University of Illinois at Urbana-Champaign
Dept. of Electrical and Computer Engineering

ECE 101: Exploring Digital Information Technologies for Non-Engineers

Ethics and Privacy

An Example of Learning by Example

Learning by example…
1. Limes are green.
2. Grapes are green.
3. Kiwis are green.
4. Apples are green.

CONCLUSION: ALL FRUITS ARE GREEN!

Humans Learn in Other Ways, Too

That's not the only way humans learn.

For example, inference:
1. Fruits are green.
2. Fruits are food.

→ **some foods are green**.

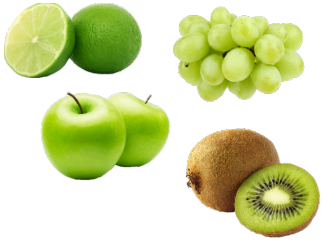

Humans Learn Continuously

And **humans keep learning**…

Here are some peaches.

→ fix model: **not all fruits are green**

Machine Learning is Mostly Learning by Example

Primary technique for machine learning:
- **start with** a set of **labeled examples**,
- **train a model** to produce the labels, then
- **evaluate** using more labeled data.

As with humans,
- **biases** in data **and absences** of data
- **produce biased / partial models**.

Unlike humans, training is much more costly than evaluation, so **learning** is typically **not continuous**.

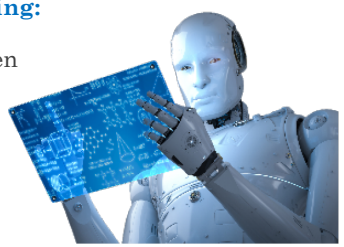

Techniques Can Support Human Analysis, Too

To be clear, **some techniques**, such as grouping points into clusters, **can be used for data science**, which helps humans to make sense of large amounts of data.

The line is a little blurry.

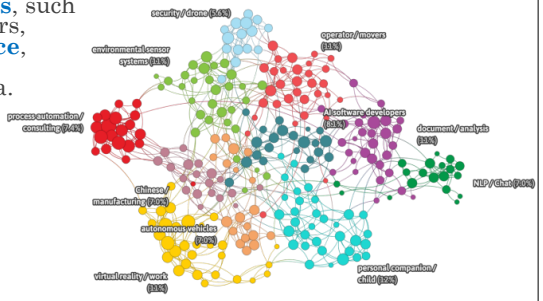

Example: Machine Prediction of Likely Variants of Interest

Example: use historical COVID data
- to **develop a model that**
- **predicts** which COVID **variants**
- are **worth investigating** in a lab.

Model only used for prediction:
- choose best 50 out of 10,000 variants to evaluate in the lab.
- **Explanations** of importance **based on lab work**.

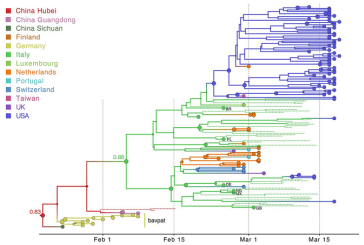

What's Better, ML or Humans? (Advantages)

For any given task, **we can compare use of ML with** use of **human employees**.

Advantages of using ML over humans:
- **Capital costs lower**.
- Usage more **flexible** (same chip executes almost any task, unlike skilled humans).
- **Operating costs** much **lower** (and go down rather than up in off-peak hours).
- **Operates 24/7** in absence of failure.

What's Better, ML or Humans? (More Advantages)

More advantages of using ML over humans:
- **Lower failure rate**.
- **No psychological issues**.
- **Minimal safety issues** (only downtime and replacement cost).
- **Replication and replacement nearly instantaneous**.
- **Lightweight and portable** (can fit on a drone or a variety of other robots that can reach places humans generally cannot).

What's Better, ML or Humans? (Disadvantages)

Disadvantages of using ML:
- **Requires** electrical **power**.
- **Results may not be as good** (balance against advantages).

Another **major disadvantage:** decisions **not explainable**.
- Can't learn from ML models.
- No "reason" for results, correct or otherwise.

Human Skills are Also Often Unexplainable

Aside from general physical fitness and experience, **why is Michael Jordan a better basketball player than I could ever be?**

Why does Yo-Yo Ma's ability with the cello exceed anything I could hope to achieve, no matter how hard I try?

I'm not being pessimistic here—humans have variations in innate skills.

I like to think I'm a better engineer than either of them, but I couldn't explain why.

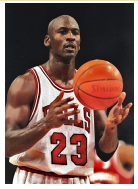

Is Machine Intelligence Ethical?

Let's talk about ethics.

What does the word mean?

Ethics: moral principles that govern … the conducting of an activity

Morality Depends Strongly on Culture

But … what are moral principles?

Across human cultures,
- there are **a few common themes**.
- For example, **don't kill people** from your village.

But **otherwise, morals are highly variable**.

Who Keeps Companies Ethical?

Corporate actions
- are **governed by people**,
- **but** individuals do **not always**
- apply **personal ethos**
- in a competitive setting.

Only Government Can Limit Companies

Bottom line: **ethical business behavior**
- **defined by law and policy**,
- **interpreted by judiciary**, and
- **enforced by government**.

In a democracy, **all** of these **start with you**.

Who is Responsible for ML Decisions?

Since results generated by ML are not explainable,

who is responsible for actions based on ML-controlled agents?

Let's say an autonomous car.

Many Answers, but Perhaps Not Many Good Ones

- the car owner (possibly in the car)
- the car passenger(s), if any
- the car manufacturer
- the company and/or person(s) who developed the ML model
  (maybe most people would stop here?)
- the company that fabricated the chip used to execute the ML model
- the company and/or person(s) who designed the chip used to execute the ML model
- the company and/or person(s) who wrote simulations for producing training data
- the company and/or person(s) who supplied data for writing simulations for training
- the company and/or person(s) who designed the algorithm for learning
- the company and/or person(s) who designed the algorithm for executing the model

Is Older "Smart" Technology So Different?

What if the element at fault was instead the Anti-lock Braking System?

Does responsibility shift?

Does the line shift?

Why?

Because ML is more complex than ABS?

Legal Decisions Made by Humans

ML failures will lead to both criminal cases and civil litigation.

In many such cases,
- **juries make decisions**.
- Jurors are **unlikely to have** much technology **expertise**.

Are humans capable of avoiding the assu that technology is unbiased and infallible, especially with (one side's) lawyers arguing that such is the case?

Humans Avoid Conflict with Technology

Example (true story): day-care billing (just software—not even AI!)

Price **per day: $100.00**

Price **for two days: $200.05**

Why?

"That's what the computer says."

Even Old Technologies Can be Problematic

Technology example: lie detectors

Usually pretty accurate, but not perfect. Results require interpretation by a trained human.

If you're in a jury, and the accused "lies" about their guilt, how much weight do you put on that result?

Do you have any way to estimate the frequency of false positives?

To know the skill of the interpreter?

New Technologies Unlikely to be Understood

Another technology (ML-based): face recognition

Phones allow it now for unlocking.

Success rate? 70% maybe?

What if someone is wearing a wearing a mask?

What if they have a bruise?

Easy to Misuse and Misrepresent Results from ML

Same technology: face recognition

Yes, **that's the person who broke into your house**.

Really?

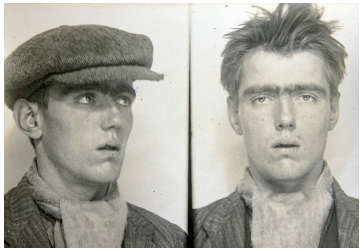

Legislators are Criminal? Well…

From https://sitn.hms.harvard.edu/flash/2020/racial-discrimination-in-face-recognition-technology/

"…**assessment** of [Amazon] Rekognition's face-matching capability **by** the American Civil Liberties Union (**ACLU**), in which 28 **members of Congress, disproportionately people of color,** were **incorrectly matched with mugshot images**."

Bias in ML is Real, and Will Cause Problems

(same site)

"A growing body of research exposes **divergent error rates across demographic groups**, with the **poorest accuracy** consistently found **in subjects who are female, Black, and 18-30** years old."

Should Use of ML be Regulated?

Perhaps we should develop standards before allowing use of ML?
- **safer** than the best humans
- **more accurate / correct** than the best humans
- **less biased** than the best humans

(These may not be easy to define and/or evaluate.)

The question of responsibility will still arise, but at least **fewer humans will be harmed**.

How Does Privacy Interact with Intelligence?

Now, let's consider privacy.

What does the word mean?

Privacy: the state or condition of being free from being observed or disturbed by other people

Privacy vs. Surveillance and Control

Why might someone else want to "observe or disturb" you?

Two reasons:
- **surveillance**: to know what you are doing (the reason could be good, bad, or in between), **and**
- **control**: to influence, limit, or manage your actions.

When are Surveillance and Control Acceptable?

To decide whether such activities are acceptable, we can start with these questions:

- **When** and to what extent **does** government or **someone** else **have the right to monitor your actions?**
- **When** and to what extent **does** government or **someone** else **have the right to control what you are allowed to do?**

"Never" is Not the Right Answer

You may be tempted to answer, "Never."
But that's not a good answer.

Am I allowed
- to come into your apartment as I please and
- use your television,
- eat your food, and
- so forth?

So you want to monitor my behavior, control me, and perhaps even disallow my visits?

You probably even expect me not to visit unless you explicitly invite me to do so.

"Your freedom ends where my nose begins." (USA, Anon)

"The right to swing my fist ends where the other man's nose begins." – Oliver Wendell Holmes Jr.

Both Culture and Politics Affect the Answers

Between
- "Who controls you in your private space?" and
- "Who controls you in another person's private space?"
- is a huge gray area.

Answers vary both culturally and politically.

Control Example: Internet Service Providers

Let's focus on a control question.

If you buy Internet service from my company, am I allowed to …
- Limit your bandwidth?
- Limit your bandwidth for specific applications?
- Prevent or restrict you from accessing some services?
- Sell your ability to access services to someone else? (In other words, only allow you to make use of services that pay me.)

Net Neutrality Akin to Freedom of Speech?

These questions are central to the issue of "**net neutrality**" that has been a topic of recent political discussions.

Do customers have the right to find content and use services that suit their tastes?

What does "free speech" mean if a small number of large companies control what can be "heard" on the Internet?

Should Control be Left to the Service Provider?

The same questions might be asked about other services, too.

Consider the last question
- in terms of search engines:
- **is a search engine allowed to sell your ability to find URLs?**

In other words, is it acceptable to put paid advertisements at the top of a search response? As the only response?

Does Payment Matter? Is Internet a Telephony Service?

Did you say "yes" because you don't pay for web search?

It's easy for an Internet service provider to include a control agreement in your contract.

The question is: **should they be allowed to do so?**

What if your cell provider only allowed you to call businesses "in their network" (in other words, businesses who pay the provider)?

Next week, we'll spend some time discussing fairness (not on Exam II).

Benefits of Strong Surveillance

Now a story about surveillance.

Minority Report is a sci-fi movie showing very believable technology about data collection and its use.

- Citizens have eye implants scanned by surveillance robots (like spiders) and security cameras
- Citizens receive tailored recommendations in shopping mall
- Do not need tickets on the subway
- Allowed to pass through security checks
- A criminal is unable to hide from law enforcement

Privacy vs. Ethics: Surveillance for your Benefit?

What if the crimes are victimless and unenforced?

10-15 years ago, rental car companies
- installed speed-monitoring chips in their vehicles.*
- Drivers who broke speed limits
  - were charged extra fees, but
  - were not reported to local authorities.

*And told customers about them in the fine print.

Customers Offered a Resounding, "NO!"

Why?
- Lower insurance costs!
- Some of the reduction
  - passed on to customers renting cars
  - (a competitive advantage).

Customers hated it, and the idea vanished.

Monitoring Includes Location and Interests

Surveillance can cause problems, even in the absence of crime.

Typically, monitoring collects things like
- location,
- motion,
- interests,
- comments/ideas, and
- events attended.

Information may be sensitive.

Surveillance Can Have Other Side Effects

In the early 2000s, speed traps in France
- automatically captured photographs of cars and
- sent them to registered addresses.

Result? Divorce rates went up… !?!

Apparently, front-seat passengers were often visible in the photographs.

The photos were eliminated, but "speed trap coming up" signsS were also posted well in advance of the traps, and were still around in 2010.

So What's Your View?

Would you agree to surveillance of your activities in return for lower prices?*

Would you agree to 24/7 monitoring of your location and actions in return for … ?

*Assume no illegal use of collected data.
But do you know what's legally allowed?

Crowdsourcing May Also Impinge on Privacy

What about crowdsourcing vs. privacy?

Perhaps you think that
- you should be able to control information
- about your vehicle's behavior,
- except when a criminal activity is involved.

But what if I gather that information?

Who Owns Information Collected About You?

Can I sell information about your car and driving?

What if my car does it for me, automatically?
- My car sees your car pass at high speed.
- Your license plate is visible.
- So I sell that information to all insurance companies, rental car companies, and so forth…
- except for my own, because I don't want them to know that I drive in high-risk areas!

Using Surveillance for (Machine) Intelligence

Another example:
travel tracking vs. identify theft

You purchase an airline ticket.

Who knows about it?
- Your **credit card company** (bill sent to them with information about passengers, etc., for fraud detection)
- Your **email provider** (confirmation email analyzed for travel plans—your name AND your companions)

Information Reduces Fraud and Improves Services

When you travel, **these companies** (and others) **can tell where you are.**

Credit card **purchases include geographical information**.

If you buy a ticket to Hawaii
- and try to make a "purchase" in-person in Florida,
- most credit cards will decline the purchase.

Any **use of** the **Internet** (WiFi) **includes a local IP address** that can be used to identify your location.

Will You Trade Your Privacy for Such Benefits?

Cell data are more difficult to use, as they are considered telephony and protected by law.

So…

Would you allow your credit card company to inspect your web history to provide better fraud detection?

Would you (do you) make use of builtin/ plugin lists of web sites that are "dangerous" in terms of likely virus attacks?

Terminology You Should Know from These Slides

- learning by example
- model bias and error
- ethics
- privacy
- surveillance
- control
- net neutrality
- crowdsourcing

Concepts You Should Know from These Slides

- advantages and disadvantages of ML
- lack of explainability for ML results
- societal assumptions about technology (and differences with reality)
- commercial and legal value of surveillance and control
- tension between surveillance and privacy
- potential benefits of allowing use of personal information
- impact of politics and governmental self-interest on privacy, law, and business