Engaging Students through Cases

By John Orlando(https://www.teachingprofessor.com/author/john-orlando-phd/) | **June 3, 2024**

In recent years, the education world has been focusing on ways to generate student engagement in learning. This is because of the well-established fact that people learn better when they are actively thinking about the material as they learn it rather than just hoping to passively absorb it. To this end there has been a steady stream of articles on how to use games, audience response systems, interactive quizzes, and the like during instruction.

While all these technologies are helpful, they are effective only if the content itself is engaging. Students need to have some interest in the content to think about it. If the content is not engaging, activities become rote that can produce only limited improvements in learning.

One of the best methods for making content engaging is to teach through cases. College teachers often start by discussing theoretical principles and then illustrate those principles with examples. But a better method is to start by asking students how they would resolve interesting cases and then draw the theoretical lessons from their answers. We are more invested in a topic when it answers our questions than when we're just listening to a discussion. Books, movies, and TED Talks never start with an overview of the topic to be covered. Instead, they open with something that grabs the audience's attention and build lessons from it. I use the following progression of cases involving AI to teach both critical reasoning and deeper ethical principles.

Whom to hit?

I start class with the following case:

> Imagine that in the future self-driving cars become common. A self-driving car is going down the road when three teenagers chasing a drone accidently run in front of it. The car cannot stop in time. It has only three options:
>
> 1. Hit the three teenagers.
> 2. Swerve left to hit two people walking on the sidewalk.
> 3. Swerve right to hit a tree.
>
> What should the car be programmed to do and why?

Most students will say that the car should hit the tree because killing the driver is less harm than killing two or three people. A second response is that the car should hit the three people because they are responsible for the situation, and others should not have to pay for their lack of attention.

I use these responses to show how they are drawn from two fundamental moral intuitions:

1. Consequentialism: that act is right that produces the best outcomes.
2. Moral responsibility: people should be held responsible for their actions.

Two bikes

I then use a **second example (https://youtu.be/ixIoDYVfKA0?si=XYoXnw9RNt346X6r)** to dive deeper into the principles:

> A car is going down a wet road and loses traction. The driver can steer it only to the right or left. On the left is a biker who is wearing a helmet, while on the right is a biker who is not wearing a helmet. Which way should the driver steer the car?

Some argue that the car should hit the biker with the helmet on grounds that they are more likely to survive, but others argue that it should hit the biker without a helmet because the conscientious biker should not have to pay for the other biker's carelessness. This further illustrates our intuitions concerning moral responsibility and consequentialism.

In the line of fire

Now I provide the following example:

> Imagine that someone is about to shoot you, and you duck so that the bullet misses you, knowing that it will hit someone behind you. Have you acted wrongly? Now imagine that you instead grab the person behind you and put them in front of you as a human shield to absorb the bullet. Have you acted wrongly?

There is universal agreement that the act of ducking is ethically permissible but putting someone in front of you is not. This is because we believe that a person has a right to defend their life by avoiding a harm but does not have a right to use an innocent life as a means of protecting their life.

Now I ask how this belief applies to the response that the car should hit the tree. It suggests that self-driving cars should be programmed to protect the driver as the driver has a right to save their own life and so should not hit the tree. Then, given the two options left, it could use consequentialism to hit the two people rather than three people. We now have a more nuanced, two-step model of ethical decision-making for self-driving cars.

Self-driving cars also lead to the question of the extent to which humans should cede control to computers. Here I start by asking whether students would want to live in a world where all cars are self-driving; most say no on grounds that computers can make mistakes. But thousands of people die every year on US roads, mostly due to driver error, such as impaired or distracted driving. Computers do not get drunk or tired or text. If computer-driven cars would lead to fewer accidents (likely the case), then someone interested in their safety would want computer-driven cars on the road.

But some people will still respond that they don't want to give up control to a computer. This leads to the topic of human control and AI.

AI respirators

> Imagine that AI respirators adjust patient oxygen levels with a 95 percent accuracy, while doctors have an 80 percent accuracy. Thus, we will save more lives and have better outcomes by telling doctors to let the AI respirators make oxygen and refrain from overriding them, even when they think that the AI respirator is

> making a mistake.

This scenario demonstrates how we can be better off not allowing ourselves to override the decisions of an AI system—an initially counterintuitive result. But we in fact use this reasoning in many situations. For instance, a couple that tattoos their names on each other does so precisely because the tattoos are permanent, which demonstrates their commitment to one another. The inability to override the decision is its very appeal.

Dr. Strangelove

Here is another case:

> In the movie *Dr. Strangelove*, the Soviet Union has created a doomsday device that will destroy the world if it detects an attack. It cannot be overridden, even by its builders. Dr. Strangelove explains that the inability to override it is critical to its deterrence purpose as an enemy cannot hope that the Soviet Union will not respond to an attack due to weakness of will.

Students generally think that this situation is unethical due to the possible consequences, but I point out that it draws on an underlying negotiation principle that is frequently used by asking the next question.

Protesters

> Imagine that you own a munitions plant that ships ammunition out of the plant by railroad. Protesters are sitting on the tracks to block the shipments. How can you get them to move?

Students generally say that you can slowly drive the train up to the protesters in a game of chicken, hoping to scare them off. But eventually someone will come up with the more effective method of getting the train rolling and jumping off so that nobody, including yourself, can stop it. Then there is no question of will and so no purpose to staying on the tracks.

This case illustrates the principle that the most powerful negotiating tactic is to take the decision-making out of your hands. Good negotiators use it all the time. For instance, a couple can use it to play good cop, bad cop with car dealers. One person's role is to accept the dealer's offer, while the other's is to absolutely refuse anything less than,

say, $300 below the offer and walk out. The remaining person says that they can't make a decision without the other's approval and so their hands are tied.

Teaching through cases can be done in nearly any field. The fascinating **Numberphile (https://youtu.be/p-xa-3V5KO8?si=_akJSdGvVqWT6N8S)** YouTube channel uses them to teach deep math principles to nonmathematicians. There is an art to finding interesting cases that lead to lively discussion, but starting with cases vastly improves student interest and retention.