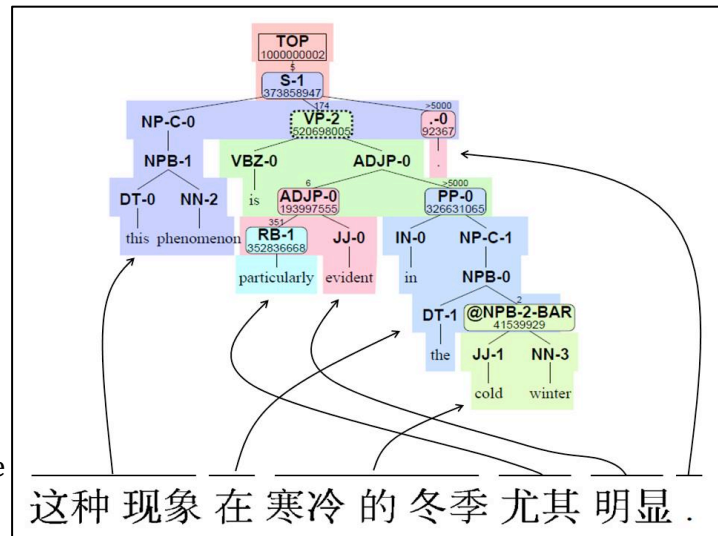


<http://cra.org/ccc/theimpactofnitrd>**Human Language Technology: What Machines Do with Text and Speech****Kevin Knight****USC Information Sciences Institute**

We do many things with human language: we communicate with each other, we store knowledge, we ask questions, we entertain. The ability to use language comes to us naturally. But computers must be programmed to understand and generate human language, and this constitutes a major challenge in artificial intelligence. Sample applications include:

- Automatic translation from one language to another (for example, Chinese to English);
- Question answering, on any topic;
- Automatic dictation for hands-free email/texting; and
- Control of any device by spoken command.



Understanding and generating human language is hard for computers because every word has many meanings. The computer must select the right one based on context, but there are an infinite number of contexts. Even then, selecting the right word meanings is not enough – we must then assemble word meanings correctly to understand whole sentences and texts.

Fifteen years ago, individual users rarely encountered human language technology. Now anyone with web access can use, for example:

- Google Translate, which has its roots in Federally-funded research on statistical machine translation at IBM's T. J. Watson Research Center;
- Online news aggregators, which have their roots in Federally-funded research at Columbia University; and
- Apple Siri, which has its roots in Federally-funded research at SRI International in Menlo Park, CA.

Military, intelligence, and corporate customers make use of an even wider range of new human language technologies, to save money and lives, and to break into new markets.

We have achieved these successes because we stopped worrying about the very hardest linguistic constructions, and instead started worrying about handling the bulk of what computers see and hear. We embraced uncertainty by building probabilistic models of language, and we trained those models on vast data collections.

There are many things we cannot yet do, and extracting meaning from speech and language in the future will enable many new applications.