

## The New York Times

This copy is for your personal, noncommercial use only. You can order presentation-ready copies for distribution to your colleagues, clients or customers, please [click here](#) or use the "Reprints" tool that appears next to any article. Visit [www.nytreprints.com](http://www.nytreprints.com) for samples and additional information. Order a reprint of this article now. »

---

October 4, 2010

SMARTER THAN YOU THINK

# Aiming to Learn as We Do, a Machine Teaches Itself

By STEVE LOHR

Give a computer a task that can be crisply defined — win at chess, predict the weather — and the machine bests humans nearly every time. Yet when problems are nuanced or ambiguous, or require combining varied sources of information, computers are no match for human intelligence.

Few challenges in computing loom larger than unraveling semantics, understanding the meaning of language. One reason is that the meaning of words and phrases hinges not only on their context, but also on background knowledge that humans learn over years, day after day.

Since the start of the year, a team of researchers at [Carnegie Mellon University](#) — supported by grants from the [Defense Advanced Research Projects Agency](#) and [Google](#), and tapping into a research supercomputing cluster provided by [Yahoo](#) — has been fine-tuning a computer system that is trying to master semantics by learning more like a human. Its beating hardware heart is a sleek, silver-gray computer — calculating 24 hours a day, seven days a week — that resides in a basement computer center at the university, in Pittsburgh. The computer was primed by the researchers with some basic knowledge in various categories and set loose on the Web with a mission to teach itself.

“For all the advances in computer science, we still don’t have a computer that can learn as humans do, cumulatively, over the long term,” said the team’s leader, Tom M. Mitchell, a computer scientist and chairman of the machine learning department.

The [Never-Ending Language Learning system](#), or NELL, has made an impressive showing so far. NELL scans hundreds of millions of Web pages for text patterns that it uses to learn facts, 390,000 to date, with an estimated accuracy of 87 percent. These facts are grouped into semantic categories — cities, companies, sports teams, actors, universities, plants and 274 others. The category facts are things like “San Francisco is a city” and “sunflower is a plant.”

NELL also learns facts that are relations between members of two categories. For example, [Peyton Manning](#) is a football player (category). The Indianapolis Colts is a football team (category). By scanning text patterns, NELL can infer with a high probability that Peyton Manning plays for the Indianapolis Colts — even if it has never read that Mr. Manning plays for the Colts. “Plays for” is a relation, and there are 280 kinds of relations. The number of categories and relations has more

than doubled since earlier this year, and will steadily expand.

The learned facts are continuously added to NELL's growing database, which the researchers call a "knowledge base." A larger pool of facts, Dr. Mitchell says, will help refine NELL's learning algorithms so that it finds facts on the Web more accurately and more efficiently over time.

NELL is one project in a widening field of research and investment aimed at enabling computers to better understand the meaning of language. Many of these efforts tap the Web as a rich trove of text to assemble structured ontologies — formal descriptions of concepts and relationships — to help computers mimic human understanding. The ideal has been discussed for years, and more than a decade ago [Sir Tim Berners-Lee](#), who invented the underlying software for the World Wide Web, sketched his vision of a "semantic Web."

Today, ever-faster computers, an explosion of Web data and improved software techniques are opening the door to rapid progress. Scientists at universities, government labs, Google, [Microsoft](#), [I.B.M.](#) and elsewhere are pursuing breakthroughs, along somewhat different paths.

For example, I.B.M.'s "question answering" machine, Watson, shows remarkable semantic understanding in fields like history, literature and sports as it plays the quiz show "Jeopardy!" Google Squared, a research project at the Internet search giant, demonstrates ample grasp of semantic categories as it finds and presents information from around the Web on search topics like "U.S. presidents" and "cheeses."

Still, artificial intelligence experts agree that the Carnegie Mellon approach is innovative. Many semantic learning systems, they note, are more passive learners, largely hand-crafted by human programmers, while [NELL is highly automated](#). "What's exciting and significant about it is the continuous learning, as if NELL is exercising curiosity on its own, with little human help," said Oren Etzioni, a computer scientist at the [University of Washington](#), who leads a project called TextRunner, which reads the Web to extract facts.

Computers that understand language, experts say, promise a big payoff someday. The potential applications range from smarter search (supplying natural-language answers to search queries, not just links to Web pages) to virtual personal assistants that can reply to questions in specific disciplines or activities like health, education, travel and shopping.

"The technology is really maturing, and will increasingly be used to gain understanding," said Alfred Spector, vice president of research for Google. "We're on the verge now in this semantic world."

With NELL, the researchers built a base of knowledge, seeding each kind of category or relation with 10 to 15 examples that are true. In the category for emotions, for example: "Anger is an emotion." "Bliss is an emotion." And about a dozen more.

Then NELL gets to work. Its tools include programs that extract and classify text phrases from the

Web, programs that look for patterns and correlations, and programs that learn rules. For example, when the computer system reads the phrase “Pikes Peak,” it studies the structure — two words, each beginning with a capital letter, and the last word is Peak. That structure alone might make it probable that Pikes Peak is a mountain. But NELL also reads in several ways. It will mine for text phrases that surround Pikes Peak and similar noun phrases repeatedly. For example, “I climbed XXX.”

NELL, Dr. Mitchell explains, is designed to be able to grapple with words in different contexts, by deploying a hierarchy of rules to resolve ambiguity. This kind of nuanced judgment tends to flummox computers. “But as it turns out, a system like this works much better if you force it to learn many things, hundreds at once,” he said.

For example, the text-phrase structure “I climbed XXX” very often occurs with a mountain. But when NELL reads, “I climbed stairs,” it has previously learned with great certainty that “stairs” belongs to the category “building part.” “It self-corrects when it has more information, as it learns more,” Dr. Mitchell explained.

NELL, he says, is just getting under way, and its growing knowledge base of facts and relations is intended as a foundation for improving machine intelligence. Dr. Mitchell offers an example of the kind of knowledge NELL cannot manage today, but may someday. Take two similar sentences, he said. “The girl caught the butterfly with the spots.” And, “The girl caught the butterfly with the net.”

A human reader, he noted, inherently understands that girls hold nets, and girls are not usually spotted. So, in the first sentence, “spots” is associated with “butterfly,” and in the second, “net” with “girl.”

“That’s obvious to a person, but it’s not obvious to a computer,” Dr. Mitchell said. “So much of human language is background knowledge, knowledge accumulated over time. That’s where NELL is headed, and the challenge is how to get that knowledge.”

A helping hand from humans, occasionally, will be part of the answer. For the first six months, NELL ran unassisted. But the research team noticed that while it did well with most categories and relations, its accuracy on about one-fourth of them trailed well behind. Starting in June, the researchers began scanning each category and relation for about five minutes every two weeks. When they find blatant errors, they label and correct them, putting NELL’s learning engine back on track.

When Dr. Mitchell scanned the “baked goods” category recently, he noticed a clear pattern. NELL was at first quite accurate, easily identifying all kinds of pies, breads, cakes and cookies as baked goods. But things went awry after NELL’s noun-phrase classifier decided “Internet cookies” was a baked good. (Its database related to baked goods or the Internet apparently lacked the knowledge to correct the mistake.)

NELL had read the sentence “I deleted my Internet cookies.” So when it read “I deleted my files,” it decided “files” was probably a baked good, too. “It started this whole avalanche of mistakes,” Dr. Mitchell said. He corrected the Internet cookies error and restarted NELL’s bakery education.

His ideal, Dr. Mitchell said, was a computer system that could learn continuously with no need for human assistance. “We’re not there yet,” he said. “But you and I don’t learn in isolation either.”

---

[Copyright 2011 The New York Times Company](#) | [Home](#) | [Privacy Policy](#) | [Search](#) | [Corrections](#) | [XML](#) | [Help](#) | [Contact Us](#)  
| [Back to Top](#)