Emerging Issues in Artificial Intelligence and Machine Learning

ARTIFICIAL INTELLIGENCE, MACHINE LEARNING AND ROBOTICS

Susan Fahringer

Co-Chair Artificial Intelligence, Machine Learning & Robotics Practice

Dean Harvey

Co-Chair Artificial Intelligence, Machine Learning & Robotics Practice

Agenda

Overview of Machine Learning Technology

Emerging Issues

- Algorithmic Discrimination
- Privacy
- Automated Decision-Making
- Antitrust Issues: Algorithmic Collusion & Price-Fixing
- Product Liability
- Demands for Access by Law Enforcement & Civil Litigants

Questions

Overview of AI and Machine Learning

ARTIFICIAL INTELLIGENCE

Artificial Intelligence is the simulation of human intelligence in machines

Traditional AI systems were programmed to attempt to simulate human intelligence (e.g., IBM’s Deep Blue)

MACHINE LEARNING

ML is a subset of AI involving a system that learns from data without rules-based programming (e.g., Google Deepmind’s AlphaGo)

Machine Learning

How does it differ from traditional software?

Key Distinctions

- Traditional software requires hand-coding with specific instructions to complete a task
- An ML system learns to recognize patterns and make predictions using large amounts of data

EXAMPLE:

- Spam the old way: “if the email contains the word ‘viagra,’ then …”
- Spam the new way: ML system learns from training data to identify if email is spam

How is a Machine Learning Model Developed?

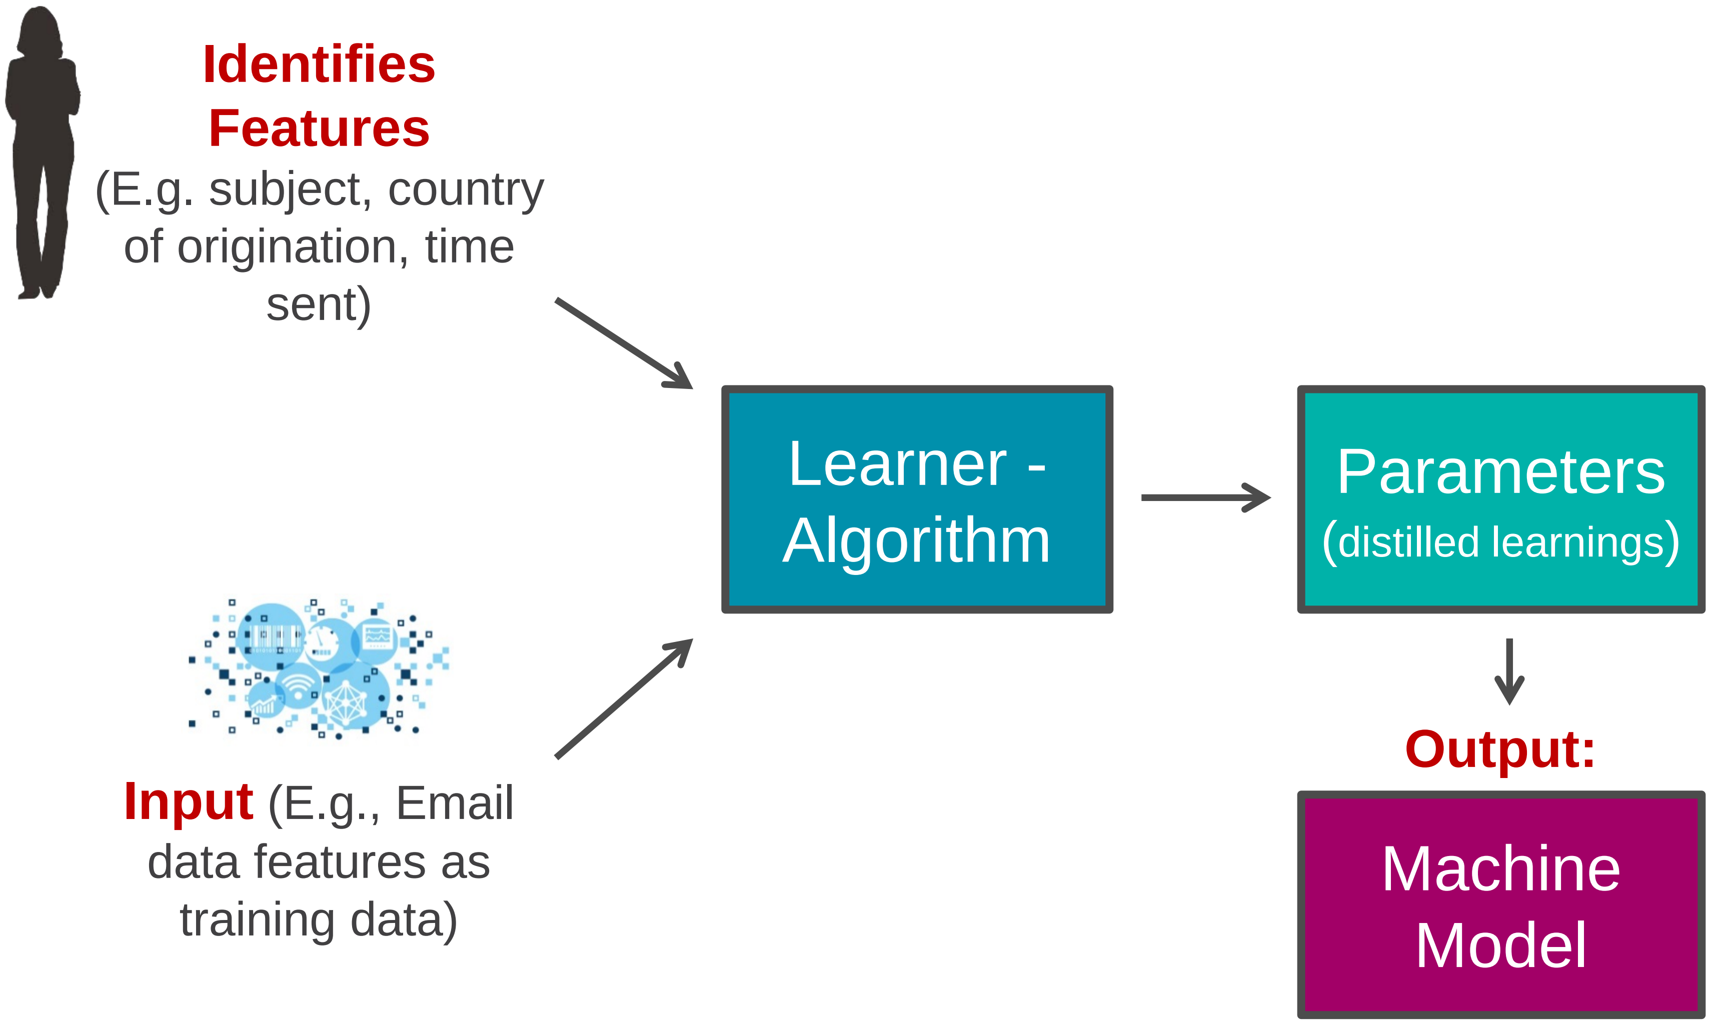

Identifies Features
(E.g. subject, country of origination, time sent)

Input (E.g., Email data features as training data)

Learner - Algorithm

Parameters
(distilled learnings)

Output:

Machine Model

How does a Supervised Machine Learning System Operate?

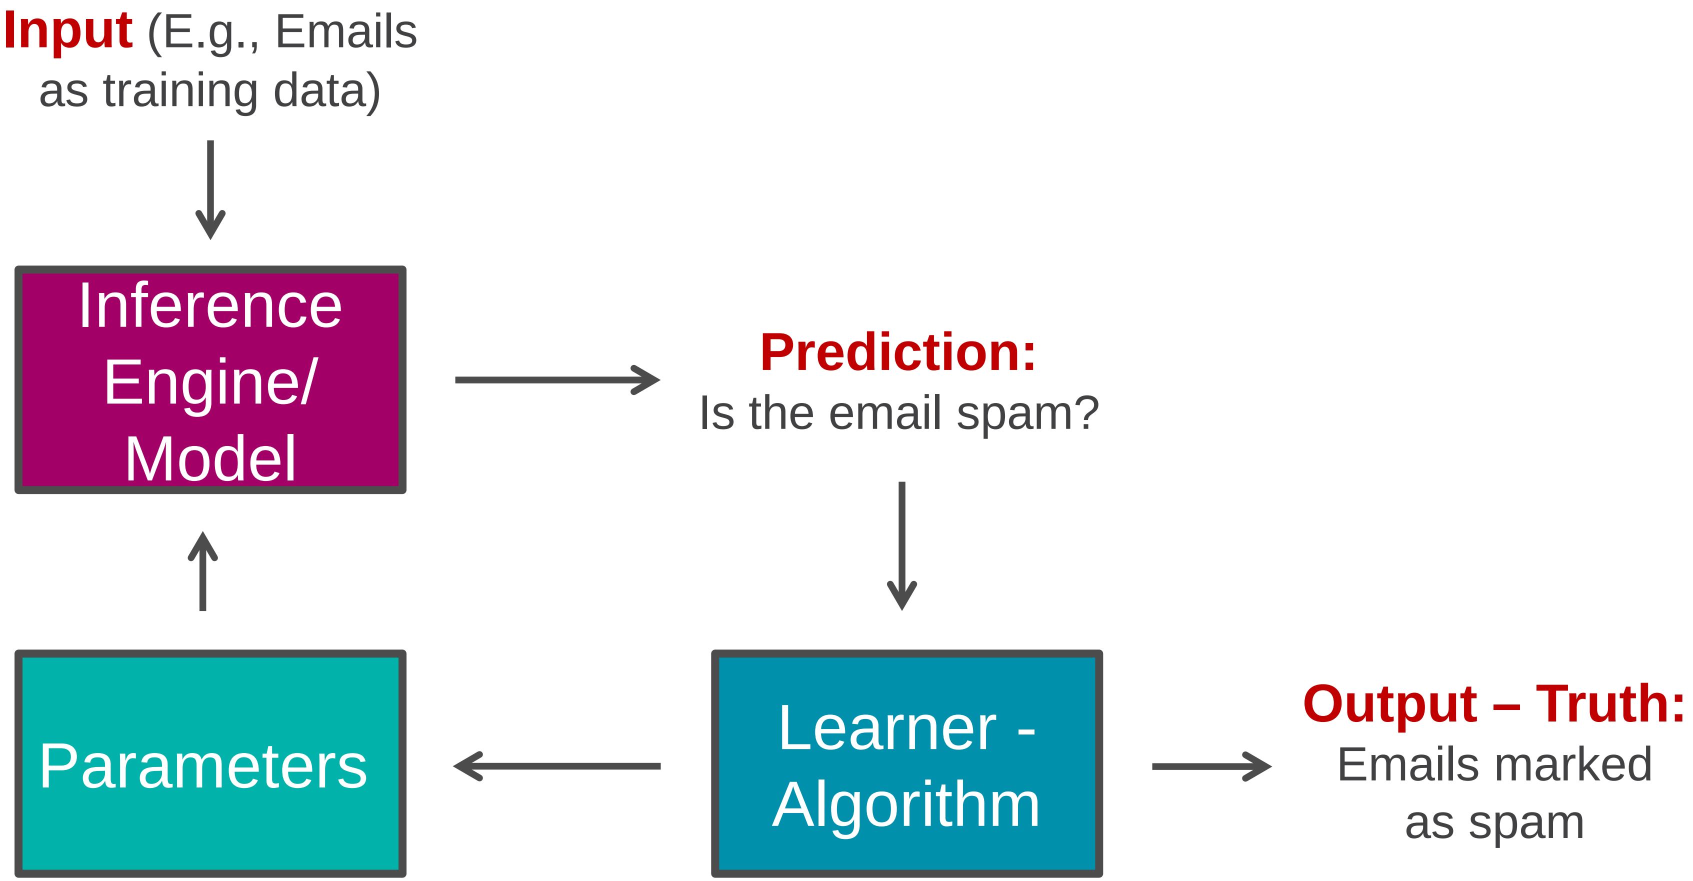

Input (E.g., Emails as training data)

Inference Engine/ Model

Prediction:
Is the email spam?

Parameters

Learner - Algorithm

Output – Truth:
Emails marked as spam

Dimensionality of Data

One Data Feature
(e.g. email subject)

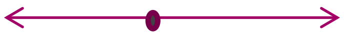

Two Data Features
(e.g. email subject, origin)

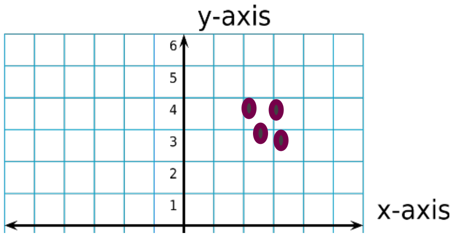

Three Data Features
(e.g. email subject, origin, time)

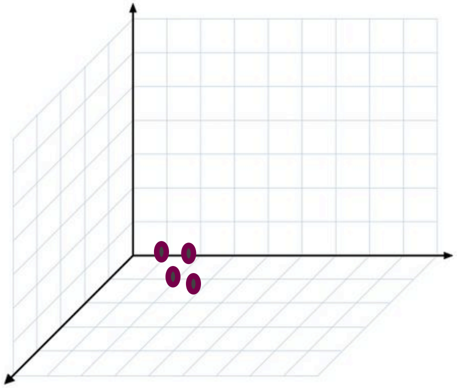

Reinforcement Learning

- Reinforcement Learning enables the machine to learn based on its environment.
- The machine learns from reinforcement signals – simple reward feedback.
- Essentially the machine is creating its training data set through trial and error.

Unity Machine Learning - Reinforcement Learning Demo

Google's DeepMind AI Just Taught Itself To Walk

What is Deep Learning?

- Subset of Machine Learning
- Deep Learning class of machine learning algorithms that use a cascade of multiple layers to build complex concepts out of simple concepts using “neural networks”
- E.g. Object recognition

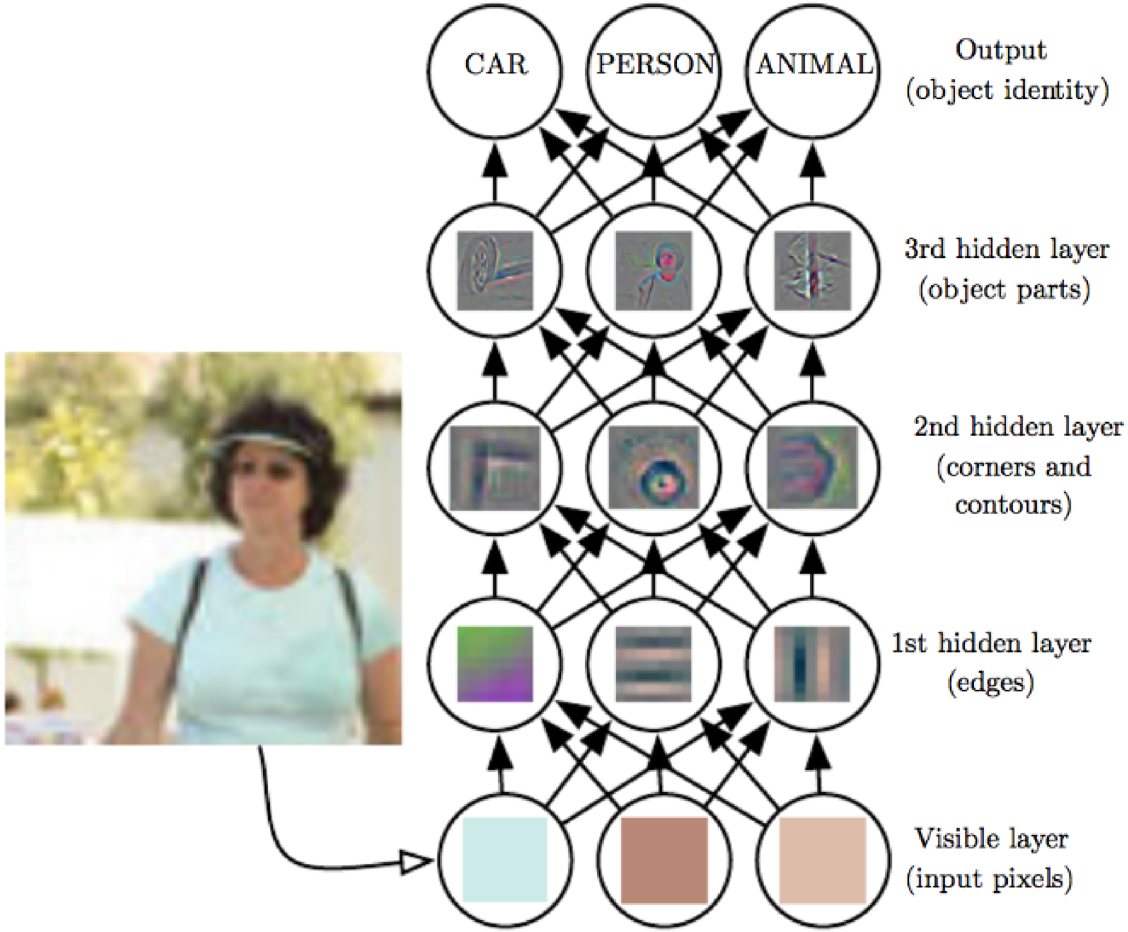

Source: Ian Goodfellow, et. al., *Deep Learning*

Google Duplex

Emerging Issue in Machine Learning

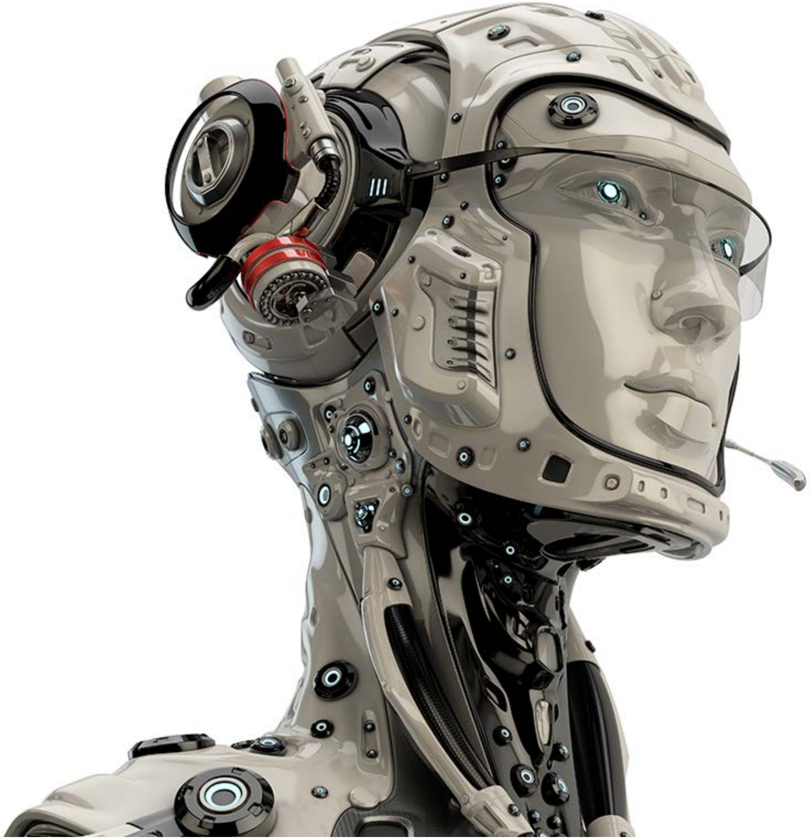

- Algorithmic Discrimination
- Privacy
- Automated Decision-Making
- Antitrust: Algorithmic Collusion
- Product Liability
- Demands for access

A Model Is Only as Good as the Training Data

- Biased Datasets = Biased Models
- Bias may be completely unintentional
- Bias may be introduced by the data scientist, or inherent in the data

Algorithmic Discrimination

Types of Biases in Data Sets

- Interaction Bias
- Latent Bias
- Selection Bias

FTC's Guidance

Big Data, A Tool for Inclusion or Exclusion?
FTC Report January 2016

- Review your data sets:
  - If they are missing information from particular populations, address it.
  - Ensure that hidden bias are not having an unintended impact.
- Correlation is not causation. Consider human oversight for important decisions, e.g., about health, credit, employment.
- Consider whether fairness or ethics require foregoing big data.
- Consider using big data to advance opportunities for underrepresented populations.

Emerging Issues in Machine Learning

Privacy

Do your public representations adequately disclose how you will use data for machine learning? Have you obtained consumer consent?

Regulatory Restrictions

- Companies should analyze whether their use of consumer data is compliant with regulatory privacy requirements.

- FTC Act – Companies should consider whether they are violating any material promises to consumers or whether they have failed to disclose material information to consumers. In addition, companies should take care to reasonably secure consumers' data.

- HIPAA – Aggregation of protected health information is only permitted without a patient authorization for Data Aggregation Purposes (i.e. assisting a covered entity with improving healthcare operations)

GDPR – Privacy and ML

> To lawfully process data for individuals that are subjects of the EU, you must obtain “freely given, specific, informed and unambiguous” consent from the data subject.

> In the context of ML, to obtain such consent you must disclose a variety of things including the existence of any automated decision-making.

> In addition, creating mechanisms to allow for erasure of the data subject’s data or the ability to restrict processing of that data is important to comply with the GDPR.

Automated Decision-Making

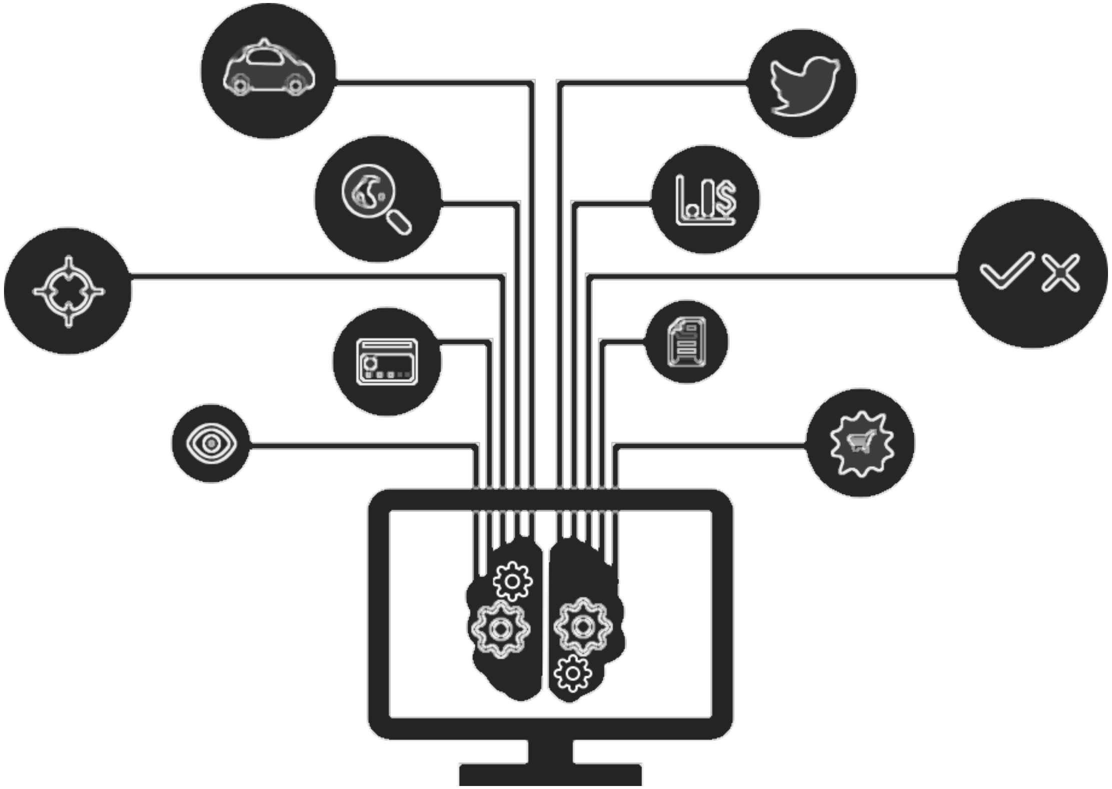

GDPR Article 22 “The data subject shall have the right not to be subject to a decision based solely on automated processing, including profiling, which produces legal effects concerning him or her or similarly significantly affects him or her.”

GDPR – Exceptions to Article 22(1)

- Where the decision is necessary for entering into, or performance of, a contract between the data subject and a data controller.

- Where the decision is based on the data subject’s explicit consent.

- Where the decision is authorized by Union or Member State law to which the controller is subject and which also lays down suitable measures to safeguard the data subject’s rights and freedoms and legitimate interests.

GDPR – Automated Decision Making

>

Article 13(2)(f) and 14(2)(g) requires data controllers to provide to data subjects, and Article 15 gives data subjects the right of access to: information about ‘the existence of automated decision-making, including profiling, referred to in Article 22(1) and (4) and, at least in those cases, ***meaningful information about the logic involved***, as well as the significance and the envisaged consequences of such processing for the data subject’ (emphasis added).

Automated Decision Making: Harms and Mitigation

Individual Harms – Illegal*

Harms	Description	Mitigation Tools
Employment Discrimination	Existing law defines impermissible outcomes, often for protected classes	• **Data methods** to ensure proxies are not used for protected classes and data does not amplify historical bias. • **Algorithmic design** to carefully consider whether to use protected status inputs & trigger manual review. • **Laws & Policies** that use data to identify discrimination.
Benefits Discrimination		
Housing Discrimination		
Education Discrimination		
Credit Discrimination		
Differential Pricing		
Individual Incarceration		

*Future of Privacy Forum, Unfairness by Algorithm: Distilling the Harms of Automated Decision Making, December 2017

Emerging Issues in Machine Learning

Antitrust: Algorithmic Collusion

(Webinar Wed., May 16)

Relevant US laws

Sherman Act Section 1

- prohibits agreements among persons or companies
- Criminal and civil penalties

FTC Act Section 5

- Prohibits unfair practices
- Does not require agreements
- Civil penalties

Algorithmic Collusion – Antitrust Division Remarks (Sept. 2017)

Algorithmic Collusion: Antitrust Division Remarks (September 2017)

- “There are **pro-competitive benefits** to technological innovations in the marketplace. **Algorithmic pricing can . . . be highly competitive**.”
- “**Our focus** in price-fixing cases . . . must remain **concerted action**”
- “Whether the concerted action is effected through direct communications or a **common understanding that competitors will use the same software to achieve the same result**, an **illegal agreement remains essential** to antitrust liability . . .”

Application to Use of Algorithms

> **Competitors** agree

- To fix prices & use algorithms to do it or
- To use same pricing algorithm

> **Algorithm supplier**
initiates and organizes price fixing agreement, assures competitors others are using same algorithm

> **Competitors** coincidentally use same algorithm, or use different algorithms that lead to higher prices

Analysis

> **Competitors** agree: Per se unlawful, algorithm or not. Criminal violation

> **Algorithm supplier**: possible civil or criminal violation for supplier and competitors

> **Competitors** using same algorithm: Appears suspicious; take care in company communications not to suggest agreement on pricing or algorithms.

Algorithmic Collusion – EU Perspective (Sept. 2017)

"Data collection may also facilitate collusion when these data are used to fix prices through the use of algorithms. . . For example, by processing all available information and thus monitoring and analyzing or anticipating their competitors' responses to current and future prices, competitors may easier be able to find a sustainable supra-competitive price equilibrium which they can agree on. Furthermore, data-crunching algorithms can also be used to implement an agreement, detect deviations and more generally let the collusive prices react in a more precise manner to changes in exogenous market conditions."

-Competition Law and Data, 10th May, 2016,
Autorité de la Concurrence, Bundeskartellamt

Emerging Issues in Machine Learning

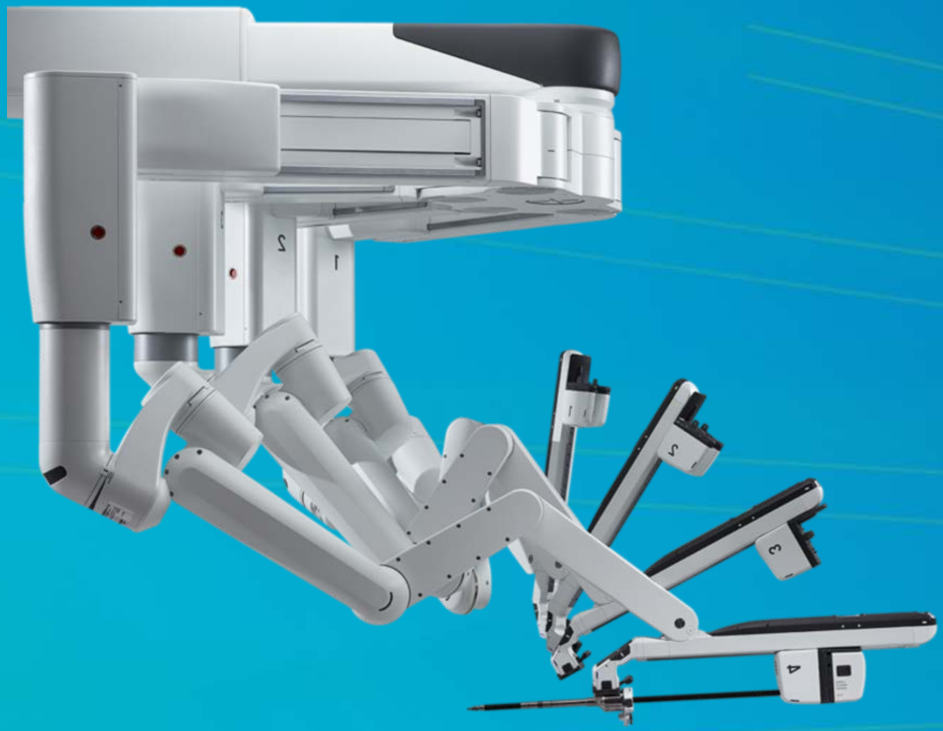

Product Liability

Product Liability – Autonomous Vehicles

- In October of 2016, Mercedes Benz executive Christoph von Hugo was quoted as saying that MB autonomous vehicles will save the car’s drivers and passengers, even if that means sacrificing the lives of pedestrians.
- A week later, MB stated that Mr. Hugo was misquoted and “neither programmers nor automated systems are entitled to weigh the value of human lives”.
- Trolley Stop Ethical Dilemma
- Narrow AI Problem

Product Liability – Machine Learning

> Machine Learning is a new technology. Strict liability makes a manufacturer or vendor responsible for all injuries that might be caused by a defective product that is unreasonably dangerous to the user, consumer or to his or her property

> Design defect: A design defect refers to a whole class of products, that are inadequately planned and therefore unreasonably hazardous for consumers.

> The manufacturer has the duty to warn the consumer against a hazardous use or instruct the consumer how to use the product properly. Additionally, the manufacturer always has a duty to warn consumers against a defect discovered after the product was sold.

Emerging Issues in Machine Learning

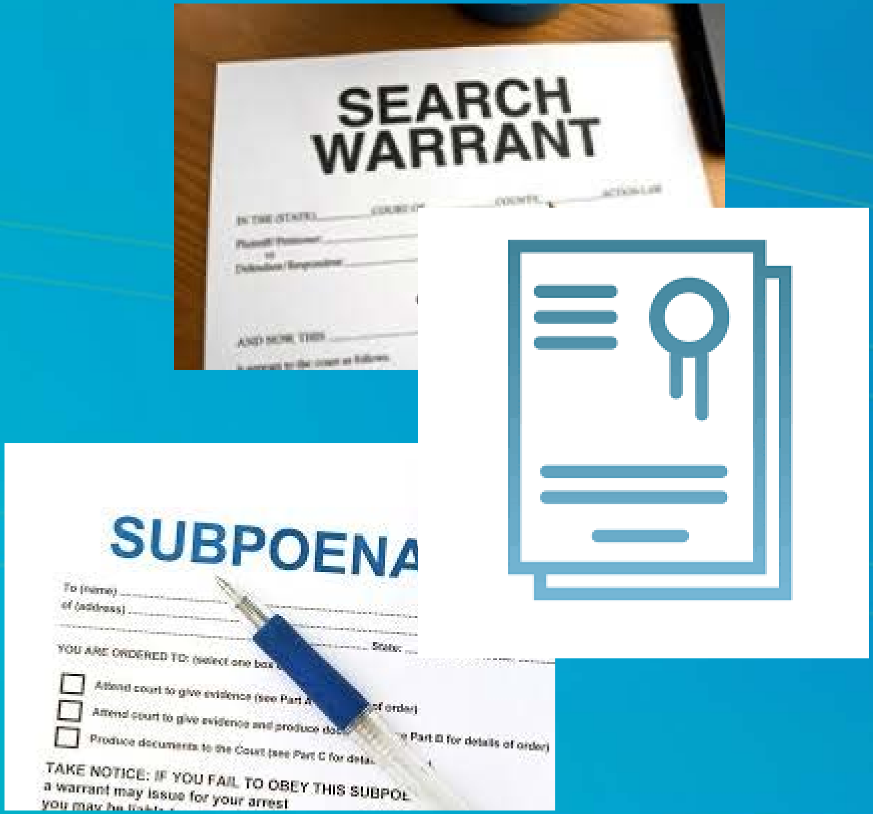

SEARCH WARRANT

IN THE (STATE) COURT OF COUNTY ACTION LAW

Plaintiff/Petitioner

vs

Defendant/Respondent

AND NOW, THIS

It appears to the court as follows:

SUBPOENA

To (name)

of (address)

State:

YOU ARE ORDERED TO: (select one box)

☐ Attend court to give evidence (see Part A

☐ Attend court to give evidence and produce do of order)

☐ Produce documents to the Court (see Part C for detai Part B for details of order)

TAKE NOTICE: IF YOU FAIL TO OBEY THIS SUBPOE a warrant may issue for your arrest you may be liable

Demands for Access
by law enforcement & Civil litigants

Demands for access by law enforcement and civil litigants

> **If you Build it, They Will Come**

> **Requests for Reporting** are expanding to new areas (child porn, terrorist content, etc.)

> **Decreased Legal Protection** as connection to user becomes more remote

Demands for access by law enforcement and civil litigants

> **Collect Only What You Need**

> **Know Your Data and Who Has Access;** limit access appropriately

> **Institutionalize Privacy:** Standardize processes

> **Centralize Legal Process Response**

Experience the Moment at the Olympic Winter Games | Intel

Emerging Issues in Machine Learning

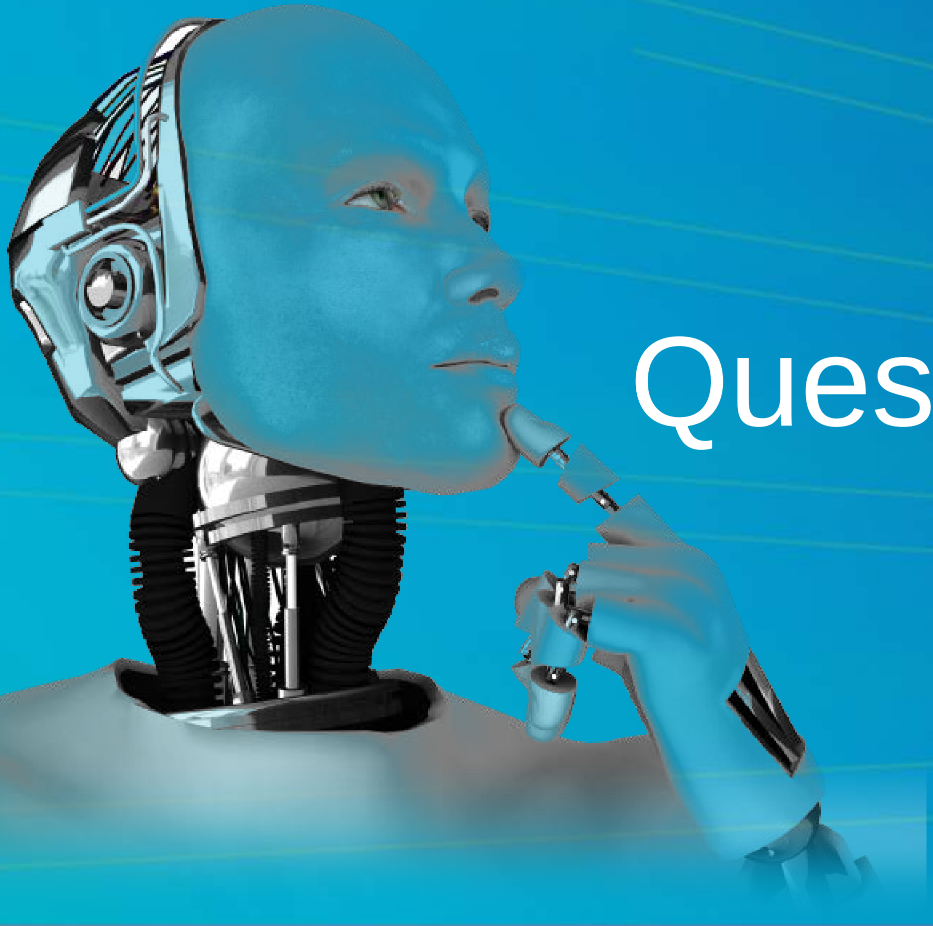

Questions?