DOI:10.1145/3359332

Vinton G. Cerf

AI Is Not an Excuse!

I keep hearing excuses for not working on difficult problems: “Eventually AI will solve this so there’s no point working on it now.” Sorry, wrong answer.

First, we should be cautious about putting too much expectation on artificial intelligence which, by most metrics, is really today’s machine learning (ML) and neural networks. There is no doubt these systems have produced truly remarkable results. The GO game story of AlphaGo and AlphaZero from Deep Mind is by now a classic example of the surprisingly powerful results these systems produce. A chess-playing version of AlphaZero learned quickly and demonstrated choices of moves unlike those of traditional chess players. I hesitate to label this *strategy* but the deep neural networks do encode in some deep way *experience* that one could consider a kind of strategic cache.

These systems are also quite brittle and can break in ways that are not always, to my understanding, predictable. Rather, we might predict there will be circumstances in which the ML system will fail but we may not know how. It’s a bit like the insurance data that indicates 1% of all males over the age of 85 will die next year—we just don’t know who is in the 1%! It seems prudent, then, to anticipate these fragile possibilities and research how they might be characterized and even identified based on the design of the neural network. I think we know, for example, that image classification systems do not work the way human classification works. The systems are much more sensitive to pixel input than humans. We abstract from raw pixel input, recognizing shapes and characteristics (“pointy ears,” for in-

stance), while the image recognition ML systems are much more sensitive to pixel-level inputs. Changing a small number of pixels can lead to significant misclassifications.

With these frailties in mind, it seems to me very important not to make too many assumptions about the power of machine learning. I do want to acknowledge, however, that considerable successes have been recorded well beyond playing board games. At Google, a ML system was trained to control the cooling system in its datacenters and saved 40% of the power needed to operate them. Machine speech recognition allows Google to converse via its Assistant and to do automatic language translation. While it can be argued the machine does not *understand* what is said, the ML system can process the input and produce useful output (“What’s the weather in Palo Alto?” “It’s 78 degrees in Palo Alto with a

high of 82 and a low of 68”). There are pathogen and disease image recognition systems helping to identify patients at risk. There is a lot to praise and admire about these applications.

In the meantime, however, I think it is not okay to ignore difficult problems on the assumption they will be solved “automagically” by ML tools. We have huge challenges ahead with *ordinary* software that we cannot reasonably assume will be solved by AI. Software analysis for potential mistakes or bugs requires tools I would not identify with traditional AI or ML. Designing systems to be updated reliably with new software doesn’t require AI but it does require careful thinking about authentication of the origin of the software update and confirmation it has retained its integrity during its journey from the source to the updated device. Security in general is not solely the purview of AI or ML. Interestingly, some aspects of security will be addressable such as fraud detection. Credit card companies are making good use of modeling to detect unusual card usage and flag anomalous events for further analysis.

Bottom line: Let’s enthusiastically explore the uses of machine learning and artificial intelligence but not use their potential to excuse ourselves from crafting high-quality, reliable software that is resistant to abuse!

Vinton G. Cerf is vice president and Chief Internet Evangelist at Google. He served as ACM president from 2012–2014.

Copyright held by author.

It seems to me very important not to make too many assumptions about the power of machine learning.