#15: CS Ethics

SAMS SENIOR CS TRACK

Last Time

Understand how **efficiency** changes how long a program takes to run.

Recognize that some problems **probably can't be solved efficiently**.

Recognize that other problems **probably can't be solved at all!**

Today's Learning Goals

Consider **privacy and security** as core components of the application-building process

Understand how **artificial intelligence and machine learning** can have unanticipated side effects

Data Privacy and Security

User Data

Most applications collect data about users from various sources

- Data provided by the user (profile information, tweets, searches, preferences)
- Data provided by the browser/system (IP address, timezone, plugins, installed applications)
- Data provided by other sources (cookies, tracking software, data-collection companies)

This collection of data isn’t always a bad thing; you probably want Grubhub to know where you live if you want to get your food! But use of user data has gotten more complicated in recent years...

Data Economy

Why are so many companies interested in data collection? **Data has become the economy of the internet**. Most websites are supported by advertising, and advertisers pay more for targeted ads.

Websites have a strong incentive to get the best data possible on their users, so they get paid more for advertisements.

Even companies that don’t rely on advertising have a use for user data- they can **sell it to other companies**. This has become common practice in recent years, but has also received pushback from consumers.

Using Data in Your Programs

When you use user data as part of your programs, you need to be conscious of **how it is being used** and **whether users would be okay with that data being shared**.

There are also legal restrictions on the treatment of certain types of data. In the EU, the **GDPR** gives users certain rights over their data; in the US, data generated by children (**COPPA**) and data generated by students (**FERPA**) is protected.

FIGURE 4 ROUGHLY HALF OF ALL CONSUMERS HIGHLY UNCOMFORTABLE WITH COMPANIES USING AND SELLING THEIR DATA IN PHYSICAL SPACES

Q. How comfortable are you with companies USING vs. SELLING your data in each of the following areas, assuming you have opted-in to their products/services.

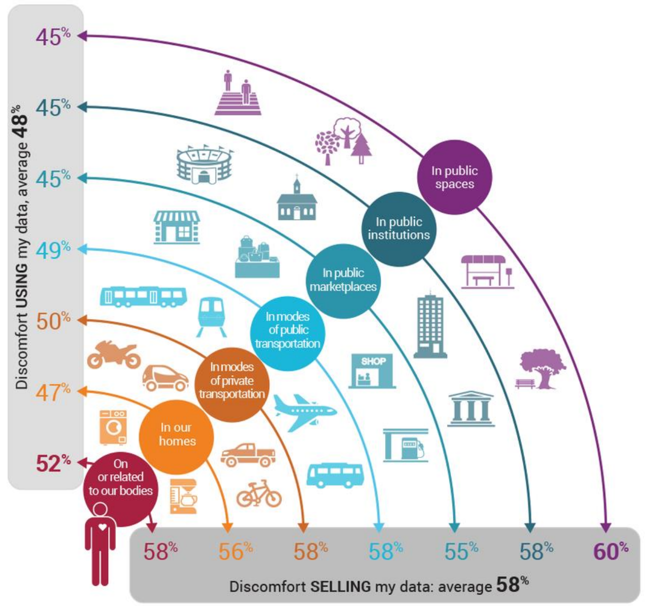

Note: These percentages reflect all respondents who, on a scale of 1-5 rated their comfort level as a 1 (Extremely uncomfortable) or 2 (Uncomfortable) with companies using vs. selling their data across each physical space.

Source: Consumer Perceptions of Privacy in the Internet of Things, Altimeter Group, 2015 Base: n=2062 respondents

ALTIMETER

Privacy

Users may want to protect their data from others for a variety of reasons.

- **Personal-** a user might have religious/cultural views on privacy.
- **Safety-** a user may identify with a persecuted minority in their country, or may have a health problem that could jeopardize their job
- **Practical-** a user probably doesn't want to share their credit card number with the world!

You should consider all of these factors when sharing *any* user data outside of a system.

Security

Even if you have good intentions towards protecting user data, that data may still be compromised if your system is not **secure**.

Computer systems are constantly under threat of attack from outside sources for economic, personal, or ideological reasons. No system is perfectly secure, but it is our responsibility as programmers to make our systems **as secure as possible**.

In the event that user data is breached, it is the responsibility of the data collector to alert users to the breach **quickly**.

Keeping Systems Secure

The field of computer security is huge! As an introduction, here are a few basic pointers to keep in mind:

Encrypt sensitive communications. Any data you send via the internet can be read by others on the network. Encrypting that data means that others will only see nonsense, not the data you send.

Never store login data in plaintext. Instead, convert the user's data using a hash function, and store that. You can still compare a user's entered password to the one stored in the database, but if an adversary gets access to the database, they don't have the password itself.

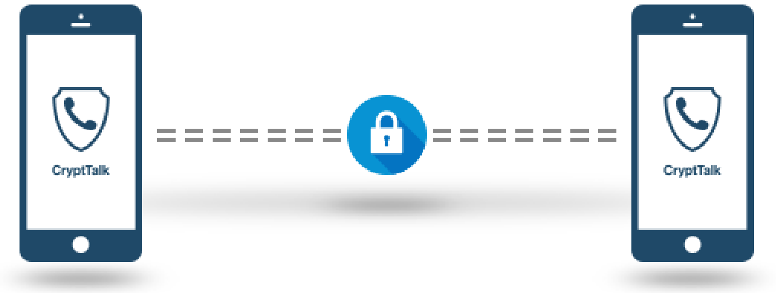

More Security Pointers

Restrict what kind of input users can provide. SQL injections are commonly used to get access to databases via text entry. Restrict input size and type to avoid this.

Be vigilant against social engineering attacks. Most hacking is not done via technical means. Instead, most hackers trick others into giving them access to systems via phishing emails or phone calls. Always be careful before clicking on links or downloading content.

Artificial Intelligence and Machine Learning

Promise of AI and ML

Artificial Intelligence: the study of how to build machines that are ‘intelligent’, that can observe their environment and act to achieve goals.

Machine Learning: one approach used to train AIs. It trains a machine to ‘learn’ a model of how a system works using many inputs and expected outputs. The AI can recognize patterns and make predictions on new data.

ML has great potential for automating tasks and improving life, but there are **potential drawbacks.**

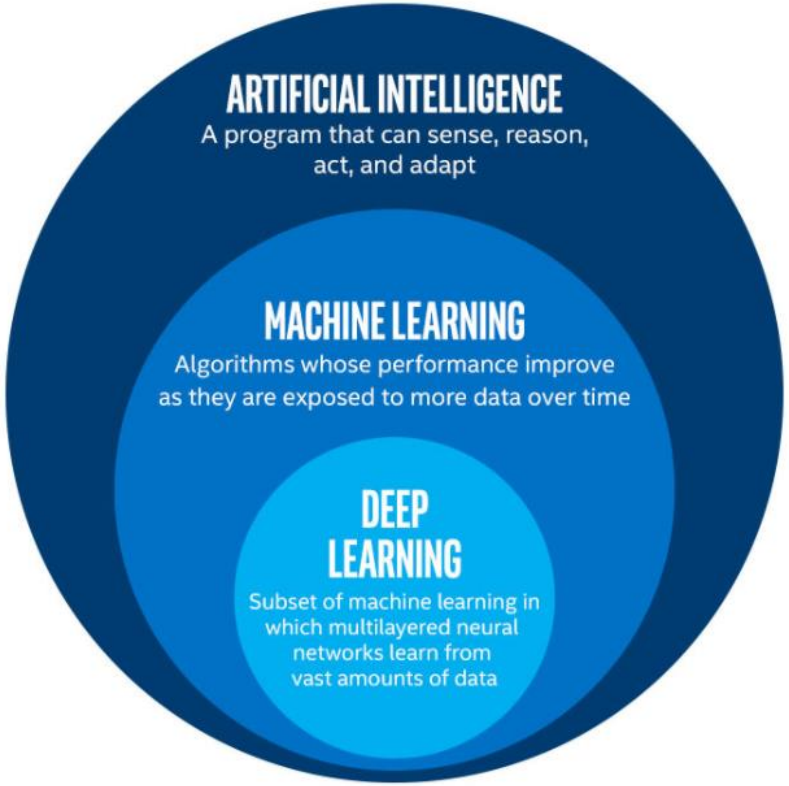
ARTIFICIAL INTELLIGENCE
A program that can sense, reason, act, and adapt
MACHINE LEARNING
Algorithms whose performance improve as they are exposed to more data over time
DEEP LEARNING
Subset of machine learning in which multilayered neural networks learn from vast amounts of data

Bias in Machine Learning

Machine Learning is highly dependent on the data that is provided to train the system. If there is bias in the data, that bias will be **propagated** into the rules the machine learns.

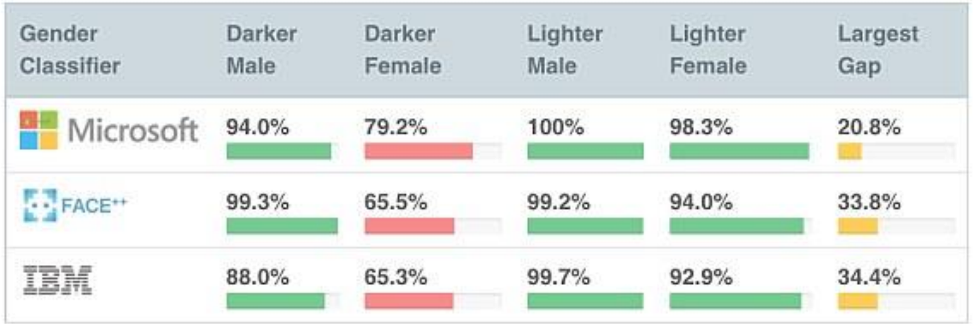

Gender Classifier	Darker Male	Darker Female	Lighter Male	Lighter Female	Largest Gap
Microsoft	94.0%	79.2%	100%	98.3%	20.8%
FACE++	99.3%	65.5%	99.2%	94.0%	33.8%
IBM	88.0%	65.3%	99.7%	92.9%	34.4%

This has caused huge problems in image recognition systems, which are often trained on data produced by the computer programmers who write the algorithms, and who are not representative of the rest of the world.

© MIT Media Lab

Bias in Machine Learning

This has also caused problems in algorithms for determining bail, which have shown systematic bias in determining recidivism rates based on race.

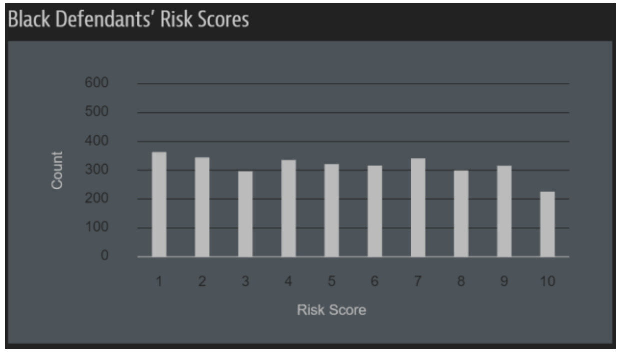

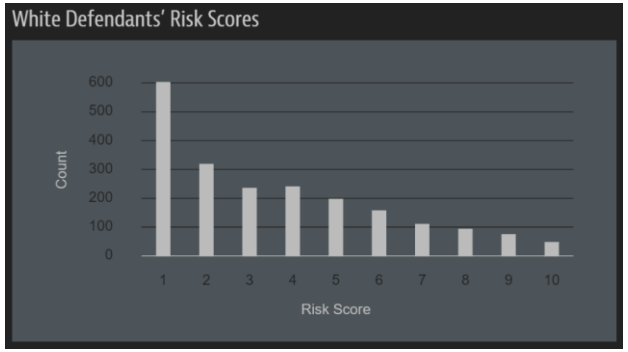

This bias is compounded by the problem of **explainability.** An AI cannot explain *why* it makes decisions; it just makes them. This is a problem when the algorithm is making an important decision about a person's life.

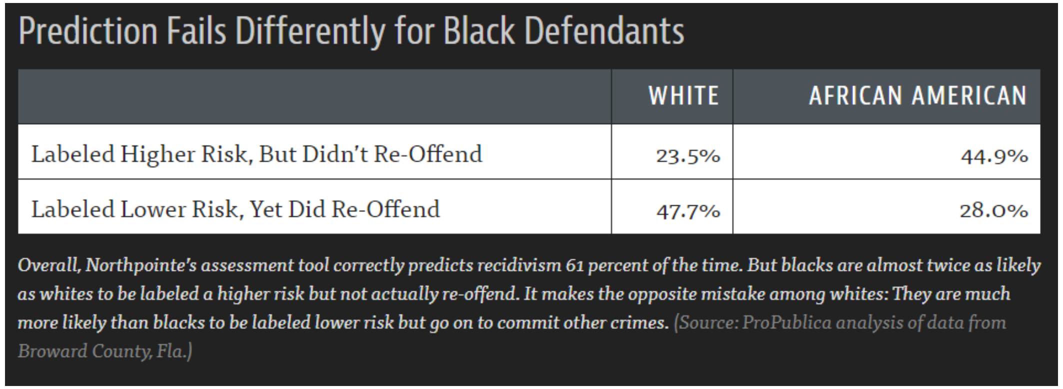

Prediction Fails Differently for Black Defendants

	WHITE	AFRICAN AMERICAN
Labeled Higher Risk, But Didn't Re-Offend	23.5%	44.9%
Labeled Lower Risk, Yet Did Re-Offend	47.7%	28.0%

Overall, Northpointe's assessment tool correctly predicts recidivism 61 percent of the time. But blacks are almost twice as likely as whites to be labeled a higher risk but not actually re-offend. It makes the opposite mistake among whites: They are much more likely than blacks to be labeled lower risk but go on to commit other crimes. (Source: ProPublica analysis of data from Broward County, Fla.)

Effect of ML on Communication

We've seen in recent elections how networks of bots can be used to spread misinformation through a community. Bots that masquerade as real people can be used to affect communities at scale, putting text communication at risk.

Recent advances in technology have also made it possible to edit videos convincingly. This may have major repercussions on public discourse. Similar work is being done in audio editing.

Effect of ML on the Environment

Many companies and researchers train machine learning algorithms on very large datasets to answer questions.

This analysis does not come without a cost. An enormous amount of energy is needed to run these algorithms, and in the US, that energy often has a carbon footprint. A recent study found that training a popular new NLP model, The Transformer, leaves a gigantic carbon footprint.

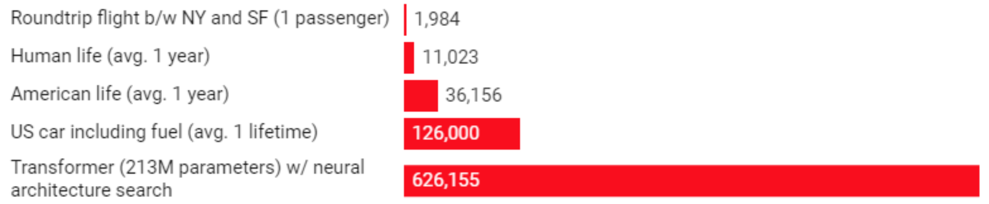

Common carbon footprint benchmarks

in lbs of CO2 equivalent

Chart: MIT Technology Review • Source: Strubell et al. • Created with Datawrapper

Instead of constantly re-training models, we should share ML models where possible, to conserve energy!

Ethics in AI Design

Finally, the programmers of Artificial Intelligence need to consider the repercussions of their design decisions while creating an AI.

Consider the Trolley Problem, and apply it to a self-driving car. Should the car protect its passenger, or should it optimize for the greatest preservation of human life?

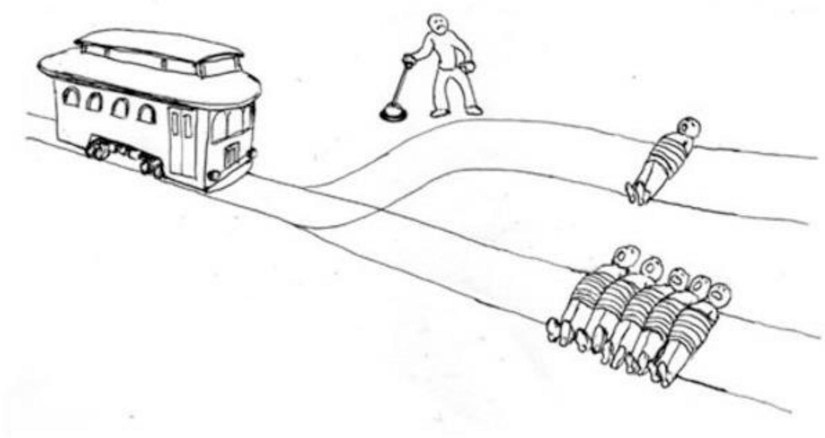

Ethics in AI Design

At a broader level, consider the effect that AI automation has on jobs.

Many of the most common jobs in the modern day are likely to be automated in the next twenty years.

Potentially the job sector will grow to find new employment opportunities for workers, but this will still cause disruption.

2.1.9 Impact of automation on jobs

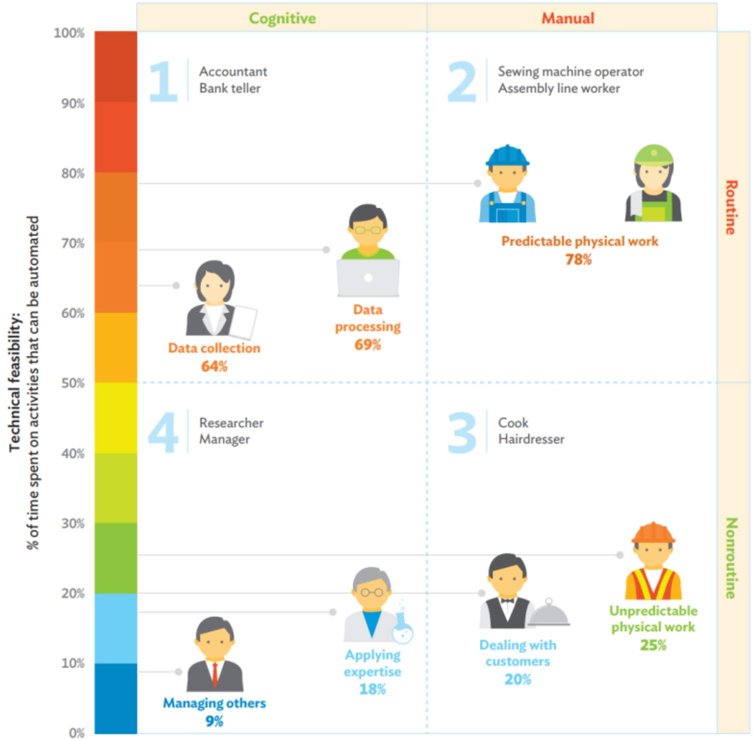

Note: Percentages are from Frey and Osborne (2017) estimates on probability of automation. Framework is based on Acemoglu and Autor (2011).

Today's Learning Goals

Consider **privacy and security** as core components of the application-building process

Understand how **artificial intelligence and machine learning** can have unanticipated side effects