'Automated dermatologist' detects skin cancer with expert accuracy

*(CNN)*Even though the phrase "image recognition technologies" conjures visions of high-tech surveillance, these tools may soon be used in medicine more than in spycraft.

A team of Stanford researchers trained a computer to identify images of skin cancer moles and lesions as accurately as a dermatologist, according to a new paper published in the journal Nature.

In the future, this new research suggests, a simple cell phone app may help patients diagnose a skin cancer -- the most common of all cancers in the United States -- for themselves.

"Our objective is to bring the expertise of top-level dermatologists to places

where the dermatologist is not available," said Sebastian Thrun, senior author of the new study, founder of research and development lab Google X and an adjunct professor at Stanford University. He added that those who live in developing countries do not have the same level of care as can be found in the US and other industrialized nations.

Melanomas represent fewer than 5% of all skin malignancies diagnosed in the US, yet they account for nearly three-quarters of all deaths related to this form of cancer. If detected early, the five-year survival rate for melanoma is 99%. When detected in its latest stage, the survival rate plummets to just 14%.

Generally, dermatologists identify whether a mole or other abnormality is cancerous by looking at it. They can confirm their diagnosis with follow-up biopsies and tests.

With a team of researchers, Thrun developed a deep learning computer system to perform the first task in detecting skin cancer: identifying it at a glance.

Essentially, the team created an automated dermatologist.

How it works

Thrun and his colleagues began by coaching a computer to develop pattern recognition skills. The method they used is an algorithm-based technique known as "deep learning."

Specifically, the research team employed a convolutional neural network.

Carl Vondrick, a Ph.D. candidate at MIT's Computer Science and Artificial Intelligence Lab, who was not involved in this study, explained the process.

"A convolutional neural net is a type of computer software that is very good at learning to recognize different concepts," he said. By downloading digital images, researchers can "tell" the computer they are images of skin cancer, or without skin cancer. The machine will basically try to learn some rules that can predict whether it's cancer.

"An algorithm is just a fancy name for a sequence of steps that the computer takes. So in this case, the algorithm refers to the whole process that they did to train the system," Vondrick said.

Andre Esteva, co-first author of the new paper with Brett Kuprel, both electrical engineering Ph.D. students at Stanford, said he, Thrun and their colleagues began by "basically teaching the algorithm what the world looks like."

"We taught it with cats and dogs and tables and chairs and all sorts of normal everyday objects look like," Esteva said. "We used a massive data set of well over a million images." This phase of learning took about a week.

Then, Esteva trained the algorithm in different skin conditions. Here, the team addressed a complex problem: Cancerous and noncancerous skin aberrations vary greatly in appearance from patient to patient.

To overcome this difficulty, the researchers presented the now-trained -- or "artificially intelligent" -- computer with an extensive dataset of 129,450 images representing more than 2,000 skin diseases. The images came from 18 doctor-curated online repositories as well as the Stanford University

Medical Center.

Since each image of a mole or abrasion had been diagnosed, the computer was fed this information as well.

Differences between AI and human intelligence

"What's really powerful about computer vision and these convolutional neural nets is, all you have to do is specify the input and output, and the machine will learn its own rules how to do that automatically," Vondrick said. "It learns kind of a series of mathematical transformations to essentially transform an image into an answer (to the question) 'Is there skin cancer or not?' "

There's an important distinction between computer visual systems and the human visual system, added Vondrick. Though humans can learn to recognize patterns from very little data, machines require thousands or even billions of examples.

"You can grow up to be in your 20s and never have seen skin cancer, and you can go to medical school and just see a few examples on slides, and all of a sudden, you're going to be pretty good at recognizing that," he said.

No computer has been trained to do the same, but computer visual systems may detect subtleties within digital photographs unseen by the human eye, the researchers pointed out.

The moment of truth came when the researchers presented previously unseen images to their algorithm. Would their artificial intelligence system be able to

recognize both the most common and the most deadly types of skin cancer: malignant carcinomas and melanomas, respectively?

"This algorithm performed as well as board-certified dermatologists at several key diagnostic tasks," Esteva said.

Notably, the computer was able to "diagnose multiple different kinds of skin cancer, not just melanoma, and we were able to do this with regular clinical images, rather than with specialized dermoscopic images," said Roberto Novoa, a co-author of the study and a dermatologist at Stanford Medicine.

He explained that doctors commonly use a dermatoscope, a specialized tool, to examine the skin for cancer. This tool enables a view from a similar distance with similar lighting and magnification, for easier diagnosis.

Though the algorithm lacked access to this expensive tool, its performance still equaled the accuracy of 21 dermatologists.

Though Thrun, Esteva and their colleagues warn that real testing in a clinical setting -- a doctor's office -- is still needed, they believe their research might be expanded to include other areas of medicine, such as ophthalmology, radiology and pathology.

With 6.3 billion smartphone subscriptions estimated to be in use by 2021 (PDF), the researchers noted, their new system, in the form of an app, could provide low-cost universal access to diagnostic care.

Paging Dr. McCoy

Science fiction is rife with visions of technology replacing human doctors. Dr. Leonard McCoy used a portable diagnostic device known as a tricorder to determine the medical condition of the USS Enterprise crew members in the "Star Trek" TV series and movies, recalled Dr. Sancy A. Leachman, a

dermatologist and chairwoman of the department of dermatology at Oregon Health and Science University, and Glenn Merlino, a senior investigator at the National Cancer Institute, in their published commentary on the new study.

Though still "fanciful," machines capable of non-invasive diagnosis are "becoming a reality," wrote Leachman and Merlino, neither of whom was involved in the new study.

That said, at least one issue still requires more investigation when it comes to a skin cancer-detecting automated system.

According to Leachman and Merlino, it is not known whether the artificial intelligence system featured in the new study can distinguish between similar-looking diseases. For example, a system would need to be able to identify melanoma versus benign seborrhoeic keratosis -- basically, a non-cancerous wart.

"Even the benign growths can be associated with certain syndromes and diseases that only a physician that has been clinically trained for detecting can diagnose," said Dr. Jill Waibel, a dermatologist and owner of the Miami Dermatology and Laser Institute.

Waibel, who did not participate in the study, added that "medical Imaging has dramatically transformed the practice of medicine," especially her field of dermatology. Yet despite many exciting developments, all of the new imaging systems are still being studied and explored for optimal uses, she said.

For many scientists, though, "Star Trek" still beckons: a future filled with machines that learn with each corrected mistake, improving their performance over time.

"This is a very specific study, and it has a very encouraging result," Thrun said. But he warned that before this research might be leveraged into anything resembling a tricorder, "we would have to run many more studies."