

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 4 Transcript**

## Module 3 Unit 1 Video 4 Transcript

THOMAS MALONE (TM): So, Regina, I thought you did a great job of presenting in your presentation about the different kinds of natural language processing that systems can do and what's easy and what's hard. But I think one of the things our students will be most interested in is getting a better sense of what actually is easy, and what's relatively hard and what's completely impossible, using today's technology for NLP. So, I want to ask you to delve a little more deeply in that. If students have some idea of something they want to do, some kind of application whether it's, you know, reading legal documents or figuring out how to reply to customers' emails or whatever it is, what are some of the things they should think about to know whether this is something that they could do with today's AI or something that they probably couldn't do?

REGINA BARZILAY (RB): Thank you, Tom, for your kind words. And, this is a great question, because every practitioner who tries to take, you know, the whole wealth of existing NLP algorithms and apply them, they have exactly the same question. Can they do it? What is plausible? And there are several considerations that have to be taken into account. The first one is what kind of tolerance your customer has for system mistakes. And, let's be truthful, most of NLP systems do generate mistakes. But sometimes we can still do useful things despite these mistakes and maybe, the most prominent example is Google Search. Even though we get a lot of irrelevant information it's very easy to go through it and find, you know, the links that you care about. Or machine translation. If you don't speak any Japanese and you need to understand the document, even though the translation is sometimes not fluent and ungrammatical, you can still get the gist of the document and perform your task. However, if your goal is to take a legal document where we really care to translate correctly every single word, clearly the existing state of the art in machine translation doesn't allow us to do so because it will translate nontrivial number of words incorrectly. Maybe you can use it as an aid for human translator but, on its own, it's not sufficient. So, understanding the capabilities of the NLP system versus the expectation of the customer is really important here.

The other important component when deciding where we can use natural language processing system is understanding what kind of training data do you have. And by training data I mean that, let's say, you want to extract sentences that describe customer complaints from, you know, the long emails that you get. And if somebody went through this document and, you know, highlighted you those sentences this is called training data. And if you have a large corpus of those this would be an easy task to solve. But if now you are, let's say, trying to use the same technology to solve some related task, but not exactly this task, let's say you want to use it for Amazon review and find places where people talk about drawback of some electronics product, it will not work. Because the distribution of the words the machine used when it used to identify, let's say customer complain for cable company, is very different from the way people will talk about the drawbacks of the product. So, you need to have a dedicated annotation for the product, which will be clearly different from the one in the cable company. And creating this kind of training examples is

sometimes very expensive, but it's really important to have it and to have access to them, so that you can design your machine learning and it can produce you reasonable results.

TM: So, I think what you're saying here is that the techniques that our students learned in the last module about machine learning are essentially at the heart of what happens with natural language processing, as well. Is that right?

RB: It's very correct. And I hope that in the NLP module they get the same idea, that pretty much everything that is done today in NLP and all the exciting products that they are using on a daily basis, like news recommendation, like machine translation, it is all machine learning. But it's a branch of machine learning that is tailored for NLP, and it can take into account very unique properties of natural language inputs and outputs.

TM: So, there's some additional software that takes advantage of things like the structure of a sentence, that it has a subject and a predicate, and things like that.

RB: So, you just design certain methods, which could capitalize on the properties. They're typically either wired in the method itself, so, like machine translation system doesn't occur directly in other contexts in machine learning. So, they were particular models that were designed, they called them neural machine translation models sequence to sequence, which they sequence in English, and then generate the sequence in French. But everything is machine learned.

TM: Okay. And in order for the machine learning to work there have to be certain features that the system can recognize as the inputs, and other kinds of features that the system knows to give as outputs. And I think what you're implying is that humans have to figure out which features are worth paying attention to.

RB: It's absolutely correct. So, it is really an imperative if you know something and you know that it makes a difference, is to include it into machine learning system. So, for instance, if you're trying to assess, you know, the reason for customer complaint and you have a text and do the most simple thing, you just throw in the words and you predict, let's say you have a class of complaints and of overcharging or malfunction or slow service and it may identify it. Now, let's say, you know that customers who have certain plan typically complain more about certain things. Machine learning model will not figure it out unless you put the information about their plan. Or if you know that customers who complain about overcharging typically have much longer complaints than customer that complain about malfunction. If you don't put the lengths of the complaint as a feature, again the machine would not figure it out. So, whatever knowledge you have, you give it as input to the model. And some of them can be weak cues, and not necessarily translate to the final prediction, but the machine will be able to figure it out, even if the correlation is weak.

TM: So, basically you need to have people who put in all the things that might be relevant, to tell the machine to look at these things. The machine can then figure out which ones actually are relevant and which ones aren't, but you need the people to do that initial level

of coding. And then I think you're also saying you need people to say what the right answer is in enough cases for the machine to learn the patterns.

RB: Exactly correct. So, for the second problem of just doing annotation, many companies today can use Mechanical Turk or other, you know, annotation services because you need, you know, maybe millions of these annotations. So, it's not worth to have an expert, sitting and annotating by hand. In most cases you can use, you know, the general annotation service and they will do a good job.

TM: So, actually, just Mechanical Turk, by the way, is an online labor market where people can hire small amounts of work from people all over the world.

RB: Yeah, and it dramatically reduces the cost of annotation. You cannot do it for every single task, because sometimes you actually need to have domain expert which understand, if this is legal documents or something like this, or there is obviously the issue of privacy. But you could typically find yourself a group of people who would annotate it. They don't need to know anything about machine learning. They need to be able to read the document and to provide you the labels. Sometimes you can be very lucky, and you actually don't need to get this kind of people, because maybe if your customer-support service, each time when they get a complaint they classify it by hand then you can just, you know, either you use what you already have, or you would pay somebody to do it. But, then, the second question is which input to put in which form. You really need to have natural language processing or machine learning expert which would identify, talking to the experts in the field, what kind of information to put and how to code it in an optimal way so that the algorithm can learn from it.

TM: So, you need people with different kinds and levels of expertise, to train the systems to do this?

RB: Absolutely.