Privacy Preserving Machine Learning: A Human Imperative?

Thomas Strohmer
Department of Mathematics
University of California, Davis

The Mathematics of Deep Learning and Data Science
Isaac Newton Institute, May 2019

Acknowledgements

This work is sponsored by the NSF-DMS and the National Geospatial-Intelligence Agency.


A Tale about
William and Kate

Overview

Powered by machine learning, a new economic system is emerging that threatens our social fabric.

This requires a new paradigm what data is and consequently how we go about AI and machine learning.

Agenda:

- Surveillance capitalism
- Dialectic dynamics in surveillance capitalism
- Need of new paradigm for understanding what data is
- Privacy-preserving machine learning

There is no "I" in AI

Current AI algorithms are powerful, but certainly not intelligent.

AI pioneer Judea Pearls:
"All the impressive achievements of deep learning amount to just curve fitting!"

Are we just seeing a
Clever Hans Effect?

Augmented Intelligence instead of Artificial Intelligence

The problem is not AI (by itself)

Potential dangers of AI:

- Massive job losses – actual challenge
- Robots take over – nonsense

Elon Musk warns against “summoning the demon”, “an immortal dictator from which we can never escape.”

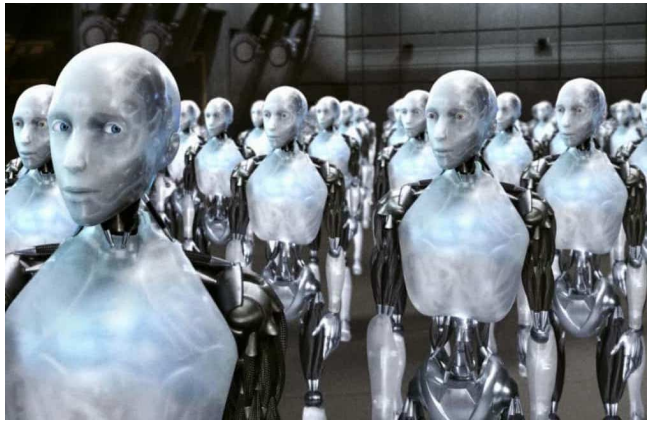

But there is a much more imminent, pervasive, and real threat.

Surveillance Capitalism

Surveillance Capitalism is the foundational framework for a new surveillance-based economic order.

Surveillance Capitalism: [S. Zuboff]

"The unilateral claiming of private human experience as free raw material for translation into behavioral data. "

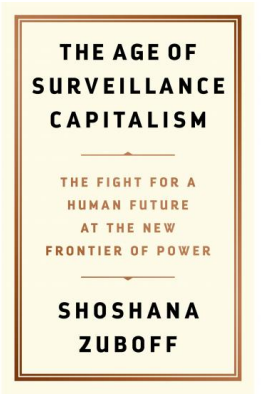

Surveillance Capitalism

Zuboff: “By providing free services that billions of people cheerfully use, it enables the providers of those services to monitor the behaviour of those users in astonishing detail - often without their explicit consent.”

The so extracted data are then fed into machine learning algorithms and packaged as prediction products and sold into behavioral futures markets.

Once we searched Google. Now it searches us

- Google's Street View mapping project scooped up passwords, e-mail and other personal information from people. According to Google, it was done unintentionally ...
- Gmail reads your email
- Google used secret code to bypass Safari's antitracking security setting
- Google can track you even if you turn off location services, stop using apps and remove your SIM card from your device. It does not allow Android users to opt out.
- Google sells this info to third parties. Eg., if you are near a specific store, that store can send you targeted advertising.

"We continue to use your data to improve the Google experience ..."

Facebook is a serial violator of privacy

- Without users' consent, Facebook has been sharing their data with more than 150 businesses, including Amazon, Microsoft, Netflix, Spotify, ...
- Facebook pretended to apply Europe's new privacy laws to all its users worldwide and then secretly switched all non-European users to the spineless US privacy law.
- Apps are sharing sensitive data with Facebook without informing users; and even when users are not logged in through Facebook, or do not have a Facebook account.
- Facebook wants banks to hand over their customers' sensitive financial data to offer better service to its users – *give us your data, we give you our users*.

The combined data from all the different apps paint a detailed and intimate picture of people's activities, interests, behaviors.

The Dead Don’t Die

Verizon used “zombie cookies” to monitor and share customers’ habits.
Zombie cookies reappear even if users clear them from their web browsers.

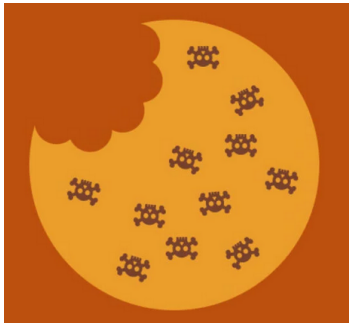

The Dead Don’t Die

Verizon used “zombie cookies” to monitor and share customers’ habits.
Zombie cookies reappear even if users clear them from their web browsers.

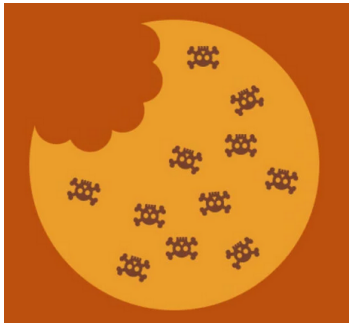

According to Verizon, people should not worry,

> *“because it is unlikely that websites and ad entities will attempt to build customer profiles.”*

The Dead Don't Die

Verizon used "zombie cookies" to monitor and share customers' habits.
Zombie cookies reappear even if users clear them from their web browsers.

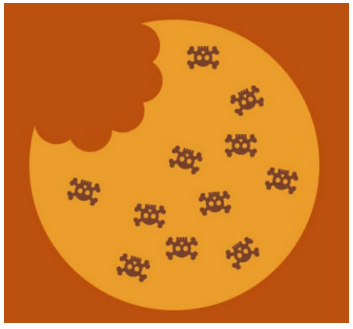

According to Verizon, people should not worry,

> *"because it is unlikely that websites and ad entities will attempt to build customer profiles."*
>
> *... the zombie cookie "has a benefit for customers, in case they mistakenly deleted it."*

I spy with my little eye ...

Companies like IBM, the Weather Channel and Foursquare obscure tracking companies (Groundtruth, Safegraph, ...), who pay to include their tracking code in common apps.

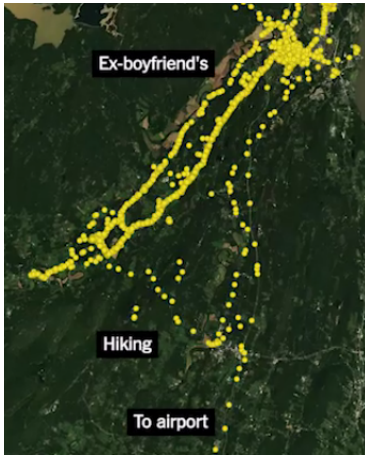

Hertz installs cameras and microphones in rental cars. But according to the company, it does not plan to use them ...
Hertz: “We do not have adequate bandwidth capabilities to the car to support streaming video at this time.”

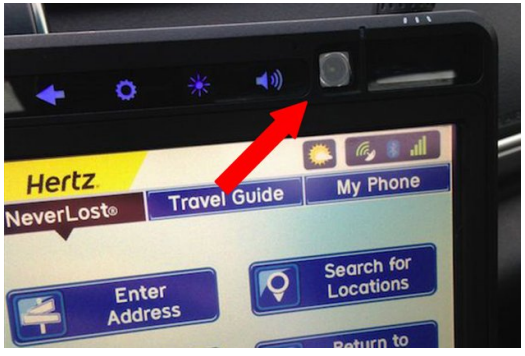

This is just for you own good ...

Companies claim this is like a doctor and patient relationship: You share information and receive personalized service.

But this is fundamentally different: Doctor-patient relationship is one of mutual dependence, symmetry, and rule of law.

It is not about what you post on Facebook or the search terms you are entering in Google.

This is just for you own good ...

Companies claim this is like a doctor and patient relationship: You share information and receive personalized service.

But this is fundamentally different: Doctor-patient relationship is one of mutual dependence, symmetry, and rule of law.

It is not about what you post on Facebook or the search terms you are entering in Google.

It is about the content of your emails, tracking your location, reading your phone contacts, the pictures you take, travels you plan, other signals of online behavior, ... – the extraction and sharing of which we never consented to and which has nothing to do with service improvement.

Surveillance capitalism

Industrial capitalism depended upon the exploitation and control of nature – with catastrophic consequences.

Surveillance capitalism depends instead upon the exploitation and control of human nature.

Private, human experience is traded and exchanged in a new kind of marketplace that is founded and operated by surveillance capitalism.

Our means of social participation have been conflated with the means through which surveillance capitalists collect their data and seek to modify our behavior.

From monitoring to control

The first phase of surveillance capitalism depended upon data extracted at scale from the Internet to produce relevant online ads.

Companies actually already move from behaviour surveillance to behaviour modification.

This requires up-close monitoring:

- Smart Home
- Wearables
- Internet-of-Things

Smart Home = Surveillance Home

Smart TVs track not only what you watch, but have built-in microphones and cameras ~~to record you~~ for personalized service.

Samsung about its Smart TVs: *"Please be aware that if your spoken words include personal or other sensitive information, that information will be among the data transmitted to a third party through your use of Voice Recognition".*

But Samsung disclaims any responsibility for the policies of third-party firms and states "You should review the privacy statements applicable to the third-party services you use".

Vizio's CTO: *"We charge a premium on 'dumb' TVs to make up for the money we lose by not spying on you."*
[Interview at CES, 2019]

Smart Home = Surveillance Home

If you install a Nest thermostat, you should review a minimum of 1000 privacy contracts, because all collected data is sent to third parties which is sent to third parties, which is sent ... .

If you do not agree to each policy at each stage, then you lose the functionality of the product and the reason that you bought it in the first place.

"Nest's on-device microphone was never intended to be a secret and should have been listed in the tech specs ... we just forgot to do so."
[Google's Nicol Addison, Feb. 2019]

Hey fridge, we're out of milk!

Roomba vacuum with its built-in cameras produces detailed maps of user's floor plans and wants to sell them to other companies such as Google, Amazon, ...

Your smart refrigerator is selling your grocery shopping habits to third parties.

Other items that spy on you in your home

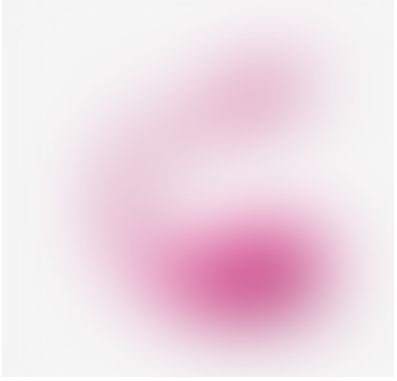

Alexa and friends

Google's Home and Amazon's Alexa, disguised as engines of "personalization," operate as complex supply chains for continuous automatic extraction of behavioral data from private human experience (that is not required for product and service improvement!)

Amazon, has a patent applications [Nr. 20190080685] for a voice sniffing algorithm that detects keywords ("like", 'love",...) via smart home devices (Alexa, door bell, thermostat, ...) to which Amazon and other companies then can respond with product and service offers.

Amazon was issued a patent [10,276,188] that would allow Alexa to decipher a user's physical characteristics and emotional state based on their voice.

But we have Privacy Policies, no?

Example: Android-based diabetes apps: [Blenner et al. 2016]

- Just downloading the app automatically authorizes collection and modification of sensitive information.
- These apps access your photos and videos
- These apps read your contact lists, modify your contacts, and activate your microphone to record your speech

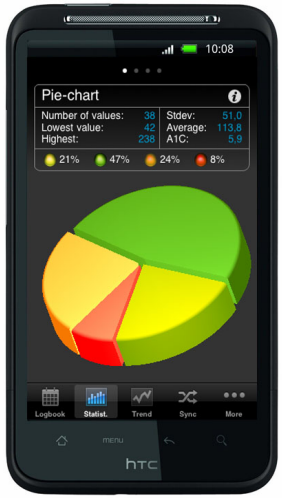

You can run, but you can’t hide

Luckily, many apps have privacy policies, so that should help ...

Example: Fitness trackers

- Among those apps **without** privacy policy, 76% shared sensitive information with third parties.

You can run, but you can't hide

Luckily, many apps have privacy policies, so that should help ...

Example: Fitness trackers

- Among those apps **without** privacy policy, 76% shared sensitive information with third parties.
- Among those apps **with** privacy policy, 79% shared sensitive information with third parties.

You Snooze, You Lose

Breathing machines purchased by people with sleep apnea are secretly sending usage data to health insurers, where the information can be used to justify reduced insurance payments. Data are sent, even if you opt out. [ProPublica, Nov. 2018]

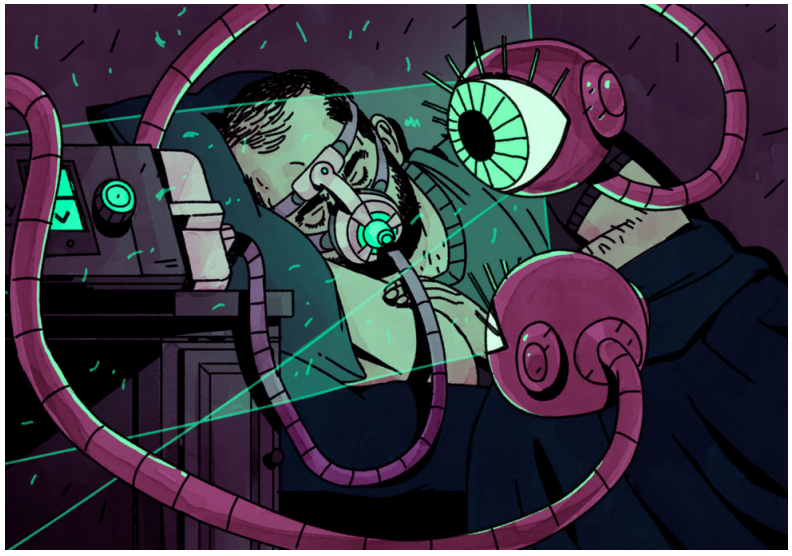

Your private medical data is for sale

In 2013 the UK National Health Services planned to sell patient data to drug and insurance firms. Patients could not opt out.

After major complains, the program was modified so that patients could opt out. Eventually the whole program was cancelled (for now).

Your private medical data is for sale

In 2013 the UK National Health Services planned to sell patient data to drug and insurance firms. Patients could not opt out.

After major complains, the program was modified so that patients could opt out. Eventually the whole program was cancelled (for now).

Major misconception:
Patients should not have to opt out in the first place, since medical data are among the most personal data of an individual.

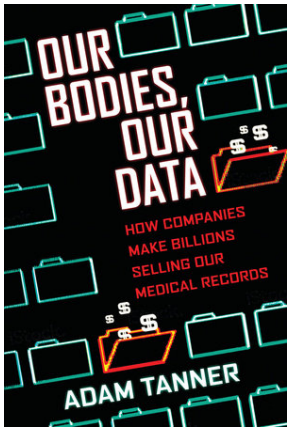

The legal right of businesses to harvest and sell the information of individual patients without their permission has been upheld by the US supreme court. [564 U.S. 552 (2011)]

Whatcha gonna do, when they come for you...

Even if you deny consent to “your” data being used, an organization can use data about other people to make statistical extrapolations that affect you.

In the US, some judges use algorithms to assess a criminal defendant’s likelihood to re-offend. COMPAS attempts to predict the likelihood that a person will commit future crimes.

- Black defendants were often predicted to be at a higher risk of recidivism than they actually were.
- White defendants were often predicted to be less risky than they were.

How do they get away with it?

Google and co. often operate in a legal vacuum, just grab personal information completely unrelated to the intended service and fail to identify its actual commercial intent.

They pursue data gathering via misdirection, euphemism, and obfuscation.

They claim that the new economic practices are an inevitable consequence of digital technology.

Companies care about privacy ...

The privacy of our customers' personal information is very important to us ... but not as important as our profits!

The surveillance capitalists use their privacy power to prevent us from being able to inspect or control their behaviour.

Zuckerberg voted down a shareholder proposal for more accountability and transparency regarding privacy.
Google's Page and Brin have voted down a similar proposal.

From behavior surveillance to behavior modification

The dynamics of surveillance capitalism have created powerful economic imperatives that are driving these companies to produce better and better behavioral-prediction products.

This requires not only amassing huge volumes of data, but actually intervening in our behavior.

Surveillance capitalists learn to tune, herd, and condition our behavior with subtle cues, rewards, and punishments that direct us toward their most profitable outcomes.

From behavior surveillance to behavior modification

Companies have always tried to shape customer behavior through priming, suggestion, and social comparison.

What distinguishes today's efforts is that not only do they extend beyond advertising, but they employ a ubiquitous digital architecture that automates the continuous comprehensive monitoring and shaping of human behavior with unprecedented accuracy, intimacy, and effectiveness.

Digital architecture allows one to do this at unprecented scale.

From behavior surveillance to behavior modification

Google's Pokémon Go was marketed a harmless foray into the world of augmented reality.

Virtual figures of Pokémon characters are placed around your area and you have to physically go there to "capture" them and get points.

In truth it is designed to guide its users towards commercial opportunities. Businesses paid to have Pokémon Go players directed to their locations.

We are not the players in Pokémon – we are the pawns.

Facebook's secret mood manipulation experiment

Facebook's data scientists skewed the News Feeds of half a million users, removing either all of the positive posts or all of the negative posts, to manipulate the users' emotional state.

"We do research to improve our services ... and to make the content as engaging as possible."

Facebook's secret mood manipulation experiment

Facebook's data scientists skewed the News Feeds of half a million users, removing either all of the positive posts or all of the negative posts, to manipulate the users' emotional state.

"We do research to improve our services ... and to make the content as engaging as possible."

Facebook promotes to other companies an AI-powered "loyalty prediction" service called FBLearner Flow that identifies "individuals who are 'at risk' of shifting their brand allegiance" and prompts advertisers to intervene swiftly.

From observation and prediction to modification

Pentland proposes the use of “unobtrusive wearable sensors” called sociometers that would help managers “infer relationships between colleagues” and coordinate collaboration between groups/individuals.

“By using these data to build a predictive, computational theory of human behavior we can hope to engineer better social systems.”

From observation and prediction to modification

The whole digital architecture becomes a global means of behavioral modification that is used to tune and herd populations in the service of advertisers.

“The internet will disappear ... It will be part of your presence all the time.” [Schmidt, Google, 2016]

> *“Conditioning at scale is essential to the new science of massively engineered human behavior.”*

Shaping the perfect citizen

China recently introduced a Social Credit Score

"It allows the trustworthy to roam freely under heaven while making it hard for the discredited to take a single step."

The system will aggregate all available data, taking into account citizens' debt, what books they read, what they buy, how long they play video games, how patriotic they are, ...

Shaping the perfect world

Facebook wants to provide a future that works for everyone and fulfills personal, emotional, and spiritual needs for purpose and hope, and moral validation. [Zuckerberg, 2017]

“I actually think most people don’t want Google to answer their questions ... They want Google to tell them what they should be doing next,” [Schmidt, Google, Wall Street Journal, 2010].

“We can know if you shouldn’t be driving, and we can just shut your car down ... we tell the TV to shut off and make you get some sleep, or the chair to start shaking because you shouldn’t be sitting so long.”

Where do we go from here?

We need an additional perspective to analyze the workings of and possible defense strategies against surveillance capitalism, while still being able to enjoy the benefits of the digital revolution and AI.

A dialectical worldview – Aporias

Conflict management:
(Using work by G. Schwarz and H. Pietschmann)

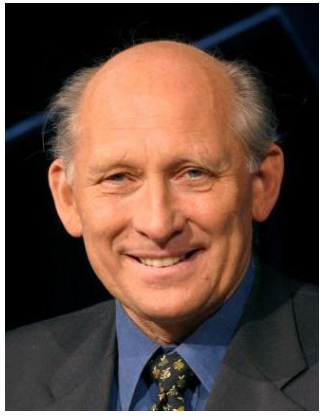

A dialectical worldview – Aporias

Conflict management:
(Using work by G. Schwarz and H. Pietschmann)

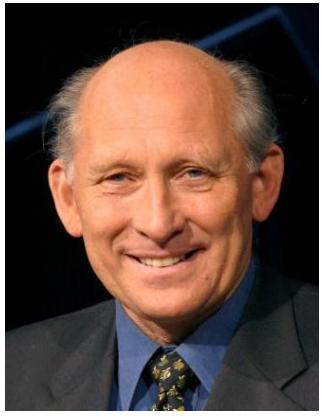

Some conflicts have a simple “right or wrong” solution.

“England has the strongest football league in the world.”

But other, very important, conflicts do not fit this logical scheme

Aporias

What is an Aporia? (Going back to Plato and Aristotle)

Aporias

What is an Aporia? (Going back to Plato and Aristotle)
An aporia is an insoluble contradiction, a paradox.

Aporia:

1. There are two contradicting statements
2. Both statements are true (there is no right or wrong)
3. The two contradicting statements induce each other

Understanding aporias is a key tool in understanding and resolving fundamental conflicts.

"The opposite of a fact is a falsehood, but the opposite of one profound truth may very well be another profound truth!"

– Niels Bohr

Example of an aporia: Freedom $\longleftrightarrow$ Order

Order preserves Freedom $\longleftrightarrow$ Order destroys Freedom

Both statements contradict each other, both statements are true, and both statements depend upon each other.

"Without order, we cannot live together."
"Without freedom, life is not worth living."

It also works in the reverse direction:

Freedom preserves Order $\longleftrightarrow$ Freedom destroys Order

Such an aporetic conflict has no right-or-wrong solution.
It cannot truly by solved by compromise, but only by consense.
In the above case, the consense strives for Autonomy.

Autonomy = "Freedom gives itself Order"

Aporias

Autonomy = “Freedom gives itself Order”

In an aporetic conflict if one side wins, both sides lose!
In a consense both sides win (in contrast to a compromise).
In conflict management, this is called a paradoxic intervention.

Such a solution of an aporetic conflict is not permanent, and needs to be “relived” every day.

Hegel: Something is living insofar as it contains a contradiction.

Recall Hegel’s dialectics (thesis-antithesis-synthesis).

Aporias and their “shadows”

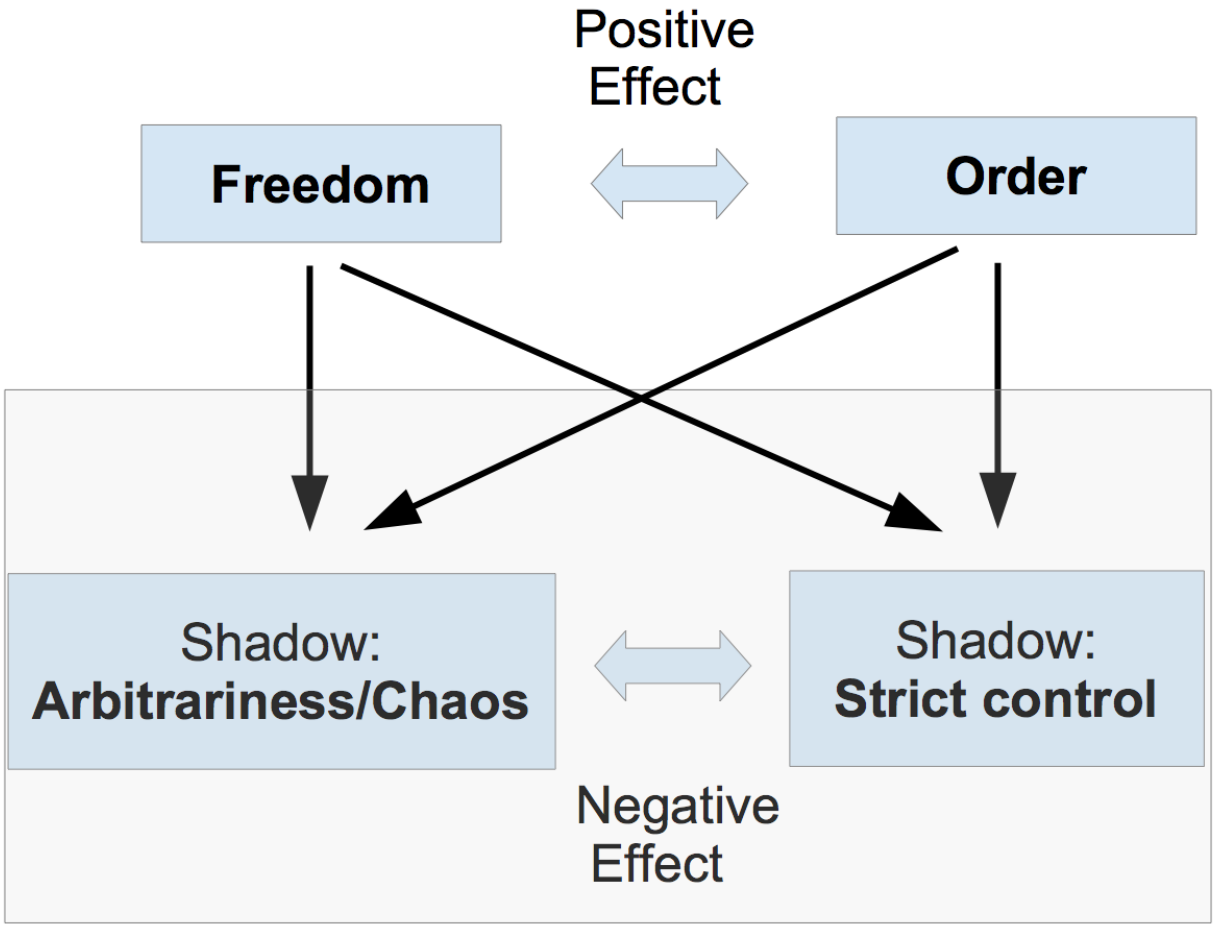

Extremizing freedom leads to chaos and arbitrariness
Extremizing order leads to strict control and regulating everything

Aporias and their “shadows”

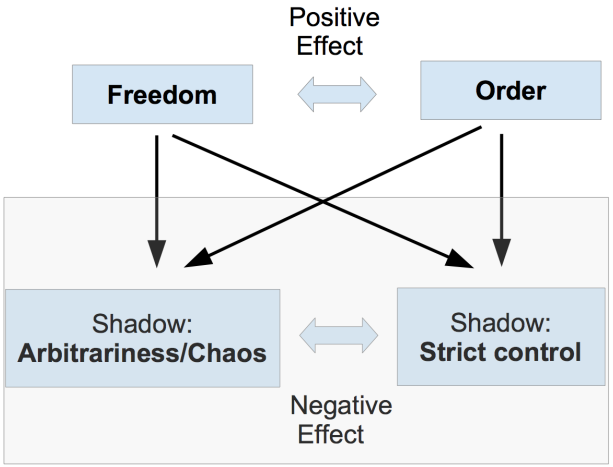

Those who favor Order fight against Chaos, and thereby fall into Control, the “shadow” of Order.

Those who favor Freedom fight against Control, and thereby fall into Chaos or Abtritrariness, the “shadow” of Freedom.
This can turn the goal of freedom into chaos, and as a result into strict control.

Aporias and their “shadows”

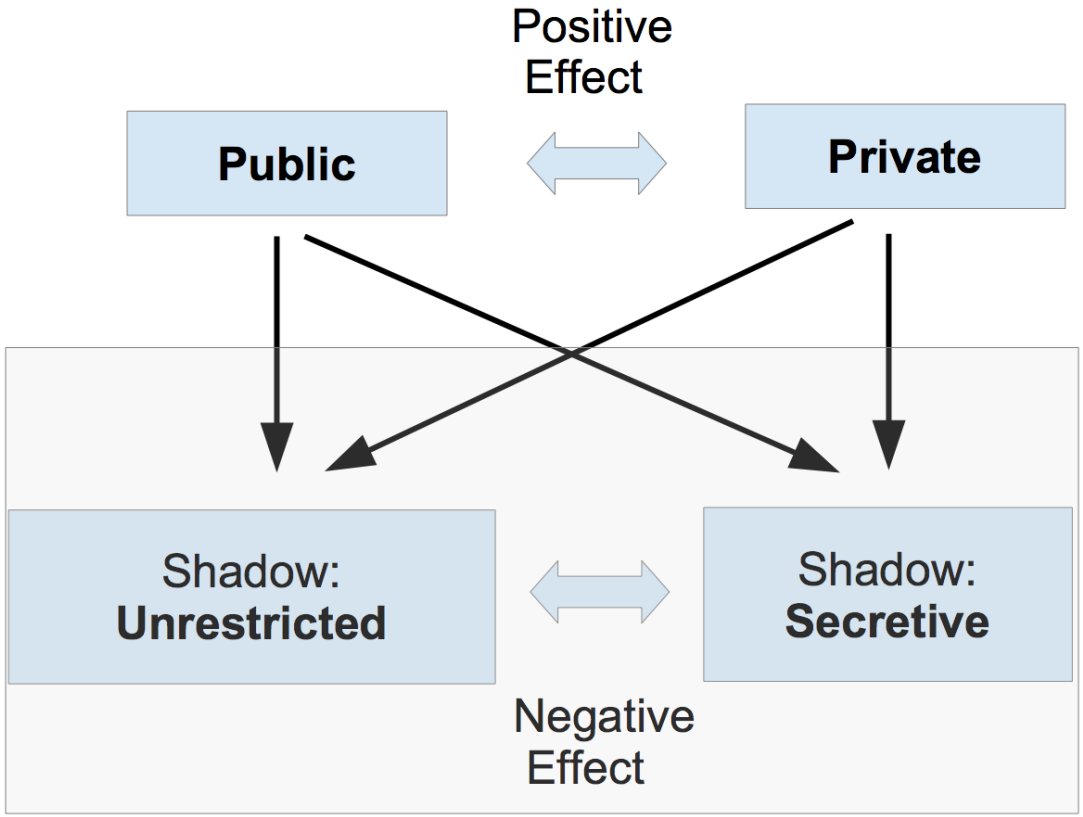

Those who favor “public”, fight against “secretive”, and thereby fall into the “unrestricted” shadow

Those who favor “private”, fight against “unrestricted” and thereby fall into “secretive” shadow

A Faustian Pact

China's Social Credit Score *"allows the trustworthy to roam freely under heaven while making it hard for the discredited to take a single step."*

Facebook wants to provide a future that works for everyone and fulfills personal, emotional, and spiritual needs for purpose and hope, and moral validation.

Google's planned "Smart City" promises to solve many problems, such as traffic congestion, rising housing prices, and environmental pollution, based on a private, fully surveilled city, so that we have more freedom to do what we want ...

They promise freedom in exchange for handing them control.

From Utopia to Dystopia

To realize their utopian vision, Zuckerberg & Co. must delete all boundaries in the service of their economic imperatives.

Instead of “public”, they embrace “unrestricted”.

They proclaim the freedom of the internet to ensure their continued data harvesting, but in fact they pursue the arbitrariness of the internet (absence of rules).

Surveillance capitalism’s means of behavioral modification at scale erodes democracy from the inside and from the outside.

Democracy and the digital revolution

If decentralization is one of the characteristica of the digital revolution, then having no global rules over who owns data and no oversight over surveillance companies, can have the opposite effect.

Freedom turns into its shadow: arbitrariness

The digital revolution, instead of increasing decentralization and democratization, can lead to a concentration of power stronger than ever seen before.

(Social) trust is based on freedom.
Algorithmic certainty is based on control.

How to take charge again?

Going into hiding?

Just don't use a smart phone, don't use the Internet. Don't use credit cards, don't fly, don't rent a car, ...

Recall: "Secrecy" is the enemy of "private".

Thus, going into hiding cannot be a solution.

Surrender?

> *"If you want to use the benefits and conveniences of the Internet and the digital age, you have to live with surveillance capitalism."*

This fatalistic approach fails to recognize that surveillance capitalism depends on the digital revolution, but the digital revolution does not need surveillance capitalism.

We need a Bill of Data Rights

Martin Tisne:
“It’s time for a Bill of Data Rights"

- A new paradigm for understanding what data is–and what rights pertain to it–is urgently needed.
- The real questions are questions about how data shapes society and individuals.
- Therefore, this also must inform our thinking about machine learning
- “The mathematics of data science” does not live in a societal vacuum

Data rights

Data rights should account for the fact that privacy is not a reactive right to shield oneself from society.
It is about freedom to develop the *'self'* away from commerce and away from governmental control.
A Bill of Data Rights should include rights like these:

1. The right of the people to be secure against unreasonable surveillance shall not be violated.
2. No person shall have his or her behavior surreptitiously manipulated.
3. No person shall be unfairly discriminated against on the basis of data.

Enforcing these rights requires foremost new laws.
But we also need mathematics, statistics, computer science.
E.g. ad (3), we need to understand how algorithms work;
ad (1), we need to boost privacy-preserving machine learning.

Data rights, data wrongs

- The very act of opening an online article on an electronic device creates data–an entry in your browser's history, cookies the website sent to your browser, an entry in the website's server log to record a visit from your IP address.
- It is essentially impossible to do anything online without leaving a "digital trace" behind.
- These digital traces cannot be owned.
- This is why data ownership is insufficient (and this is why Pentland's "New Deal on Data" is fundamentally flawed.)

The trialectic nature of data

Trialectics: Instead of two contradictions (dialectics), we face three interrelated contradictions and associated aporias.

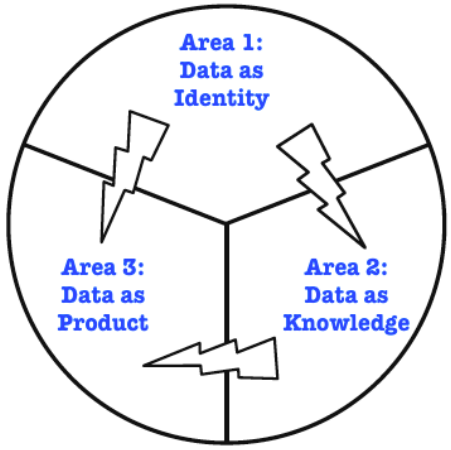

- Area 1: Data as identity is the synthesis of data as knowledge and data as product
- Area 2: Data as knowledge is the synthesis of data as product and data as identity
- Area 3: Data as product is the synthesis of data as identity and data as knowledge

How to take charge again?

1. We need public outrage – stop being *comfortably numb*.
2. We need to muster the resources of our democratic institutions in the form of law and regulation.
3. For this we need a new paradigm for understanding what data is and what rights pertain to it. This needs to be done with the aporia public–private in mind.
4. We need to develop competitive concepts.
   This is a great opportunity for an alliance of companies to found an alternative ecosystem.
5. If we do not want to give up on the potential benefits of AI (not going into hiding), we need to invest massively in privacy preserving machine learning – with new data paradigm in mind.

What can we scientists do?

- No fatalistic attitude – “there is nothing we can do ...”
- We do not want to fall into the secrecy mode.
- Surveillance capitalism relies on algorithms and sensors, machine intelligence and platforms, but it is not the same as any of those.
- Surveillance capitalism needs machine learning, but machine learning does very well without surveillance capitalism.
- AI offers many benefits, it needs data to work.
  New paradigm for understanding what data is
  (we may need to adopt a trialectic viewpoint of data)
- Privacy preserving machine learning

Privacy preserving machine learning

- Privacy preserving machine learning is not the silver bullet that makes all problems go away.
- It is one important necessary (but not sufficient!) step toward balancing the powers.
- What privacy preserving machine learning is supposed to accomplish, must be understood in the bigger context of surveillance capitalism and its consequences
- For example, an incremental loss in privacy, taken by itself, may do little harm (like, if I pollute the environment a little, it does not do any noticable harm)
- But in aggregate, it can have calamitous collective consequences and cause fundamental damage to the social fabric.

Privacy-preserving machine learning

Differential privacy:

An algorithm is differentially private if an observer seeing its output cannot tell if a particular individual's information was used in the computation Captures the increased risk to one's privacy incurred by participating in a database. But note that miniscule loss in privacy, taken by itself, may do little harm, however in aggregate, it may be very damaging.

Machine learning on encrypted data:

Homomorphic encryption allows for computations to be done on encrypted datay. The results can be revealed only by the owner of the secret key. Unfortunately at this point, it is highly impractical and computationally not feasible.

Privacy-preserving machine learning

Distributed/federated learning:

Enables devices to collaboratively learn a shared prediction model while keeping all the training data on device, decoupling the ability to do machine learning from the need to store the data in the cloud.

On-device machine learning:

Offers benefits such as increased privacy and security, low latency, and more autonomy. A major challenge is that limited memory, weak processors, and scarce energy supply.

What Happens on the Edge, Stays on the Edge

Another paradox: The emergence of the Internet-of-Things can be used to improve our privacy.
Problem of the IoT: Too many devices, too little bandwidth.
Moreover, edge devices often have little memory and little computing power.
Many companies invest in on-device machine learning.
A good step towards keeping AI private.

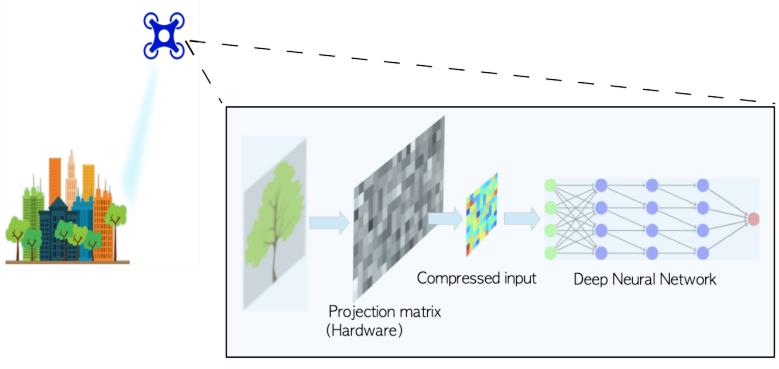

What Happens on the Edge, Stays on the Edge: Toward Compressive Deep Learning [T.S. and Yang Li, 2019]

Conclusion

Powered by AI, surveilleance capitalism emerges as a new economic system that threatens our social fabric.

We need to respond to this threat without sacrificing the benefits of the digital revolution $\rightarrow$ understand aporias!

A logical worldview clashes with a dialectic worldview.

Need a new paradigm for understanding what data is and what rights pertain to it.

Mathematics of data science does not live in a societal vacuum.

Conclusion

Powered by AI, surveilleance capitalism emerges as a new economic system that threatens our social fabric.

We need to respond to this threat without sacrificing the benefits of the digital revolution $\rightarrow$ understand aporias!

A logical worldview clashes with a dialectic worldview.

Need a new paradigm for understanding what data is and what rights pertain to it.

Mathematics of data science does not live in a societal vacuum.

"Never waste the opportunity offered by a good crisis."