

8 ASPIRATIONAL AI

PJD's Opening Remarks

The AI systems we have discussed so far can be ranked by the kinds of problems they are able to solve. For example, rule-based systems are good at diagnoses but cannot distinguish an image of a dog from an image of a cat. Neural networks can make those distinctions but cannot learn to play Chess or Go. Unsupervised learning systems can, without any external supervision, become world class game players or can classify data into clusters. Human-Machine teams can solve problems that neither human nor machine can solve alone.

In this course, we have chosen “learning power” as the criterion for ranking machines. Each level of machines has greater learning power than the previous level. However, it is also possible to call the distinguishing factor “storage power” because each level has a more sophisticated way of storing and retrieving information. We use “learning power” because learning is a sign of intelligence and is the most easily quantified of various factors that signal intelligence. But no one calls any of these machines intelligent.

For many years, AI researchers has aspired to build machines that perform intelligently. This has always been an elusive goal because every success, however impressive, did not strike anyone as intelligent. The successes came from well-designed machines doing their jobs well. But there is much more to do. Many AI researchers continue to pursue the goal of intelligent machines.

One of the most talked about aspirations – and fears – is a super-intelligent machine. Ray Kurzweil, an AI pioneer, argues that the continued exponential growth of computing power will lead, by 2030 or 2040, to machines with all the complexities of a brain in the physical space of brain. These machines will develop consciousness and intelligence. Because no one knows what concerns those machines will have, or how they will regard homo sapiens, the event at which a machine becomes conscious is called The Singularity. It is impossible to stare into a singularity and see what lies within.

Today Professor Neil Rowe will talk about the aspirations of AI and what appears possible or likely. He is a pioneer in the AI field and has been teaching at NPS for over 30 years. In 1988 her authored a book on PROLOG, perhaps the most famous language for building expert systems. He knows more about AI than anyone on campus and probably in the universe.