

## 1 A HIERARCHY OF AI MACHINES

PJD's opening remarks

There is a large number of AI machines, software and hardware. They have all sorts of functions and capabilities. How do we make sense of it all?

When I was a young computer scientist, I learned about a hierarchy of computers postulated by Noam Chomsky, the famous MIT linguist. He spoke of three main levels of computing machines, each more powerful than the last.

The simplest one was "finite state automata", universally used for computer logic circuits. The next one was "pushdown automata", universally used in compilers to translate programs into machine code. The highest was the Turing Machine, a very general model of computing invented by the British mathematician and code-breaker in 1936. The motto of the field soon became "Whatever can be computed, can be computed by a Turing machine."

Recently, in trying to sort through AI hype, my colleague Ted Lewis and I wondered if we could do the same for AI machines. (Ted, now retired, is former CS chair and Director of the NPS Center for Homeland Security and Defense.). A hierarchy would enable people to understand the power and limitations of each kind of AI machine.

The question was: how to rate their power?

Rating by computing power was already done by Chomsky, and most AI machines can be jiggered into Turing machines anyway. Since all AI machines are touted as machines that can learn, we decided to use learning power as a criterion for distinguishing AI machines into a small hierarchy of categories.

The first section of this course is a series of 8 lectures on the different categories of machines, characterized by learning power. I will introduce them to you today.

The professors in the coming lectures will talk to you in more depth about these machines and what we can use them for.