Imperial College London

Alan Turing and the development of Artificial Intelligence

Stephen Muggleton
Department of Computing
Imperial College, London

August, 2012

Alan Turing (1912-1954)

Theory of Computation (1937)

Development of Automated Decryption (1942)

Foundations of Artificial Intelligence (1950)

Foundations of Systems Biology (1952)

Turing's 1950 paper in *Mind*

A.M. Turing, "Computing Machinery and Intelligence", *Mind*, 59:433-460, 1950.

- First major academic paper on the foundations of Artificial Intelligence
- Suggested Computational Logic as a basis for Artificial Intelligence
- First major academic paper on foundations of Machine Learning
- Describes a 50 year research program for implementing Artificial Intelligence using Logic-based Machine Learning

Structure of Turing's 1950 paper

1. The Imitation Game [1.5 pages]
2. Critique of the New Problem [1.5]
3. The Machines Concerned in the Game [1.2]
4. Digital Computers [4]
5. Universality of Digital Computers [3]
6. Contrary Views on the Main Question [13]
7. Learning Machines [7]

Section 6 of the paper

Quote 1 We may now consider the ground to have been cleared and we are ready to proceed to the debate on our question, “Can machines think?”

Quote 2 It will simplify matters for the reader if I explain first my own beliefs in the matter. ... I believe that in about fifty years’ time it will be possible, to programme computers, with a storage capacity of about 10^9 binary digits, to make them play the imitation game so well that an average interrogator will not have more than 70 per cent chance of making the right identification after five minutes of questioning.

Section 7 - Learning Machines
Version 1: AI by programming

Storage capacity For human-level intelligence less than 10^{15} bits needed. Advances in engineering will be adequate for the requirements. [Moore's law 1964-]

Programming approach to AI. Sufficient storage capacity will be available by 2000 but needs prohibitively large human resources to program. [1960s-1980s AI vs 1992- WWW]
[Machine Intelligence = Reasoning, Perception, Physical action and Programming Languages]

Section 7 - Learning Machines
Version 2: AI by *ab initio* machine learning

***ab initio* machine learning** - the Child Machine - simple program which can learn from rewards and punishment. [1968 Boxes - Michie+Chambers, 1980 Machine Learning - Michalski, Mitchell and Carbonell; 1986 Connectionism - Rumelhart and J.L. McClelland; 1995 SVMs - Vapnik]

Need for more than pos/neg examples .. the amount of information .. does not exceed the total number of rewards and punishments applied. By the time a child has learnt to repeat "Casablanca" he would probably feel very sore indeed.. [1984- COLT - Valiant, Haussler, Kearns, Ehrenfeucht]

Section 7 - Learning Machines
Version 3: AI using logic, probabilities, learning and background knowledge - Inductive Logic Programming (ILP)

Logic-based learning with background knowledge .. one might have a complete system of logical inference "built in." .. the store would be largely occupied with definitions and propositions. The propositions would have various kinds of status, e.g., well-established facts, conjectures, mathematically proved theorems, statements given by an authority, expressions having the logical form of proposition but not a belief-value. [1959 McCarthy - Logic-based common sense, NO LEARNING(ie Turing's Version 1); 1980 Logic Programming - Kowalski, NO LEARNING (ie Version 1); 1991 ILP - Muggleton (Version 3, earlier related work in 1970 Plotkin and 1980 Shapiro)]

Section 7 - Learning Machines
The challenge of “Supercriticality” in AI

Supercriticality - analogy nuclear chain reaction One could say that a man can “inject” an idea into the machine, and that it will respond to a certain extent and then drop into quiescence, like a piano string struck by a hammer. .. An idea presented to a subcritical mind will on average give rise to less than one idea in reply. A smallish proportion are supercritical. An idea presented to such a mind may give rise to a whole “theory” consisting of secondary, tertiary and more remote ideas. Animals’ minds seem to be very definitely subcritical. Adhering to this analogy we ask, “Can a machine be made to be supercritical?”

Conclusion of Turing's 1950 paper

We can only see a short distance ahead, but we can see plenty there that needs to be done.

AAAI archive

- AAAI's Virtual Archive of AI project
- Organised by Bruce Buchanan
- Funded by the AI Jnl Foundation and the NSF
- Online digitised versions of Machine Intelligence series
- Also Teddington Conference
- Aims to contain 1700 titles
- http://aitopics.org/classic-articles-books
- buchanan@cs.pitt.edu