

MIT SLOAN SCHOOL OF MANAGEMENT

MIT COMPUTER SCIENCE AND ARTIFICIAL INTELLIGENCE LABORATORY (CSAIL)

ARTIFICIAL INTELLIGENCE: IMPLICATIONS FOR BUSINESS STRATEGY

ONLINE SHORT COURSE

MODULE 3 UNIT 1
Video 2 Transcript

Module 3 Unit 1 Video 2 Transcript

REGINA BARZILAY: Let me just describe in general where natural language processing today is used in industry. So search comes to mind, this will be maybe the main application. Information extraction, where the machines take free text and build databases from it. Machine translation. Text generation, which is used today to take databases of events and translate them into a text, like for instance if you need to do some regulatory filings, which are very formulaic, machines can do it for you. And sentiment analysis, where machines can assess the sentiment of the stories or reviews.

The point that I hopefully illustrated in the previous slides, that machines can do very well on the tasks which are perceived as very hard by humans, and can do a really miserable job of the things which are very simple for humans. And it's very important for all of us to keep in mind what is real capacity when you decide to use certain natural language processing tools.

So what you see on this slide is Hathaway, an actress, and Warren Buffett, who owns Berkshire Hathaway. So some time back, a blogger noticed when the actress Hathaway has a movie, the stocks of Warren Buffett's Berkshire Hathaway go up, and you can clearly see that automatic traders that use the sentiment near the word Hathaway misrepresented and misunderstood the reference to the actress Hathaway versus Berkshire Hathaway.

One can even, you know, utilize it to bet on stocks, but this is just one example of when things can go wrong. So whenever you decide which technology you use, you should make sure that it actually provides you the accuracy that you need to deliver reasonable performance.

The next part relates to the question that many of you may be asking, is why we're still not doing a fantastic job in natural language processing. And here you can see some examples of sentences like, hospitals are sued by seven foot doctors, or Iraqi head seeks arms. All of these sentences, the reason that they are so funny is because they are ambiguous. You can have an appropriate reading of the sentence: ban on nude dancing on governor's desk – versus that somebody is banned from dancing nude on governor's desk. This is ambiguous for us when we are reading it, when a machine reads a sentence it doesn't select just between the two interpretations, it can be hundreds and sometimes thousands of possible interpretations.

So I want to illustrate to you different types of ambiguity, using a sentence that was actually an ad from a computer company. So the ad said: "At last, a computer that understands you like your mother." First of all, you can see that if you are looking at the syntactic structure you can parse it in different ways. You can parse it that the computer understands you like your mother does, versus understands that you like your mother. Those are two different interpretations which results in different syntactic representations.

There is also ambiguity at the semantics level. So let's look again at the same sentence, talking about mothers, and you can say okay, mother – let's select the word which is the

most unambiguous, which is mother. So there is one definition of mother that you will find in the dictionary, which is female parent. Another interpretation of the same word is a stringy, slimy substance consisting of yeast cells and bacteria. Clearly for any human, it is clear that we are talking here about the female parent, but the machine has to do this kind of word-sense disambiguation to generate the correct interpretation.

And finally, there is also ambiguity when you stick many sentences together. If you have the sentence, say, Alice says that they've built a computer that understands you like your mother, but she doesn't know any details. Versus Alice says they've built a computer that understands you like your mother, but she doesn't understand me at all. You can see that she here refers to different subjects, and this is called anaphora resolution, so a machine, when it sees he, she, it, needs all the time to decide what exactly do we mean by it.

And finally, there is a huge ambiguity related to human languages. So for instance, the trivial task in English is to decide what is a word, you can say it just has spaces around it and commas and so on – in Semitic languages, like Hebrew and Arabic, you can actually put several words together into a single word, and as a result if you want to do the search or translation you first have to decide that the word in the country has to be broken into three units. So there are a lot of difficulties which are very much language-specific and make the task more difficult.

THOMAS MALONE: 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.