

Computer Science Ethics

15-110 – Wednesday 12/04

Learning Goals

Recognize why **code maintenance** is important and possible repercussions of badly-maintained code.

Understand the current extent of **data collection** on the internet and its possible repercussions.

Identify the societal impact of **machine learning** in terms of bias in data and responsibility for decisions made by AIs.

Code Maintenance

Coding for Real Projects

You'll leave this course with a basic working knowledge of programming, which you may want to apply to your own projects. But if you plan to write code for real projects, you'll need to treat that code like an artifact that others will use. This comes with a new set of recommendations and rules for coding.

We'll focus on three main rules: **comment**, **test**, and **attribute**.

Commenting Code

Up until now, we've primarily used comments to give instructions on assignments, and maybe to comment out non-working code.

In real projects, comments should be used to add **documentation** to your code. This makes it possible for other people to understand what your code does by scanning the comments, instead of trying to parse the code.

In general, it's good practice to have a big comment at the top of every file, and smaller comments on every function that describe what they do.

Let's check out documentation on a real-life project:

<https://github.com/trending/python?since=daily>

Testing Code

In this class, we've provided test cases for you to check your code. In real life, you'll need to write tests on your own, to make sure your code does what it's supposed to.

Test cases are primarily useful for **refactoring**- that is, making sure you don't break your code if you update it later on.

It's good practice to make sure you test for expected cases, multiple possible results (both True and False), and edge cases (empty lists/strings). Use the **assert(bool)** statement, which does nothing if bool is True and crashes if bool is False.

Attributing Code

Finally, if you borrow code from somewhere online, **cite it!** You can do this by adding a comment with a link to the original source.

This follows best practices ethically, but also will make debugging easier if something breaks down later on.

Coding with Other People

You might occasionally need to write code with another person, or with a team of people. When this happens, you may need to use **style guides** for the code you write together.

Why do we need style guides? Let's look at an example of code without and with good style...

What does the following program do?

```
def f(a):  
    f = False  
    if a < 2 :    return f  
    for i in range (2,a) :  
        if (((a)%(i))== 0)==True:  
            return not f  
    return f
```

```
def is_prime(num):  
    if num < 2:  
        return False  
    for factor in range(2, num):  
        if num % factor == 0:  
            return True  
    return False
```

Coding with Style

A **style guide** for coding is like a style guide for writing- it's a set of rules that describe how you should format the code that you write.

Style guides let you standardize format across multiple people, so that everyone can easily write and modify each other's code.

Python's official style guide is [PEP 8](#), but different organizations and companies may have their own style guides.

Real Life Implications

Why does all of this matter? Computer science is a very open-source field, and people share and use each other's code all the time.

However, you can't write code once, share it with others, and then be done with it. Code lives in an environment that is constantly changing- languages update, new OS versions are released, and expectations change. **Modules regularly need to be updated to fix bugs and support new systems.**

For an example, let's look at the Numpy release logs:

<https://docs.scipy.org/doc/>

Security Concerns with Legacy Code

Many companies (and governments) rely on old code systems that have not been updated in decades. This makes it difficult to upgrade systems, and also leaves organizations open to security threats.

[A recent analysis](#) of the US government showed that that government spends 80% of its IT budget on legacy code maintenance, and that ten different systems across different agencies are a critical security risk.

Of those systems, three have been used for over 30 years!

Open Source Cautionary Tale

Finally, always be aware of code that is open-source one day may not be available forever!

In March 2016, Javascript developers around the world suddenly started receiving an error message when they attempted to run their code. This had happened because a developer decided to remove his open-source code from an online database, and that code (specifically, a function called left-pad) had been used in several popular modules, including React.js.

Read more here:

<https://qz.com/646467/how-one-programmer-broke-the-internet-by-deleting-a-tiny-piece-of-code/>

Data Collection

User Data

Most applications collect data about users from various sources. We'll discuss three main categories: data provided by the **user**, data provided by the **browser/system**, and data provided by **other sources**.

As a user of the internet and various applications, you already voluntarily share a lot of data with the world!

- Internet – profile information, tweets, searches
- Applications – preferences, locations, images
- Real life – purchase history, contact info, location

Browser/System Data

Behind the scenes, your browser or phone/computer is sending additional information to the services you use.

This is not done maliciously- services can put this information to good use. However, you may be surprised by some of the data being shared.

Check out the data your browser shares here:

<https://webkay.robinlinus.com/>

There are plugins you can install that limit the information your browser sends, but this may also limit functionality of websites.

Other Data Sources

Cookies are used by websites to store temporary information about people using their services (like which items you've put in a shopping cart). A cookie is a small packet of data that is sent back and forth between the website and your browser.

Cookies that are shared between two or more websites are called **tracking cookies**, or just trackers. These cookies attempt to collect a portfolio of information on you, the user, by gathering information on the websites you visit. This is commonly done through ads that are placed on websites.

You can check what kinds of trackers your browser stops here:

<https://panopticlick.eff.org/>

Data Economy – Data Collection

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. This has led to **hyper-targeting** in ads, with ads attempting to reach more niche populations.

If you have a Facebook account, try going to [Settings > Ads > Your Information > Your Categories](#). This will show you the niche groups Facebook thinks you might be a part of.

Data Economy – Selling Data

Even companies that don't rely on advertising have a use for user data- they can **sell it to other companies**. This data is aggregated by companies that can then sell portfolios of individuals to advertisers or insurance companies.

Even when companies promise not to sell individual data, it still isn't entirely private. For example, consider online DNA services like 23andMe. This site (and many others) sell **aggregated data**; though this data does not have a user's name or address attached, the genetic information is still shared.

Data and the Government

There are also concerns around how companies share data with police forces and the government. For example, the smart doorbell company Ring recently [formed a partnership with police forces across the US](#) to share video data, with homeowner permission.

Some governments have gotten more directly involved in large-scale data collection; in particular, China has recently instituted the [Social Credit System](#), where data collected on individuals by a large number of surveillance cameras has direct impacts on their ability to interact in everyday life.

Data Privacy

Why care about all of this? You may want to protect your data from others for a variety of reasons.

Some people have religious or cultural values that emphasize privacy. Others may want to keep their data private for their own safety, if they identify with a persecuted minority or hold subversive views.

Most governments have not yet determined how to handle questions surrounding privacy and the sale and collection of data. However, there are some legal restrictions on the treatment of specific types of data.

Restrictions on Data Sharing

In the EU, the **GDPR** gives all users certain rights over their data; they must be told when data is being collected, data must be stored securely, and users have the right to obtain their data and/or ask for it to be deleted. The EU also has the **Right to Be Forgotten**, which lets users request that certain pages be removed from search results after time has passed.

In the US, data generated by children (**COPPA**) and data generated by students (**FERPA**) is protected. There is no general law about data privacy yet, but California recently passed the **CCPA**, which institutes some regulations for that state.

Protecting Your Data

If you want to protect your data online, you have a lot of options! Most browsers let you block cookies and can request that websites do not track you. You can also restrict permissions given to websites and applications on your devices.

For advanced protection, you can also use a VPN (Virtual Private Network) to connect to the internet. A VPN provides an additional layer of security to communication through 'tunneling'.




In tunneling, the request your browser sends is hidden inside an external request; the external request directs the packets to the VPN router, which then unpacks the real request and sends it on. VPNs are commonly used by companies to let employees work from home, and are used for circumventing firewalls.

Machine Learning

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.**

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.

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





© MIT Media Lab

Bias in Machine Learning

Bias has also caused problems in [algorithms for determining bail](#), which have shown systematic bias in determining recidivism rates based on race.

A similar problem was observed in an [algorithm to hire engineers for Amazon](#), which showed bias towards hiring employees based on gender.

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.)

Effects of ML on the Environment

Even if we can avoid bias, machine learning can still have unintended side effects.

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

Ethics in AI Design

Even at a theoretical level, there are still big ethical questions about how we should go about programming AIs.

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? And how should this be treated legally?



Ethics in AI Design

We could put aside the big questions about AI design, but there are smaller day-to-day actions AIs take that have effects on the world around them.

For example, Google has become a gatekeeper for much of the information in the world. If Google's search algorithm puts a small business on the second page, that could have a drastic effect on the business's revenue.

At a local level, more companies are deploying robots with AIs to move about the world autonomously. This has led to questions about the [legality of these machines](#), and even [concerns about accessibility](#) here in Pittsburgh.

The Role of Machine Learning

This all leads to an important question: which problems should be solved using machine learning, and which problems should be left alone? What are the ethical responsibilities of computer scientists?

This question has come up recently at Carnegie Mellon as well, with [protests by students against recruitment by Palantir](#), which provides software to ICE (US Immigration and Customs Enforcement).

The professional field of computer science has only recently adopted a [code of ethics](#). We still have much to debate over what the responsibilities of computer scientists are.

Learning Goals

Recognize why **code maintenance** is important and possible repercussions of badly-maintained code.

Understand the current extent of **data collection** on the internet and its possible repercussions.

Identify the societal impact of **machine learning** in terms of bias in data and responsibility for decisions made by AIs.