

MIT SLOAN SCHOOL OF MANAGEMENT

MIT COMPUTER SCIENCE AND ARTIFICIAL INTELLIGENCE LABORATORY (CSAIL)

ARTIFICIAL INTELLIGENCE: IMPLICATIONS FOR BUSINESS STRATEGY

ONLINE SHORT COURSE

MODULE 1 UNIT 3
Video 2 Transcript

MIT AI M1 U3 Video 2 Transcript

TOM MALONE: The next role that machines can play is that of peer. In this case, the machines do things very much like the things that people do. In fact, in some cases, the machines can solve a whole case or problem by themselves, but people often have to solve some or even many of the cases that machines can't solve.

For instance, there's an insurance company called Lemonade, where customers can use mobile apps to file insurance claims. If the claim is within a set of parameters that the system recognizes, the AI system can automatically pay the claim within seconds, but if it's unusual in any way, the claim gets referred to a human adjuster.

One of my favorite examples of computers acting as peers comes from a research project I did with a student here named Yiftach Nagar. We were trying to help make better predictions of things like when a competitor might introduce a new product or when a terrorist might attack, or various other kinds of things. But we wanted to use simpler example, so we tried to predict what would be the next play in an American football game, whether it would be a run or a pass. To do that, we showed people pictures of a football game and stopped the video just before each play began.

We asked people to predict which the next play would be, run or pass, in a format called a prediction market, where people buy and sell predictions of what events would happen. We also had some simple computer agents making the same predictions, based on simplified information about the game, like what yard line the ball was on and how many yards to the next first down. In that case, the computer agents were also participating in a separate prediction market with other computer agents.

And, finally, we had one condition where people and computer agents were both participating in the same prediction market, buying and selling predictions about whether the next play would be a run or a pass. And they didn't even know whether the last sale was made by a person or a computer. What we found was that the combination of people and computers together was both more accurate at making these predictions and more robust to various kinds of errors. So, having people and computers participate as peers in this way led to better results than either could have done alone.

The last role that computers can play relative to people is that of manager. Now, some people worry when they hear about a machine being a manager, but, in fact, we already do it all the time. For example, when a traffic light directs traffic, rather than a human police officer, that machine is essentially managing the people who are driving cars through the intersection, but we don't find anything wrong with that. And, I think, there're, in fact, many cases when machines, computers, can manage people in ways that will be completely unobjectionable. For instance, managers' roles include things like assigning tasks to people and evaluating and training people. And there're certainly ways machines can help with both of those kinds of things.

One of my favorite examples of computers operating as managers comes from a research system called CrowdForge, developed by researchers at Carnegie Mellon University. They were using this CrowdForge system to get a group of people to write things like

encyclopedia articles. For instance, one of the main examples was writing an encyclopedia article about New York City.

To do this, they recruited people using the Amazon Mechanical Turk online labor market, where people do what's called micro tasks, very small things, often only a few minutes, paying only a few cents. In this case, the researchers had one group of people create outlines for an encyclopedia article about New York City with sections like what are the attractions, what are, the history of New York, et cetera.

Then they had for each section in these outlines, they asked other people to generate facts or find facts that were relevant to that section. Then for each section they took all the facts that people had generated and gave those to still other people who wrote coherent paragraphs using those facts. Then the machine at the end put together the paragraphs for the different sections of the article. And that was the final article.

When they had other observers rate the quality of these articles, they rated them as better in quality than articles written by a single person for the same amount of money and roughly equal in quality to the articles in the simple English version of Wikipedia.

So, here's a case where the humans were doing all the work of gathering information and writing the articles, but the machine was playing a critical role in coordinating that work by assigning different tasks to different people. In this case, the machine's coordination was very straightforward. Perhaps you wouldn't even want to call it AI, but researchers are also working with much more sophisticated ways of deciding who will do which tasks based on how well they've done similar tasks in the past.

So, I think it's easy to imagine a future in which artificial intelligence systems like this will do a very good job of managing some of the workflow and task assignment in organizations in a way that no one will find particularly objectionable.

Another role managers often play is evaluating people. An example of that would be the Cogito system, about which you'll hear more later in this course. Cogito is a system that helps telephone representatives improve their interactions with customers on the phone by giving them feedback about the emotional quality of the conversations they've had and are having. The Cogito system doesn't begin to know enough to actually carry on the conversation by itself, but by listening to various subtle characteristics of tone of voice and timing and so forth, it can give telephone reps very useful feedback about how well they're synchronizing and relating with the people they're talking to on the phone and, thus, help them improve their interpersonal skills.

So, we've talked about how to divide tasks between people and computers. The other key question we need to think about is how the human-computer system can improve over time. In the past in business, for instance, when people did business process re-engineering in the 1990s, they often thought of automating some of the tasks in an organization as a major project. And once you had finished doing that, you didn't want to think about making any more changes for quite a long time. But with AI, as the computer programs improve and, in many cases, learn from their own experience, the trade-off between what people and computers can do will keep changing.

So, increasingly, we'll need to think of the combination of people and computers as an ever-evolving system, learning from experience to be better and better over time. So, in fact, one way of describing this is as a cyber-human learning loop. We need to think of the human-computer system as one that gets better and better as it learns from experience, time after time. There are at least three ways this can happen.

One is: Humans figure out how to do better the things that they are already doing. An example of this would be the kind of things people did in business process redesign where they realize some of the steps are no longer needed and eliminate them. Another way this can happen is that programmers improve the machines. For instance, the Google programmers are constantly tweaking and improving the Google search algorithms, same with the Facebook programmers with many of the Facebook algorithms.

A third way this can happen is that the machines learn from their own experience with various kinds of machine learning. For instance, we talked about the KLM chatbot earlier where humans review and sometimes correct the suggested responses that come from the AI programs. When they do that, the AI programs are also learning so that they are more likely to be able to make the response the human would have made when they see a similar situation in the future.

So, in conclusion, in this course we'll see many examples of different ways of combining people and computers. And in order to make effective strategic use of AI to create more intelligent organizations, we need to think carefully about how to divide the tasks between people and computers at any given time, and how to improve the whole process with experience.

Did you understand all the concepts covered in this video? If you'd like to go over any of the sections again, please click on the relevant button.