Artificial Intelligence

It's Not What You Think

BY CHRIS HORN

If you're concerned about robots with artificial intelligence conquering the world anytime soon, just watch a video of the latest RoboCup, the annual international robotics soccer competition. The tournament, which represents current robotic and AI technology, features squads of plodding, humanoid robots whose spastic play is downright laughable.

But don't laugh too hard. The engineers developing these robots have a stated goal of fielding a team that can compete with the World Cup winners — the human ones — by 2050. Artificial intelligence-equipped computers have already surpassed human performance in several arenas: chess, trivia, image recognition and speech recognition.

"A number of people have expressed concern about AI taking over the world. Even Stephen Hawking has expressed that opinion, even the guy who runs Tesla," says **Michael Huhns**, a distinguished professor emeritus of computer science at USC whose research has focused on various aspects of AI. "But I don't think that will be a concern within anyone's lifetime."

A group of scientists from major institutions like MIT, Harvard, Microsoft and Berkeley agree. Last year they wrote the first installment of the One Hundred Year Study on Artificial Intelligence, "Artificial Intelligence and Life in 2030," a crystal ball prognosis of how AI will affect life in a typical North American city. Their verdict?

"The frightening, futurist portrayals of Artificial Intelligence that dominate films and novels, and shape the popular imagination, are fictional. In reality, AI is already changing our daily lives, almost entirely in ways that improve human health, safety and productivity. Unlike in the movies, there is no race of superhuman robots on the horizon or probably even possible."

Moreover, while the scientists recognize that AI has the potential to be abused, they see a much larger upside, with new technologies improving safety, education and quality of life.

Significant advancements in machine learning have surged in the past 20 years, and in many subtle ways humans are already ceding control to those devices, Huhns says. Think of debit card readers, personal assistants such as Siri and Alexa, and the dozens of sensors on modern cars that assist in driving.

Those are mostly innocuous examples of smart machines. But the autonomously driven cars that Google and Tesla have been testing for several years bring a new level of complexity.

"It's impossible to program for every situation, so an autonomously driven car will have to make decisions like we do," Huhns says. "The computer has to decide, 'Do I drive over a cliff and kill my driver to avoid hitting pedestrians in the road?' If there's a 90-year-old and a toddler in the road, which one would the autonomously driven car choose to run over if there was no alternative?"

USC computer science professor **Marco Valtorta** points to the fatal crash of an autonomously driven car this summer in which Tesla's navigation system failed to distinguish a tractor-trailer rig crossing a Florida highway.

"The car's visual recognition system was trained to expect underpasses — which is what the trailer crossing the road might have resembled — more commonly than a semi turning around on a highway," Valtorta says.

The death of the Florida passenger was tragic, of course, but the overall track record of self-driving cars is pretty darned good. And Valtorta and Huhns say the AI technology fueling that trend will improve even more over time. It won't ever be perfect, they say, but the technology doesn't have to be perfect to do far better than human performance — the determining factor in most of the 38,000-plus traffic fatalities last year in the U.S.

But forget about self-driving cars and robots for a moment. How far has AI advanced in really cool stuff such as computers that think like humans?

"We're not so far from computers with personality," Huhns says. "Some can express emotion and sound sympathetic or happy. They can even make sounds of laughter, depending on the sophistication of the programming."

For example, if an AI-programmed computer heard you say the word "funeral" or "cancer," it could respond in a lower, more measured tone, sounding more sympathetic than happy. "There's logic involved, not thinking, per se, but some reasoning ability," Huhns says.

And that's the thing: what we think of as "thinking" is still beyond the grasp of AI. Consider Watson, IBM's computer that famously competed on "Jeopardy!" The computer appeared to be pretty smart, Huhns says, "but if you asked it the same 10 questions, again and again, it would process them the same way 10 times. And if you talked with Watson, each sentence would be a discrete interaction, not the thread of a larger overall conversation."

It boils down to what thinking is. "You can study thinking by modeling the brain, that's what neuroscientists do, or you can study properties of thinking, like aerodynamicists study flight," Valtorta says. "While we don't have a clear understanding of thinking, we do know that context is important, the idea of bringing to bear what we already know to a situation."

An example of that is found in language processing. If someone says, "The laptop didn't fit in the backpack because it was too small," we know that "it" refers to the backpack, not the laptop. That kind of context can be programmed, as well as a sense of probability, Valtorta says. Both are integral components of thinking, but they don't give an AI-equipped computer the full measure of human thinking, with all of its subtle nuances and discernment.

"You can study thinking by modeling the brain, that's what neuroscientists do, or you can study properties of thinking, like aerodynamicists study flight."
- MARCO VALTORTA

"We're not so far from computers with personality. Some can express emotion and sound sympathetic or happy. They can even make sounds of laughter ..."
- MICHAEL HUHNS