


The evolution of computational linguistics and where it's headed next

31 May 2017, by Andrew Myers



Natural language processing is a key part of the artificial intelligence revolution. Credit: Stefani Billings

Earlier this year, Christopher Manning, a Stanford professor of computer science and of linguistics, was named the Thomas M. Siebel Professor in Machine Learning, thanks to a gift from the Thomas and Stacey Siebel Foundation.

Manning specializes in [natural language processing](#) – designing computer algorithms that can understand [meaning](#) and sentiment in written and spoken [language](#) and respond intelligently. His work is closely tied to the sort of voice-activated systems found in smartphones and in online applications that translate text between human languages. He relies on an offshoot of [artificial intelligence](#) known as [deep learning](#) to design algorithms that can teach themselves to understand meaning and adapt to new or evolving uses of language.

Siebel, a pioneer in numerous areas of information technology and known for his ability to see and understand emerging trends in computer science and beyond, has long held an interest in precisely this kind of work. "AI and machine learning change everything about computing. Advances in natural language processing change everything about human-computer interaction, and Professor

Manning's innovative natural language processing research is making some of those advancements possible," said Siebel, CEO of C3 IoT.

We spoke to Manning recently about his lifetime pursuit that is only now starting to gain wider public awareness.

How do you explain your work to a complete stranger?

My focus is on natural language processing, otherwise known as computational linguistics. It's getting computer systems to respond intelligently to textual material and human languages. I focus on what the words mean. If you say something to Google or Siri or Alexa, the speech recognition is incredibly good, but half the time it doesn't understand what you mean and asks, "Would you like me to do a web search for that?"

What got you interested in language?

When I was younger, I was totally interested in computer science and artificial intelligence but I also found human languages fascinating with their grammatical structures and complexity. I was interested in the linguistics and the study of language itself, but also in the computational side, too. That side appeals to the kind of people who are more interested in how sentences are put together to create meaning and sentiment and how metaphors arise.

How has the field evolved over the years?

In the early days, the field involved writing out symbolic rules of grammar – subject, noun, clauses, predicates with a verb, perhaps followed by a noun phrase, perhaps followed by a prepositional phrase and so forth. People worked for years trying to replicate grammar and lexicons. It worked fine in very small contexts, but never really extended to understanding meaning.

Then in the '90s the first revolution came when masses of language became available online, digitally. It was then that people started to explore statistical methods of analyzing all that data and building probabilistic models of which words are likely to appear together to create meaning and sentiment.

The grammatical rules have very little to do with communication, what it means to sound natural. You could walk up to someone and say, "Good morning, how are you this morning?" It's perfectly correct, but no one says that. They say, "Hey, how's it going?" Most language is comprised of these softer decisions as to how people use the language.

In natural language processing, you need the computer to understand the world. How do you do that? Well, a very good way is through the enormous amount that has already been written about the world. Every day, writers across the planet are writing about our world and how it all works. We create computational models that assign mathematical values to words and groups of words and use them to successfully read text and derive meaning.

Can natural language processing adapt to new uses of language, to new words or to slang?

That was one of the big limitations of early approaches. Words tend to pick up different usages and even meanings over time, often very remarkably. The word "terrific" used to have a highly negative meaning – something that terrifies. Only recently has it become a positive term.

That's one of the areas I've been involved in. Natural language processing, even if trained in the earlier meaning, can look at a word like "terrific" and see the positive context. It picks up on those soft changes in meaning over time. It learns by examining language as it is used in the world.

Would we have to repeat this learning process for every language?

We've worked a lot in machine translation. How can we translate automatically between two different

[human languages](#)? What you learn is that there is a large amount of already translated text lying around that provides context from which we can build probabilistic models to translate new text. Google Translate did this around 2005 or so. It wasn't great, but it helped. You could sort of figure out what a page was about, but you certainly couldn't expect to get good grammatical sentences out of the translation.

Recently, there has been a big leap in translation quality from deep learning or neural machine translation approaches, which have been explored at Stanford and also at the University of Montreal. Google, Microsoft and Baidu now all use neural machine translation systems for their translations but it is still far from a solved problem. I'm sure people will still be working on this in 20 years.

How is your work being incorporated in voice-based systems like Siri, Alexa, and Google Voice?

We call these dialogue agents – programs that try to understand the meaning of your spoken language and to perform tasks based on what you've said. It's easy to create a chatbot that just repeats random stuff that's half-connected to what you asked about, but how do you build dialogue agents that can understand and do tasks that people are asking for?

One way to get there is to build that sort of contextualization we talked about before based on data freely available from sources like Freebase and Wikidata with their structured facts. Unfortunately, the reality is that those structures must be built by hand right now. The dream is a technology that builds them automatically by reading Wikipedia or online newspapers.

Provided by Stanford University

APA citation: The evolution of computational linguistics and where it's headed next (2017, May 31)
retrieved 15 June 2021 from <https://phys.org/news/2017-05-evolution-linguistics.html>

This document is subject to copyright. Apart from any fair dealing for the purpose of private study or research, no part may be reproduced without the written permission. The content is provided for information purposes only.