Sample Classroom Exercise

Collaborative filtering demo (from Introduction to Ethics of Artificial Intelligence)
Explanation:
Systems that recommend movies/books/search results (like Netflix, Amazon, and Google) often use a machine learning algorithm called *collaborative filtering*. The way it works is by recommending things to you that other users who are judged to be similar to you liked. Similarity of users can be calculated by comparing what you liked to what other people liked. What a user "likes" can be measured several ways: 5 star ratings, purchases, watching a video to the end, clicking "like", etc.

Instructions:
We'll start by building a dataset. We need some data to train our recommender system on, and some data to test it on to see how well it works.
Everybody write down whether you liked the following movies (Yes/No/blank):

Training Set:	**Test Set:**
Deadpool	Avengers: Endgame
La La Land	Green Book
Blade Runner 2049	The Secret Life of Pets
Finding Dory	Bohemian Rhapsody
Moana	Fantastic Beasts and Where to Find them
Ghostbusters	Arrival
Star Wars: The Last Jedi	Black Panther
Snowden	A Star is Born
Jurrasic World	Wonder Woman
Pitch Perfect	Minions

We need a volunteer to be the user looking for recommendations. Write your ratings for the Training Set on the blackboard.
Is there anyone in the class who agreed perfectly with all of those ratings? (Ignore any that either of you marked as blank.) You're the closest match. If we're confident that our system works well, we could recommend the movies you liked to the user (and vice versa). But first, we'll check whether the system is working well.
Write down your ratings for the Test Set on the blackboard. Check whether the original user agrees with all of those ratings.
Repeat with another volunteer or two.

Discussion Questions:
Notice that the system doesn't pay any attention to the content of the media. It just matches people. Does this seem like a good way of figuring out what to recommend?
What kinds of things do you think could go wrong with this kind of recommendation system? Have you ever bought a gift, looked something up, or let someone else use your computer or phone, then started getting recommendations that aren't relevant to you?
How well do you think collaborative filtering would work if:

- you have an unusual combination of interests?
- you hate anything popular?
- you are only interested in brand new, cutting edge things?

Can an algorithm be biased?