their operations are global, and decisions they take unilaterally can affect entire societies. Denmark has gone so far as to send an ambassador to Silicon Valley. When Jeff Bezos, the CEO of Amazon, met recently with the Canadian prime minister, the occasion was covered in the press like a state visit.

[1] At the time of writing, Amazon was valued at $966B, Microsoft $988B, Apple $974B, Facebook $558B, and Google (Alphabet) $824B.

The emergence of this tech oligopoly reflects a profound shift in our society, the migration of every area of commercial, social, and personal life into an online realm where human interactions are mediated by software.

To an extent that has no precedent, the daily activities of most Americans are now tracked and permanently recorded by automated systems. It is likely that every person in this hearing room carries with them a mobile phone that keeps a history of their location, is privy to their most private conversations, and contains a rich history of their private life. Some of you may even have an always-on microphone in your car or home that responds to your voice commands.

Emerging technologies promise to afford these systems even more intimate glimpses into our private lives—phones that monitor our facial expressions as we read, and connected homes that watch over us while we sleep. Scenarios that were once the province of dystopian dime fiction have become an unremarkable consumer reality.

The sudden ubiquity of this architecture of mass surveillance, and its enshrinement as the default business model of the online economy, mean that we can no longer put off hard conversations about the threats it poses to liberty.

Adding to this urgency is the empirical fact that, while our online economy depends on the collection and permanent storage of highly personal data, we do not have the capacity to keep such large collections of user data safe over time.

The litany of known data breaches is too long to recite here, but includes every one of the top five tech companies, as well as health and financial firms and government agencies. Every year brings new and more spectacular examples of our inability to protect our users. At Yahoo, an internet giant at the time with a world-class security team, over 3 billion user accounts were compromised in a 2013 breach. In 2015, the US Office of Personnel Management allowed unauthorized access to the records of over four million people, including many with highly sensitive security clearances. And in 2017, Equifax exposed data, including social security numbers, on 147 million Americans, nearly half the US population.

While many individual data breaches are due to negligence or poor practices, their overall number reflects an uncomfortable truth well known to computer professionals—that our ability to attack computer systems far exceeds our ability to defend them, and will for the foreseeable future.

The current situation, therefore, is not tenable. The internet economy today resembles the earliest days of the nuclear industry. We have a technology of unprecedented potential, we have made glowing promises about how it will transform the daily lives of our fellow Americans, but we don't know how to keep its dangerous byproducts safe.

Two Views of Privacy

Discussing privacy in the context of regulation can be vexing, because the companies doing the most to erode our privacy are equally sincere in their conviction that they are its champions.

The confusion stems from two different ways in which we use the word privacy, leading us to sometimes talk past each other.

In the regulatory context, discussion of privacy invariably means data privacy—the idea of protecting designated sensitive material from unauthorized access.

Laws like the Health Insurance Portability and Accountability Act (HIPAA) and the Gramm-Leach-Bliley Act (GLBA) delimit certain categories of sensitive information that require extra protection, and mandate ways in which health and financial institutions have to safeguard this data, or report when those safeguards have failed. The Children's Online Privacy Protection Act of 1998 extends similar protection to all data associated with children.

We continue to use this framework of data privacy today, including in the recently enacted General Data Protection Regulation (GDPR).

It is true that, when it comes to protecting specific collections of data, the companies that profit most from the surveillance economy are the ones working hardest to defend them against unauthorized access.

But there is a second, more fundamental sense of the word privacy, one which until recently was so common and unremarkable that it would have made no sense to try to describe it.

That is the idea that there exists a sphere of life that should remain outside public scrutiny, in which we can be sure that our words, actions, thoughts and feelings are not being indelibly recorded. This includes not only intimate spaces like the home, but also the many semi-private places where people gather and engage with one another in the common activities of daily life—the workplace, church, club or union hall. As these interactions move online, our privacy in this deeper sense withers away.

Until recently, even people living in a police state could count on the fact that the authorities didn't have enough equipment or manpower to observe everyone, everywhere[2], and so enjoyed more freedom from monitoring than we do living in a free society today.

A characteristic of this new world of ambient surveillance is that we cannot opt out of it, any more than we might opt out of automobile culture by refusing to drive. However sincere our commitment to walking, the world around us would still be a world built for cars. We would still have to contend with roads, traffic jams, air pollution, and run the risk of being hit by a bus.

Similarly, while it is possible in principle to throw one's laptop into the sea and renounce all technology, it is no longer be possible to opt out of a surveillance society.

---

[2] The record for intensive surveillance in the pre-internet age likely belongs to East Germany, where by some estimates one in seven people was an informant. (https://archive.nytimes.com/www.nytimes.com/books/first/k/koehler-stasi.html)

When we talk about privacy in this second, more basic sense, the giant tech companies are not the guardians of privacy, but its gravediggers.

The tension between these interpretations of what privacy entails, and who is trying to defend it, complicates attempts to discuss regulation.

Tech companies will correctly point out that their customers have willingly traded their private data for an almost miraculous collection of useful services, services that have unquestionably made their lives better, and that the business model that allows them to offer these services for free creates far more value than harm for their customers.

Consumers will just as rightly point out that they never consented to be the subjects in an uncontrolled social experiment, that the companies engaged in reshaping our world have consistently refused to honestly discuss their business models or data collection practices, and that in a democratic society, profound social change requires consensus and accountability.

Behavioral Data

Further complicating the debate on privacy is the novel nature of the data being collected. While the laws around protecting data have always focused on intentional communications—documents that can be intercepted, conversations that can be eavesdropped upon—much of what computer systems capture about us is behavioral data: incidental observations of human behavior that don't seem to convey any information at all.

Behavioral data encompasses anything people do while interacting with a computer system. It can include the queries we type into a search engine, our physical location, the hyperlinks we click on, whether we are sitting or standing, how quickly we scroll down a document, how jauntily we walk down a corridor, whether our eyes linger on a photo, whether we start to write a comment and then delete it—even the changes in our facial expression as we are shown an online ad.

This incidental data has proven to be such a valuable raw material that an entire industry now specializes in finding ways to mine it. The devices used to spy on us include our computers, cell phones, televisions, cars, security cameras, our children's toys, home appliances, wifi access points, even at one point trash cans in the street[3].

Privacy and Consent

[3] Campbell-Dollaghan, Kelsey. "Brave New Garbage: London's Trash Cans Track You Using Your Smartphone". *Gizmodo*. (Aug 9, 2013) https://gizmodo.com/brave-new-garbage-londons-trash-cans-track-you-using-1071610114

The extent to which anyone consents—or *can* consent—to this kind of tracking is the thorny question in attempting to regulate the relationship between people and software.

The General Data Protection Regulation (GDPR), enacted in May of 2018, is the most ambitious attempt thus far to regulate online privacy. It takes a very traditional view of the relationship between people and data.

In the eyes of the GDPR, people own their data. They make an affirmative choice to share their data with online services, and can revoke that choice. The consent they give must be explicit and limited to a specified purpose—the recipient does not have *carte blanche* to use the data as they please, or to share it with third parties, with some complicating caveats.

People have the right to request a full download of their data from the services they have entrusted it to, and they have the right to demand that it be permanently erased.

The GDPR imposes a notification requirement for data breaches, and requires affirmative consent for the sale of user data. It also restricts the movement of data to outside jurisdictions (though in the case of the United States, this restriction is superseded by the US-EU Privacy Shield framework).

Finally, the GDPR mandates that privacy safeguards like data tokenization and encryption be built in to new systems, and that companies appoint a dedicated privacy officer.

The GDPR is not a simple regulation, and many of its most potentially significant provisions (such as the scope of a data controller's 'legitimate interests', or what the right to erasure means in the context of a machine learning model) await interpretation by regulators.

What limits, if any, the GDPR will place on the application of machine learning is a particularly important open question. The law on its face prohibits automated decision making that has a "legal or similarly significant effect" on data subjects, but the definition of "significant effect" is not clear, nor is it clear whether having a human being simply countersign an algorithmic decision would be enough to satisfy regulators that the decision process is not fully automated.

Impacts

As it is so new, the GDPR's ultimate impact on online privacy in the EU is unclear. Some of the dramatic early impacts (like major US newspapers going offline) have proven to be transient, while many of the biggest impacts hinge on future decisions by EU regulators.

Enough has happened, however, to draw some preliminary conclusions.

The GDPR so far has made life hard on internet users. It is not clear that this is the GDPR's fault.

The plain language of the GDPR is so plainly at odds with the business model of surveillance advertising that contorting the real-time ad brokerages into something resembling compliance has required acrobatics that have left essentially everybody unhappy.

The leading ad networks in the European Union have chosen to respond to the GDPR by stitching together a sort of Frankenstein's monster of consent, a mechanism whereby a user wishing to visit, say, a weather forecast page[4] is first prompted to agree to share data with a consortium of 119 entities, including the aptly named "A Million Ads" network. The user can scroll through this list of intermediaries one by one, or give or withhold consent *en bloc*, but either way she must wait a further two minutes for the consent collection process to terminate before she is allowed to find out whether or not it is going to rain.

This majestically baroque consent mechanism also hinders Europeans from using the privacy preserving features built into their web browsers, or from turning off invasive tracking technologies like third-party cookies, since the mechanism depends on their being present.

For the average EU citizen, therefore, the immediate effect of the GDPR has been to add friction to their internet browsing experience along the lines of the infamous 2011 EU Privacy Directive ("EU cookie law") that added consent dialogs to nearly every site on the internet.

The GDPR rollout has also demonstrated to what extent the European ad market depends on Google, who has assumed the role of *de facto* technical regulatory authority due to its overwhelming market share[5]. Google waited until the night before the regulation went into effect to announce its intentions, leaving ad networks scrambling.

It is significant that Google and Facebook also took advantage of the US-EU privacy shield to move 1.5 billion non-EU user records out of EU jurisdiction to servers in the United States. Overall, the GDPR has significantly strengthened Facebook and Google at the expense of smaller players in the surveillance economy.

The data protection provisions of the GDPR, particularly the right to erase, imposed significant compliance costs on internet companies. In some cases, these compliance costs just show the legislation working as intended. Companies who were not keeping adequate track of personal data were forced to retrofit costly controls, and that data is now safer for it.

But in other cases, companies with a strong commitment to privacy also found themselves expending significant resources on retooling. Personally identifying information has a way of seeping in to odd corners of computer systems (for example, users will sometimes accidentally paste their password into a search box), and tracking down all of these special cases can be challenging in a complex system. The requirements around erasure, particularly as they interact with backups, also impose a special burden, as

---

[4] This is an actual example.

[5] Google has at least a 70% advertising market share in Europe, though this figure is averaged over the ten year period 2006-2016 and likely far higher today. Laurent, Lionel. "Europe Is Changing Google for the Better". Washington Post (March 20, 2019). https://www.washingtonpost.com/business/europe-is-changing-google-for-the-better/2019/03/20/691aaff4-4b2e-11e9-8cfc-2c5d0999c21e_story.html

most computer systems are designed with a bias to never losing data, rather than making it easy to expunge.

A final, and extremely interesting outcome of the GDPR, was an inadvertent experiment conducted by the New York Times. Privacy advocates have long argued that intrusive third-party advertising does not provide more value to publishers than the traditional pre-internet style of advertising based off of content, but there has never been a major publisher willing to publicly run the experiment.

The New York Times tested this theory by cutting off all ad networks in Europe, and running only direct sold ads to its European visitors. The paper found that ad revenue increased significantly, and stayed elevated into 2019, bolstering the argument that surveillance-based advertising offers no advantage to publishers, and may in fact harm them.[6]

The Limits of Consent

While it is too soon to draw definitive conclusions about the GDPR, there is a tension between its concept of user consent and the reality of a surveillance economy that is worth examining in more detail.

A key assumption of the consent model is any user can choose to withhold consent from online services. But not all services are created equal—there are some that you really can't say no to.

Take the example of Facebook. Both landlords and employers in the United States have begun demanding to see Facebook accounts as a condition of housing or employment[7]. The United States Border Patrol has made a formal request to begin collecting social media to help vet people arriving in the country[8]. In both those contexts, not having a Facebook account might stand out too much to be a viable option. Many schools now communicate with parents via Facebook; Facebook groups are also the locus for political organizing and online activism across the political spectrum.

Analogous arguments can be made for social products offered by the other major tech companies. But if you can't afford to opt out, what does it mean to consent?

---

[6] Davies, Jessica. "After GDPR, the New York Times cut off ad exchanges in Europe—and kept growing ad revenue". Digiday (Jan 6, 2019)., https://digiday.com/media/gumgumtest-new-york-times-gdpr-cut-off-ad-exchanges-europe-ad-revenue/

[7] Dewey, Caitlin. "Creepy startup will help landlords, employers and online dates strip-mine intimate data from your Facebook page." Washington Post, (Jun 9, 2016). https://www.washingtonpost.com/news/the-intersect/wp/2016/06/09/creepy-startup-will-help-landlords-employers-and-online-dates-strip-mine-intimate-data-from-your-facebook-page/

[8] 81 FR 40892 https://www.federalregister.gov/documents/2016/06/23/2016-14848/agency-information-collection-activities-arrival-and-departure-record-forms-i-94-and-i-94w-and#h-11

Opting out can also be impossible because of how deeply the internet giants have embedded themselves in the fabric of the internet. For example, major media properties in the EU use a technology called ReCaptcha on their GDPR consent forms[9]. These forms must be completed before a user can access the website they are gathering consent for, but since the ReCaptcha service is run by Google, and the form cannot be submitted without completing the Google-generated challenge (which incidentally performs free image classification labor for the company), a user who refuses to give Google access to her browser will find herself denied access to a large portion of the internet.

While this specific example may change when it comes to the attention of an EU regulator, the broader issue remains. The sheer reach of the tech oligopoly makes it impossible to avoid using their services. When a company like Google controls the market-leading browser, mobile operating system, email service and analytics suite, exercises a monopoly over search in the EU, runs the largest ad network in Europe, and happens to own many of the undersea cables that connect Europe to the rest of the world[10], how do you possibly say 'no'?

Informed Consent

Beyond one's basic ability to consent, there is the question of what it means to give informed consent. Presumably we are not opting in or out of the services we use for capricious reasons, but because we can make a rational choice about what is in our interest.

In practice, however, obtaining this information is not possible, even assuming superhuman reserves of patience.

For example, anyone visiting the popular Tumblr blogging platform from a European IP address must first decide whether to share data with Tumblr's 201 advertising partners, and read five separate privacy policies from Tumblr's several web analytics providers.

Despite being a domain expert in the field, and spending an hour clicking into these policies, I am unable to communicate what it is that Tumblr is tracking, or what data of mine will be used for what purposes by their data partners (each of whom has its own voluminous terms of service). This opacity exists in part because the intermediaries have fought hard to keep their business practices and data sharing processes a secret, even in the teeth of strong European regulation.

Organizations like the Interactive Advertising Bureau Europe (IABE) defeat the spirit of the GDPR by bundling consent and requiring it across many ad-supported properties in Europe. If regulators block the bundling in its current incarnation, it will no doubt rise from the dead in a modified form, reflecting the undying spirit of surveillance advertising. But at no point will internet users have the information they

---

[9] The purpose of ReCaptcha is to prevent automated submissions, and ensure that a human being is filling out the form.

[10] Zimmer, Jameson. "Google Owns 63,605 Miles and 8.5% of Submarine Cables Worldwide." Broadband Now, (September 12, 2018) https://broadbandnow.com/report/google-content-providers-submarine-cable-ownership/

would need to make a truly informed choice (leaving aside the ridiculousness of requiring a legal education and two hours of sustained close reading in order to watch a cat video).

Consent in a world of inference

Finally, there is a sense in which machine learning and the power of predictive inference may be making the whole idea of consent irrelevant. At this point, companies have collected so much data about entire populations that they can simply make guesses about us, often with astonishing accuracy[11].

A useful analogy here is a jigsaw puzzle. If you give me a puzzle with one piece missing, I can still assemble it, reconstruct the contours of the missing piece by looking at the shape of the pieces around it and, if the piece is small compared to the whole, easily interpolate the missing part of the image.

This is exactly what computer systems do to us when we deny them our personal information. Experts have long known that it takes a very small amount of data to make reliable inferences about a person. Most people in the United States, for example, can be uniquely identified by just the combination of their date of birth, gender, and zip code[12].

But machine learning is honing this ability to fill in the blanks to surprising levels of accuracy, raising troubling questions about what it means to have *any* categories of protected data at all.

For example, imagine that an algorithm could inspect your online purchasing history and, with high confidence, infer that you suffer from an anxiety disorder. Ordinarily, this kind of sensitive medical information would be protected by HIPAA, but is the inference similarly protected? What if the algorithm is only reasonably certain? What if the algorithm knows that you're healthy now, but will suffer from such a disorder in the future?

The question is not hypothetical—a 2017 study[13] showed that a machine learning algorithm examining photos posted to the image-sharing site Instagram was able to detect signs of depression before it was diagnosed in the subjects, and outperformed medical doctors on the task.

The paradigm of automatic ownership of personal data does not mesh well with a world where such private data can not only interpolated and reconstructed, but independently discovered by an algorithm!

And if I can infer such important facts about your life by applying machine learning to public data, then I have deprived you of privacy just as effectively as I would have by direct eavesdropping.

---

[11] The line of argument in this section is adapted from the work of Dr. Zeynep Tufekci, UNC Chapel Hill. For example, "Think You're Discreet Online? Think Again" (April 21, 2019) https://www.nytimes.com/2019/04/21/opinion/computational-inference.html

[12] Sweeney, Latanya. (2000) "Simple Demographics Often Identify People Uniquely", Carnegie Mellon University, Data Privacy Working Paper. https://dataprivacylab.org/projects/identifiability/paper1.pdf

[13] Reece, Andrew and Danforth, Cristopher. (2017) "Instagram photos reveal predictive markers of depression" *EPJ Data Science* https://epjdatascience.springeropen.com/articles/10.1140/epjds/s13688-017-0110-z

In order to talk meaningfully about consent in online systems, the locus of regulation will need to expand beyond data collection, to cover how those data collections, and the algorithms trained on them, are used. But to do this, we will first need far greater visibility into the workings of surveillance-dependent tech companies than they have so far been willing to grant us.

As it stands, the consent framework exemplified in the GDPR is simply not adequate to safeguard privacy. As much as we would like to be the masters of our data, we are not. And the real masters aren't talking.

Goals for Privacy Regulation

Absent a clear understanding of how our data is being used, and the role it plays in surveillance-based business models, it is hard to lay out a specific regulatory program.

Nevertheless, there are some general goals we can pursue based on the experience of regulation attempts in Europe, and what we know about the surveillance economy.

Clarity

Privacy regulation should be understandable, both for users of the technology, and for the companies the regulations govern. Users especially should not be required to make complex and irrevocable decisions about privacy. To the extent possible, intuitions about privacy from the human world ("a casual conversation between friends is not recorded forever") should carry over into the digital world.

Privacy

At the risk of sounding tautological, privacy regulation should not punish people for seeking privacy. It should not be necessary to turn on invasive tracking technologies in one's browser in order to express the desire to not to be tracked.

Retention Limits on Behavioral Data

Knowing that we lack the capacity to keep data collections safe over time, we can reduce the potential impact of any breach by setting strict lifetimes for behavioral data.

Google has demonstrated the feasibility of this approach with their recent announcement that users will be able to set their account to automatically delete location data after three or 18 months[14]. This

[14] Monsees, David and McGriff, Marlo. "Introducing auto-delete controls for your Location History and activity data." (2019, May 1). https://www.blog.google/technology/safety-security/automatically-delete-data/

demonstrates that permanent retention of behavioral data is not critical to surveillance-based business models. Such limits should be enforced industrywide.

Moving to a norm where behavioral data is kept briefly instead of forever will mark a major step forward in data security, both reducing the time data is potentially exposed to attackers, and reducing the total volume of data that must be kept safe.

Time limits on behavioral data will also reduce consumers' perception that they are making irrevocable privacy commitments every time they try a new product or service.

Right To Download

The right to download is one of the most laudable features in the GDPR, and serves the important secondary purpose of educating the public about the extent of data collection.

This right should, however, be expanded to include the right to download, and correct, all information that third-party data brokers have provided about a user, in a spirit similar to the Fair Credit Reporting Act.

Fairness

Tech startups in the highly regulated areas of health, finance and banking should be required to compete on the same regulatory footing as established businesses in those areas. In particular, they should not be allowed to do an end run around existing data privacy laws by using machine learning and algorithmic inference.

For example, the use of a machine learning algorithm should not allow a loan company to evade consumer protections against discrimination in fair lending laws.

(For a fuller discussion of this point, see the addendum on machine learning at the end of this document).

Positive Regulation

While the above suggestions seek to impose limits and restrictions, there is an important way that privacy regulation can create new ground for innovation.

What is missing from the regulatory landscape is a legal mechanism for making credible and binding promises to users about privacy practices.

Today, internet startups in the U.S. who want to compete on privacy have no mechanism to signal their commitment to users other than making promises through their terms of service (which usually include a standard legal clause that they may change at any time).

Except in the case of the most egregious violations, which sometimes attract the attention of the Federal Trade Commission, these terms of service carry little weight.

As the owner of a company that markets itself to privacy-conscious people, I would derive enormous benefit from a legal framework that allowed me to make binding privacy promises (for example, a pledge that there is no third-party tracking on my website), and imposed stiff fines on my company if I violated these guarantees (including criminal liability in the case of outright fraud).

Such a legal mechanism would not only enable competition around privacy-enhancing features, but it would also give future regulators a clearer idea of how much value consumers place on data privacy. It is possible that the tech giants are right, and people want services for free, no matter the privacy cost. It is also possible that people value privacy, and will pay extra for it, just like many people now pay a premium for organic fruit. The experiment is easy to run—but it requires a modest foundation in law.

Academic research in computer science is full of fascinating ideas that could serve as the seed for business built around user privacy. Results in fields like homeomorphic encryption, differential privacy, privacy-preserving machine learning, and zero-knowledge proofs all await a clever entrepreneur who can incorporate them into a useful product or service. It is very hard to compete against companies like Amazon or Facebook on price, but it is not hard to beat them on privacy. With a minimum of regulatory scaffolding, we might see a welcome new burst of innovation.

Preserving Liberty

The final, and paramount goal, of privacy regulation should be to preserve our liberty.

There is no clearer warning of the danger of building up an infrastructure of surveillance than what is happening today in China's Xinjiang Uygur Autonomous Region. Claiming to be concerned about the possible radicalization of a Muslim minority, Chinese authorities have imposed a regime of total surveillance over a population of twenty-five million people.

As recent reporting by Human Rights Watch has shown, a computer system called the Integrated Joint Operations Platform (IJOP) monitors the location and movement of all people in the province (based on phone data), as well as their gas and electricity consumption, which apps they use, where they worship, who they communicate with, and how they spend their money. This surveillance information is fed into machine learning models that can bin people into one of thirty-six suspect categories, bringing them to the closer attention of the police[15]. Never before has a government had the technical means to implement this level of surveillance across an entire population. And they are doing it with the same off-the-shelf commercial technologies we use in America to get people to click on ads.

---

[15] Human Rights Watch, "China's Algorithms of Repression", (May 1 2019) https://www.hrw.org/report/2019/05/01/chinas-algorithms-repression/reverse-engineering-xinjiang-police-mass-surveillance

The latent potential of the surveillance economy as a toolkit for despotism cannot be exaggerated. The monitoring tools we see in repressive regimes are not 'dual use' technologies—they are single use technologies, working as designed, except for a different master.

For sixty years, we have called the threat of totalitarian surveillance 'Orwellian', but the word no longer fits the threat. The better word now may be 'Californian'. A truly sophisticated system of social control, of the kind being pioneered in China, will not compel obedience, but nudge people towards it. Rather than censoring or punishing those who dissent, it will simply make sure their voices are not heard. It will reward complacent behavior, and sideline troublemakers. It's even possible that, judiciously wielded, such a system of social control might enjoy wide public support in our own country.

But I hope you will agree with me that such a future would be profoundly un-American.

There is no deep reason that weds the commercial internet to a business model of blanket surveillance. The spirit of innovation is not dead in Silicon Valley, and there are other ways we can grow our digital economy that will maintain our lead in information technology, while also safeguarding our liberty. Just like the creation of the internet itself, the effort to put it on a safer foundation will require a combination of research, entrepreneurial drive and timely, enlightened regulation. But we did it before, and there's no reason to think we can't do it again.

Addendum: Machine Learning and Privacy

Machine learning is a mathematical technique for training computer systems to make accurate predictions from a large corpus of training data, with a degree of accuracy that in some domains can mimic human cognition.

For example, machine learning algorithms trained on a sufficiently large data set can learn to identify objects in photographs with a high degree of accuracy, transcribe spoken language to text, translate texts between languages, or flag anomalous behavior on a surveillance videotape.

The mathematical techniques underpinning machine learning, like convolutional neural networks (CNN), have been well-known since before the revolution in machine learning that took place beginning in 2012. What enabled the key breakthrough in machine learning was the arrival of truly large collections of data, along with concomitant computing power, allowing these techniques to finally demonstrate their full potential.

It takes data sets of millions or billions of items, along with considerable computing power, to get adequate results from a machine learning algorithms Before the advent of the surveillance economy, we simply did not realize the power of these techniques when applied at scale.

Because machine learning has a voracious appetite for data and computing power, it contributes both to the centralizing tendency that has consolidated the tech industry, and to the pressure companies face to maximize the collection of user data.

Machine learning models poses some unique problems in privacy regulation because of the way they can obscure the links between the data used to train them and their ultimate behavior.

A key feature of machine learning is that it occurs in separable phases. An initial training phase consists of running a learning algorithm on a large collection of labeled data (a time and computation-intensive process). This model can then be deployed in an exploitation phase, which requires far fewer resources.

Once the training phase is complete, the data used to train the model is no longer required and can conceivably be thrown away.

The two phases of training and exploitation can occur far away from each other both in space and time. The legal status of models trained on personal data under privacy laws like the GDPR, or whether data transfer laws apply to moving a trained model across jurisdictions, is not clear.

Inspecting a trained model reveals nothing about the data that went into it. To a human inspecting it, the model consists of millions and millions of numeric weights that have no obvious meaning, or relationship to human categories of thought. One cannot examine an image recognition model, for example, and point to the numbers that encode ‘apple’.

The training process behaves as a kind of one-way function. It is not possible to run a trained model backwards to reconstruct the input data; nor is it possible to "untrain" a model so that it will forget a specific part of its input.

Machine learning algorithms are best understood as inference engines. They find structure and excel at making inferences from data that can sometimes be surprising even to people familiar with the technology. This ability to see patterns that humans don't notice has led to interest in using machine learning algorithms in medical diagnosis, evaluating insurance risk, assigning credit scores, stock trading, and other fields that currently rely on expert human analysis.

The opacity of machine learning models, combined with this capacity for inference, also make them an ideal technology for circumventing legal protections on data use. In this spirit, I have previously referred to machine learning as "money laundering for bias". Whatever latent biases are in the training data, whether or not they are apparent to humans, and whether or not attempts are made to remove them from the data set, will be reflected in the behavior of the model.

A final feature of machine learning is that it is curiously vulnerable to adversarial inputs. For example, an image classifier that correctly identifies a picture of a horse might reclassify the same image as an apple, sailboat or any other object of an attacker's choosing if they can manipulate even one pixel in the image[16]. Changes in input data not noticeable to a human observer will be sufficient to persuade the model. Recent research suggests that this property is an inherent and ineradicable feature of any machine learning system that uses current approaches[17].

In brief, machine learning is effective, has an enormous appetite for data, requires large computational resources, makes decisions that resist analysis, excels at finding latent structure in data, obscures the link between source data and outcomes, defies many human intuitions, and is readily fooled by a knowledgeable adversary.

---

[16] Su, Jiawei, Vargas, Danilo, and Kouichi, Sakurai. "One Pixel Attack for Fooling Deep Neural Networks" (2017, Oct 24). https://arxiv.org/pdf/1710.08864.pdf

[17] Wang, Xianmin, Li, Jing, Kuang, Xioahui, Tan, Yu-an. Journal of Parallel and Distributed Computing
(August 2019) "*The security of machine learning in an adversarial setting: A survey"*