

# Computer Science Ethics

15-110

# Learning Goals

- Critically reflect on individual ethical concerns in relation to Computer Science
- Highlight examples of ethics-related instances in the real world
- Understand the significance of ethics from the perspective of your own academic disciplines and desired careers
- Engage in discussions of ethical questions with peers

# Ethics in Computer Science

When we move from theoretical concepts of computer science to applying those theories in real life, the decisions we make have consequences.

These consequences can be a **direct** result of the technology itself or can have an **indirect** social impact that the technology may facilitate. As responsible designers and developers of technology, it is important to consider both cases.

Today, we will cover some topics in computing that highlight how **data**, **machine learning**, and **AI algorithms** can impact individuals both directly and indirectly.

# Ethics in Computer Science

The professional field of computer science has only recently adopted a [code of ethics](#), and the code is not yet uniformly taught to new computer scientists or programmers. There is still much to debate over what the responsibilities of computer scientists are.

As its primary ethical principle, the code of ethics states that computing professionals should **contribute to society and to human well-being**, acknowledging that **all people are stakeholders in computing**.

When reflecting on **ethical responsibilities**, it may be helpful to ask yourself: “What should we value broadly as a society? And what should we do to realize those values?”

# Data

# Collecting User Data

Most applications collect data about users from various 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

Behind the scenes, your browser or phone/computer is sending additional information to the services you use. Check out the data your browser shares here: <https://webkay.robinlinus.com/>

# 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.

Even companies that don't rely on advertising have a use for user data – they can **sell data 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 Economy

By simply participating as a member of society, you may contribute to this data economy, whether you are aware of it or not. Some companies are trying to make this process [more transparent for its users](#).

Take a moment to consider the following:

- Who are the **stakeholders** in the data economy?
- Is this a fair representation of those who **participate** in it?
- Should the people whose data is being used have **ownership** and a say over its usage?



# Bias in Machine Learning & Real World Examples

# Data Bias

Another important use of data is for training and improving AI algorithms. In many cases, the performance of a machine learning or AI algorithm can only be as good as the data it's given.



Data is most effective when it is representative of some ideal distribution or population. However, a dataset that disproportionately represents some sampled population is said to be **biased**.

And when training data is biased, it then impacts the machine learning algorithms and creates a **bias in machine learning** as well.

# Bias in Facial Recognition Algorithms

[Recent studies](#) that test facial recognition algorithms have shown huge variation in performance. One found that many facial recognition algorithms are "10 to 100 times more likely to inaccurately identify a photograph of a black or East Asian face, compared with a white one". Even among the best algorithms there are notable differences in recognition performance across race and gender.

One factor that could lead to this difference is **bias in the data used to train the algorithms**. An [analysis](#) showed that two popular training sets were overwhelmingly composed of lighter-skinned subjects. This is supported by the above study, which showed that algorithms developed in Asian countries performed better on Asian faces.





# Bias in Other AI Algorithms

Sometimes, bias can also arise when the underlying distribution is skewed (in other words, **systemic bias** prevalent in a culture or society).

For example, bias has caused problems in [algorithms for determining bail](#), which have shown bias in predicting a person's likelihood to commit future crimes 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**.



## 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.)

# Misinformation

Questions about AI and responsibility extend to the algorithms social media networks use to decide which posts should be promoted. Studies have shown these algorithms can lead to [the spread of false information](#).

Furthermore, a user may not even be aware of what data is being used to make a recommendation.

# Responsibility and AI

It is important to consider all the ways in which a machine learning algorithm can make things easier for some while making things more difficult for others.

When algorithms affect certain communities disproportionately, those communities should be able to have a say in those decisions.


Reflecting on ethical responsibilities, remember to ask yourself: “What should we value broadly as a society? And what should we do to realize those values?”

# AI Algorithm Design Activity

# AI Algorithm Design Activity

Suppose that you are part of a team that is deciding whether to adopt an AI hiring algorithm for your company to choose potential job candidates.

- What questions might you have for the designers?
- What concerns might you have about the data used by the algorithm?
- How will you ensure the adoption of the technology is not causing harm or reinforcing existing structural inequalities?





# Learning Goals

- Critically reflect on individual ethical concerns in relation to Computer Science
- Highlight examples of ethics-related instances in the real world
- Understand the significance of ethics from the perspective of your own academic disciplines and desired careers
- Engage in discussions of ethical questions with peers