Chapter 2: The Intellectual History of Artificial Intelligence

Top Ten Salient Sentence Strings

1. "The speeds and memory capacities of present computers may be insufficient to simulate many of the higher functions of the human brain, but the major obstacle is not lack of machine capacity, but our ability to write programs taking full advantage of what we have… Probably a truly intelligent machine will carry out activities which may best be describes as self improvement… A fairly attractive and yet clearly incomplete conjecture that the difference between creative thinking and unimaginative competent thinking lies in the injection of some randomness. The randomness must be guided by intuition to be efficient. In other words, the educated guess or the hunch include controlled randomness in otherwise orderly thinking." All these somewhat off-the-cuff remarks presaged important areas of study within the field.

2. But perhaps the most remarkable, albeit overlooked, result of the Dartmouth proposal is the improbable and most likely unintentional success of the term *artificial intelligence* in attracting interest and attention far beyond its academic roots. Nothing in McCarthy's life suggests that he harbored a hidden interest or talent for coining brilliant marketing slogans, yet his choice of this particular moniker has sparked an enduring fascination by the press, public, and entertainment media – an achievement that alludes all but the most accomplished advertising professionals.

3. Today, expert systems are no longer considered an active area of research within AI, much less an investment opportunity, for a number of reasons. Foremost among them is that dramatic increases in computer power, storage, and networking have led to an explosion of data in readily accessible electronic form, which opened the door to a completely different approach to incorporating expertise into computer programs – one that eliminated the need to painstakingly encode the knowledge and skills of a human practitioner by hand.

4. Heuristic reasoning tackles a common, if not universal, problem plaguing the symbolic systems approach – that the number of possible sequences of steps can be very large (called a "combinatorial explosion"), so you can't simply examine all options, as discussed in chapter 1 with respect to the game of chess.

5. We know a fair amount about the gross structure of the brain – which layers and regions are typically involved in various activities, such as seeing, getting hungry, doing arithmetic, adjusting your heart rate, recognizing faces, and wiggling your big toe. But surprisingly little is understood about the intermediate structure – how the neurons are connected to perform these tasks. In other words, we don't know much about how the brain is wired (metaphorically speaking). And of course this is precisely the area of interest to AI researchers building artificial neural networks.

6. Perhaps the most remarkable application of machine learning systems is some recent work in which the techniques are used not to simulate the brain but to reverse engineer it. A group of scientists led by Jack Gallant at the Henry H. Wheeler Jr. Brain Imaging Center of the University of California at Berkeley is succeeding in using machine learning techniques to read minds. Really. The researchers train a machine learning system to look for patterns in an array of brain sensors while they show test subjects pictures of various objects. Then they put a new subject into the test rig and show him or her a picture. Once trained, their program can correctly identify what the subject is looking at with significant accuracy.

7. The plain fact is that symbolic reasoning and machine learning have different strengths and weaknesses. In general, symbolic reasoning is more appropriate for problems that require abstract reasoning, while machine learning is better for situations that require sensory perception or extracting patterns from noisy data. For instance, suppose you want to build a robot that can ride a bike. Representing this problem in symbolic terms may be possible, but imagine trying to interview a human expert in an effort to build an expert system to do this. There certainly are experts at riding bikes, but the nature of their expertise simply doesn't lend itself to description in words. Clearly, knowledge and expertise can take forms that resist codification into human language or any explicitly symbolic form.

8. By contrast, using machine learning techniques, this problem is a ride in the park, so to speak. For a single example, a recent research project by some graduate students at the Georgia Institute of Technology accomplished this task using neural network techniques; the system succeeded in learning stunts such as wheelies, the "bunny hop," front wheel pivot, and back hop. But there are other issues for which machine learning techniques aren't well suited. To state the obvious, machine learning is not useful for problems where there's no data, just some initial conditions, a bunch of constraints, and one shot to get it right.

9. In short, if you have to stare at a problem and think about it, a symbolic reasoning approach is probable more appropriate. If you have to look at lots of examples or play around with the issue to get a "feel" for it, machine learning is likely to be more effective. So why did the focus of work shift from the former to the latter? […] So four trends – improvements in computing speed and memory, the transition from physically to electronically stored data, easier access (mainly due to the Internet), and low-cost high-resolution digital sensors – were prime drivers in refocusing of effort from symbolic reasoning to machine learning.

10. A group of researchers at Google's DeepMind division applied their machine learning algorithms to the ancient game of Go […] Go swamps chess with respect to the number of possible moves, making it resistant to solution by many other AI approaches, such as the heuristic search techniques IBM's Deep Blue mainly used to be Kasparov at chess. The Google program, called AlphaGo, scored a decisive win over Lee Sedol, a top-ranked international Go player, winning 4 out of a 5-game series in South Korea in March of 2016. The win was certainly a significant technical achievement, but what it means for machine intelligence and its relationship to human intelligence is unclear at best. Fei-Fei Li, Director of the Stanford AI Lab, put this well. She was quoted in the *New York Times* as saying "I'm not surprised at all. How come we are not surprised that a car runs faster than the fastest human?"