4) On *Evolution* (Advancement) of AI Research

Jiming Liu

Dean of Science & Chair Professor in Computer Science

Hong Kong Baptist University

Era of *Pervasive Artificial Intelligence (AI)*

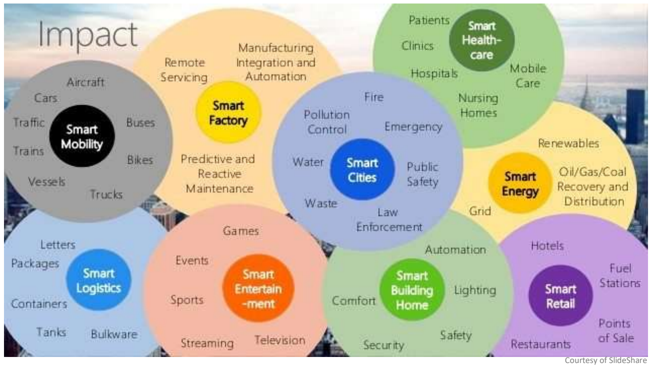

Courtesy of SlideShare

"Computing Machinery and Intelligence" A. Turing

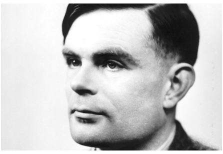

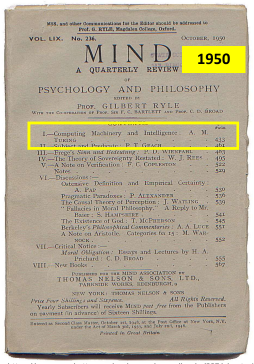

MSS. and other Communications for the Editor should be addressed to Prof. G. RYLE, Magdalen College, Oxford.

VOL. LIX. No. 236. OCTOBER, 1950

MIND

A QUARTERLY REVIEW OF PSYCHOLOGY AND PHILOSOPHY

EDITED BY PROF. GILBERT RYLE

WITH THE CO-OPERATION OF PROF. SIR F. C. BARTLETT AND PROF. C. D. BROAD

CONTENTS.	PAGE
I.—Computing Machinery and Intelligence: A. M. Turing	433
II.—Subject and Predicate: P. T. Geach	461
III.—Frege's *Sinn und Bedeutung*: P. D. Wienpahl	483
IV.—The Theory of Sovereignty Restated: W. J. Rees	495
V.—A Note on Verification: F. C. Copleston	522
Notes	529
VI.—Discussions:—	
Ostensive Definition and Empirical Certainty: A. Pap	530
Pragmatic Paradoxes: P. Alexander	536
The Causal Theory of Perception: J. Watling	539
"Fallacies in Moral Philosophy." A Reply to Mr. Baier: S. Hampshire	541
The Existence of God: T. McPherson	545
Berkeley's *Philosophical Commentaries*: A. A. Luce	551
A Note on Aristotle. Categories 6a 15: M. Warnock	552
VII.—Critical Notice:—	
Moral Obligation: Essays and Lectures by H. A. Prichard: C. D. Broad	555
VIII.—New Books	567

PUBLISHED FOR THE MIND ASSOCIATION BY THOMAS NELSON & SONS, LTD., PARKSIDE WORKS, EDINBURGH, 9

NEW YORK: THOMAS NELSON & SONS

Price Four Shillings and Sixpence. *All Rights Reserved.*

Yearly Subscribers will receive MIND *post free* from the Publishers on payment (in advance) of Sixteen Shillings.

Entered as Second Class Matter, October 1st, 1948, at the Post Office at New York, N.Y., under the Act of March 3rd, 1933, and July 2nd, 1946.

Printed in Great Britain

https://www.manhattanrarebooks.com/pages/books/227/alan-turing/computing-machinery-and-intelligence/?soldItem=true

- "We may hope that machines **will eventually** **compete with men in *all purely intellectual* fields**. But which are the best ones to start with? Even this is a difficult decision. Many people think that a very abstract activity, like the **playing of chess**, would be best."

- "I believe that in **about *fifty years*' time** it will be possible, to programme computers, with a **storage capacity of about *10^9*** [one gigabyte], to make them play the imitation game so well that an average **interrogator will not have more than *70%* chance of making the right identification** after five minutes of questioning."

Q1: So far, how much has AI accomplished? [L]

Q1: So far, how much has AI accomplished? [L]

A1: A lot, but with the exceptions of some *most challenging* ones [M]

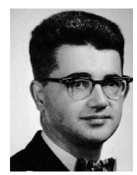

A 1955 Workshop Proposal by:
- **John McCarthy, Dartmouth**
- **Claude E. Shannon**, Bell
- **Marvin L. Minsky**, Harvard
- **Nathaniel Rochester**, IBM

Topics:
1. Automatic Computers
2. How Can a Computer be Programmed to Use a Language
3. Neuron Nets
4. Theory of the Size of a Calculation
5. **Self-Improvement**
6. **Abstractions**
7. **Randomness and Creativity**

A Proposal for the

DARTMOUTH SUMMER RESEARCH PROJECT ON ARTIFICIAL INTELLIGENCE

June 17 - Aug. 16

We propose that a 2 month, 10 man study of artificial intelligence be carried out during the summer of 1956 at Dartmouth College in Hanover, New Hampshire. The study is to proceed on the basis of the conjecture that every aspect of learning or any other feature of intelligence can in principle be so precisely described that a machine can be made to simulate it. An attempt will be made to find how to make machines use language, form abstractions and concepts, solve kinds of problems now reserved for humans, and improve themselves. We think that a significant advance can be made in one or more of these problems if a carefully selected group of scientists work on it together for a summer.

The following are some aspects of the artificial intelligence problem:

1) Automatic Computers

If a machine can do a job, then an automatic calculator can be programmed to simulate the machine. 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 inability to write programs taking full advantage of what we have.

2) How Can a Computer be Programmed to Use a Language

It may be speculated that a large part of human thought consists of manipulating words according to rules of reasoning

https://www.bnext.com.tw/article/47176/ai-turing-test-the-brief-history-of-intelligence
https://www.semanticscholar.org/paper/The-Dartmouth-College-Artificial-Intelligence-The-Moor/

… aspired to develop "Strong AI"
(or Artificial *General* Intelligence)

"Our problem then is to find out **how to programme these machines to play the game.** At my **present rate of working** I produce about a thousand digits of programme a day, so that about ***sixty* workers**, working steadily through the ***fifty* years** might accomplish the job, if nothing went into the wastepaper basket. Some more **expeditious method** seems desirable."

Alan Turing, 1950

https://www.bnext.com.tw/article/47176/ai-turing-test-the-brief-history-of-intelligence
https://www.semanticscholar.org/paper/The-Dartmouth-College-Artificial-Intelligence-The-Moor/

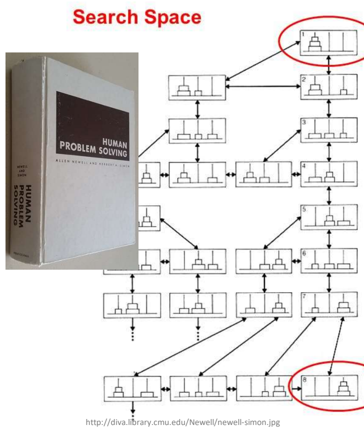

Solution

Problem Solving is *a search problem*

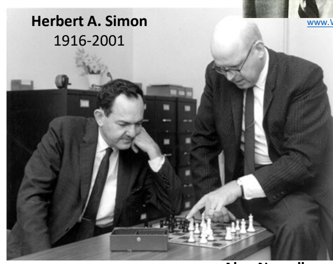

Herbert A. Simon 1916-2001

Alan Newell 1927-1992

Satisficing

From Wikipedia, the free encyclopedia

Satisficing is a decision-making strategy or cognitive heuristic that entails searching through the available alternatives until an acceptability threshold is met.[1] The term *satisficing*, a portmanteau of *satisfy* and *suffice*,[2] was introduced by Herbert A. Simon in 1956,[3] although the concept was first posited in his 1947 book *Administrative Behavior*.[4][5] Simon used satisficing to explain the behavior of decision makers under circumstances in which an optimal solution cannot be determined. He maintained that many natural problems are characterized by computational intractability or a lack of information, both of which preclude the use of mathematical optimization procedures. He observed in his Nobel Prize in Economics speech that "decision makers can satisfice either by finding optimum solutions for a simplified world, or by finding satisfactory solutions for a more realistic world. Neither approach, in general, dominates the other, and both have continued to co-exist in the world of management science".[6]

Simon formulated the concept within a novel approach to rationality, which posits that rational choice theory is an unrealistic description of human decision processes and calls for psychological realism. He referred to this approach as bounded rationality. Some consequentialist theories in moral philosophy use the concept of satisficing in the same sense, though most call for optimization instead.

Q2: How to model *human cognition*? [L]

Herbert A. Simon

Previous (Herb Brooks)

Herbert Alexander Simon (June 15, 1916 – February 9, 2001) was an American researcher in the fields of cognitive psychology, computer science, public administration, economic sociology, and philosophy (sometimes described as a "polymath"). In 1978, he received the Nobel Prize in Economics for his pioneering research into the decision-making process within economic organizations. His later work involved artificial intelligence, developing computer simulations of problem-solving. Simon was able to develop more complex models of economic decision-making by bringing psychological concepts into play, thus leading to models that more closely resembled human social behavior.

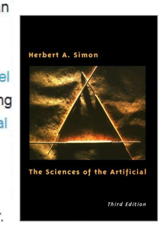

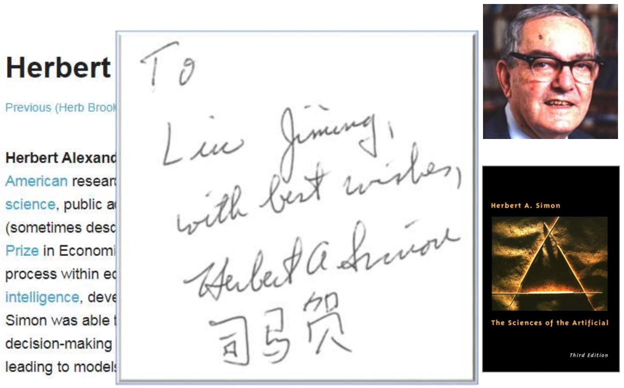

Q2: How to model *human* ***cognition***? [L]

A2: *Means-End approach* (state-space) lies the key [S]

"Playing of Chess"
as envisioned in 1950's

... has now been "achieved" via
Deep Learning

19 Oct. 2017: *DeepMind* Introduced AlphaGo Zero

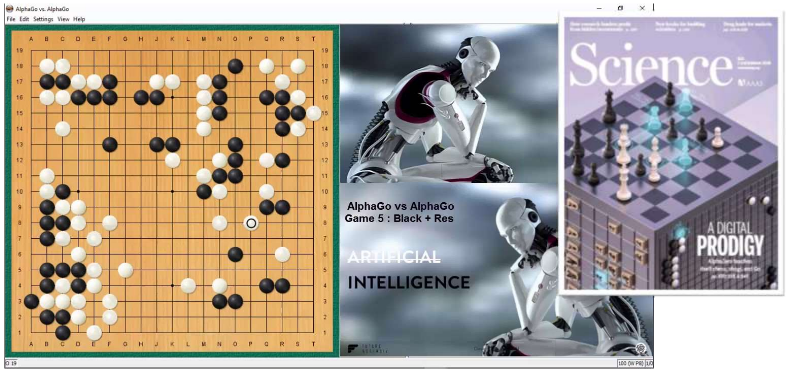

David Silver, *et al.*, "Mastering the game of Go without knowledge," *Nature*, 19 Oct. 2017, pp.354-359.

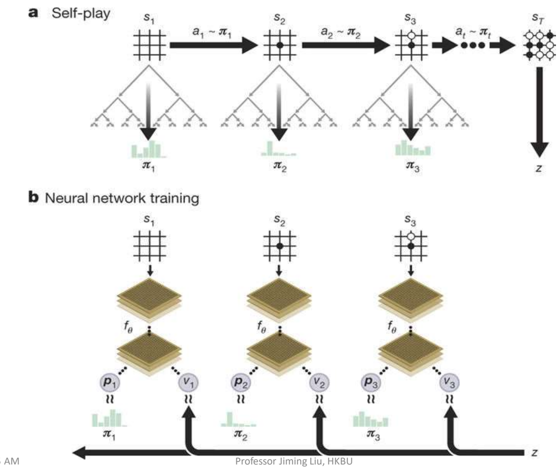

Early *Related* Developments

Knowledge-Based AI

Knowledge-Based Systems

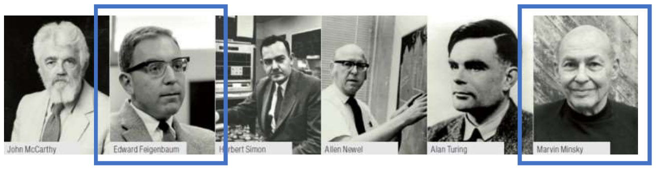

https://businessworld.cz/cio-bw-special/prukopnici-informacniho-veku-john-mccarthy-8728

- Knowledge-representation
  - attempt to capture *expertise* of human experts
  - build knowledge-based systems, more powerful than just algorithms and code
  - "*In the knowledge lies the power*" (Ed Feigenbaum, Turing Award: 1994)
  - first-order logic
    - $\forall p\ \text{vegetarian}(p) \leftrightarrow (\forall f\ \text{eats}(p,f) \rightarrow \neg\exists m\ \text{meat}(m) \wedge \text{contains}(f,m))$
    - $\forall x,y\ \text{eat}(joe,x) \wedge \text{contains}(x,y) \rightarrow \text{fruit}(y) \vee \text{vegetable}(y)$
    - $\therefore$ vegetarian(joe)
  - inference algorithms
    - satisfiability, entailment, modus ponens, backward-chaining, unification, resolution

http://slideplayer.com/slide/4380716/

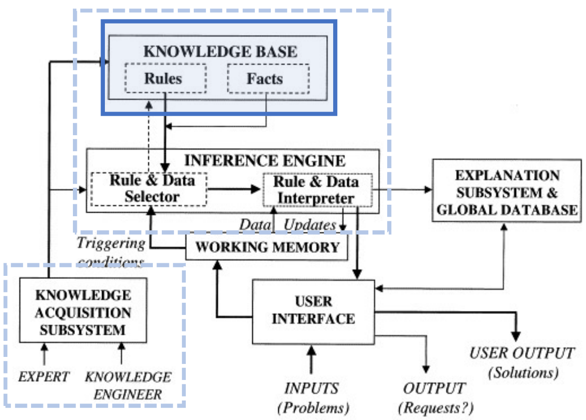

https://www.sciencedirect.com/science/article/abs/pii/S096386870000041X

Intelligence without representation*

Rodney A. Brooks

MIT Artificial Intelligence Laboratory, 545 Technology Square, Rm. 836, Cambridge, MA 02139, USA

Received September 1987

Brooks, R.A., Intelligence without representation, Artificial Intelligence 47 (1991), 139–159.

* This report describes research done at the Artificial Intelligence Laboratory of the Massachusetts Institute of Technology. Suppo[rt for the] research is provided in part by an IBM Faculty 9 Development Award, in part by a grant from the Systems Development Foundation, i[n part by] the University Research Initiative under Office of Naval Research contract N00014-86-K-0685 and in part by the Advanced [Research] Projects Agency under Office of Naval Research contract N00014-85-K-0124.

Abstract

Artificial intelligence research has foundered on the issue of representation. When intelligence is approached in an incremental manner, with strict reliance on interfacing to the real world through perception and action, reliance on representation disappears. In this paper we outline our approach to incrementally building complete intelligent Creatures. The fundamental decomposition of the intelligent system is not into independent information processing units which must interface with each other via representations. Instead, the intelligent system is decomposed into independent and parallel activity producers which all interface directly to the world through perception and action, rather than interface to each other particularly much. The notions of central and peripheral systems evaporateeverything is both central and peripheral. Based on these principles we have built a very successful series of mobile robots which operate without supervision as Creatures in standard office environments.

IBM Deep Blue (1997)

IBM Watson (2011)

"…**I should be surprised if more than 10^9 was required for satisfactory playing of the imitation game**, at any rate against a blind man. (Note: The capacity of the *Encyclopedia Britannica*, 11th edition, is *2 X 10^9*)."

A. Turing, 1950

https://www.zdnet.com/article/wolfram-alpha-vs-ibms-watson-how-they-think/

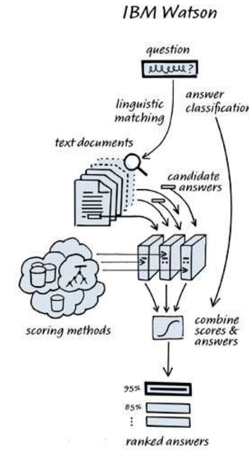

Data-Driven AI

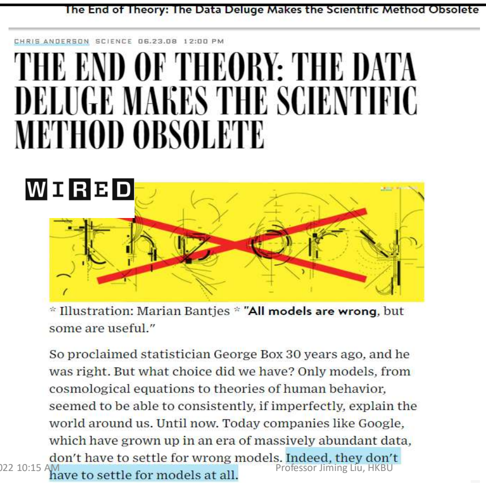

The End of Theory: The Data Deluge Makes the Scientific Method Obsolete

CHRIS ANDERSON SCIENCE 06.23.08 12:00 PM

THE END OF THEORY: THE DATA DELUGE MAKES THE SCIENTIFIC METHOD OBSOLETE

WIRED

* Illustration: Marian Bantjes * **"All models are wrong**, but some are useful."

So proclaimed statistician George Box 30 years ago, and he was right. But what choice did we have? Only models, from cosmological equations to theories of human behavior, seemed to be able to consistently, if imperfectly, explain the world around us. Until now. Today companies like Google, which have grown up in an era of massively abundant data, don't have to settle for wrong models. Indeed, they don't have to settle for models at all.

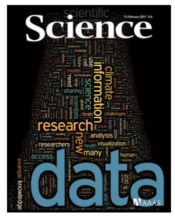

Deep Learning Algorithms

providing lift for classification and forecasting models

Deep Neural Networks

feature extraction and classification of images

Convolutional Neural Networks

for sequence of events, language models, time series, etc.

Recurrent Neural Networks

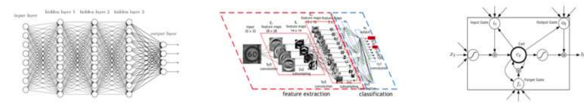

https://mapr.com/blog/demystifying-ai-ml-dl/

Future Developments

Data-Driven Systems

e.g., in Healthcare

Machine Learning-Driven *Precision* Cardiology Through *Multiscale Biology* and *Systems Medicine*

- Heterogeneous data sources:
  - experimental evidence, bioinformatics databases,
  - lifestyle measurements, electronic health records,
  - environmental influences,
  - biobank findings
- Aim: To identify **causal disease networks**, **stratify patients**, and ultimately predict more **efficacious therapies**.

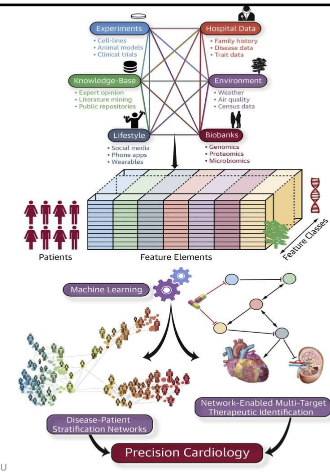

Johnson, K.W. et al., J Am Coll Cardiol Basic Trans Science, 2017;2(3):311-27.

Deep Explanation

Deep Explanation

- Aim: To develop modified/hybrid deep learning **<u>techniques</u>** that learn more
  - **explainable features,**
  - **explainable representations, or**
  - **explanation generation facilities**

towards "Applied AI"

... BUT?

"Computing Machinery and Intelligence" A. Turing

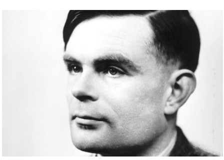

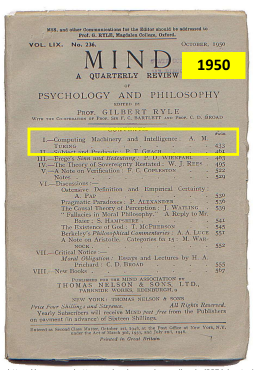

MSS. and other Communications for the Editor should be addressed to Prof. G. RYLE, Magdalen College, Oxford.

VOL. LIX. No. 236. OCTOBER, 1950

MIND

A QUARTERLY REVIEW OF PSYCHOLOGY AND PHILOSOPHY

EDITED BY PROF. GILBERT RYLE

WITH THE CO-OPERATION OF PROF. SIR F. C. BARTLETT AND PROF. C. D. BROAD

PUBLISHED FOR THE MIND ASSOCIATION BY THOMAS NELSON & SONS, LTD., PARKSIDE WORKS, EDINBURGH, 9

NEW YORK: THOMAS NELSON & SONS

Price Four Shillings and Sixpence. *All Rights Reserved.*

Yearly Subscribers will receive MIND *post free* from the Publishers on payment (in advance) of Sixteen Shillings.

Entered as Second Class Matter, October 1st, 1948, at the Post Office at New York, N.Y. under the Act of March 3rd, 1933, and July 2nd, 1946.

Printed in Great Britain

https://www.manhattanrarebooks.com/pages/books/227/alan-turing/computing-machinery-and-intelligence/?soldItem=true

- "We may hope that machines **will eventually compete with men in all purely intellectual fields**."
- "The original question, **"Can machines think?"** I believe to be too meaningless to deserve discussion."
- "**By observing the results of its own behaviour it can modify its own programmes so as to achieve some purpose more effectively.** These are possibilities of the near future, rather than Utopian dreams."

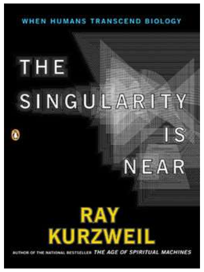

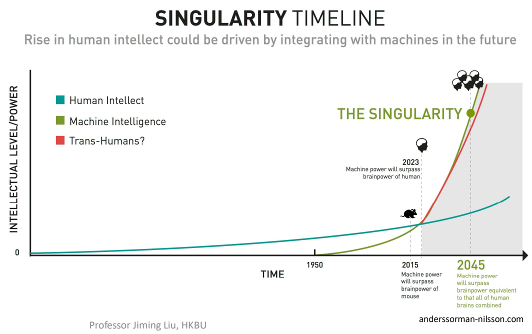

9/16/2022 10:15 AM Professor Jiming Liu, HKBU

Is "Strong AI" *near* ?

9/16/2022 10:15 AM Professor Jiming Liu, HKBU

Q3: Any perspectives on future AI? [L]

9/16/2022 10:15 AM Professor Jiming Liu, HKBU

Knowledge-Based Systems

2018 J. Liu
https://businessworld.cz/cio-bw-special/prukopnici-informacniho-veku-john-mccarthy-8728

Turing Keynote
Chaired by Prof. Jiming Liu
(9:30a.m.-10:20a.m., December 5, 2012)

Edward Feigenbaum,
Professor at Stanford University, USA
(1994 Turing Award Winner)

Extreme Ture-Ing In The Spectrums Of AI

Abstract

"I became obsessed with the Turing machine… I had no access to a scientific library. I looked up in my English-Greek dictionary the verb "to ture"… In the end, I managed to piece together the puzzle, how a man named Alan Turing crafted his machine in order to answer the paramount question of his time: what can be computed, and what cannot."

Christos H. Papadimitriou
Communications of the ACM
Vol. 55 No. 9, (August 2012) Page 42

In this talk, we "ture" some abstractions (or "spectrums") of the work of AI scientists in the first Turing Century, looking for points that have been little explored, but might have great impact and value in the second Turing Century.

The landscape to be "tured" contains these three spectral dimensions: Cognitive -> Perceptual; Knowledge -> Search; and What -> How. On our "ture" we will glimpse ranges of Applications, peaks of Creativity, and unexpectedly large watersheds of Learning and rivers of Expertise. At the end of this "ture", we should have more insight into, and power over, how we and our robots compute "what can be computed".

Biography

Edward Feigenbaum was born in Weehawken, New Jersey, in 1936. He holds a B.S. (1956) and Ph.D. (1960), both from Carnegie Mellon University. His dissertation was supervised by legendary computer pioneer Herb Simon and explored a pioneering computer simulation of human learning.
Feigenbaum is a pioneer in the field of artificial intelligence and is often known as "the father of expert systems." He founded the Knowledge Systems Laboratory at Stanford University and is currently a professor emeritus of computer science there.
Feigenbaum joined the Stanford computer science faculty in 1965 as one of its founding members. That same year, he and Nobel laureate Joshua Lederberg started the DENDRAL project. Later joined by eminent chemist Carl Djerassi and others, this project produced the world's first expert system (1965-1982). DENDRAL's groundbreaking accomplishments inspired an evolution of expert systems, moving artificial intelligence out of the laboratory and into the structure of countless software applications. As important, it changed the frame-

Turing Keynote
Chaired by Prof. Jiming Liu
(9:30a.m.-10:20a.m., December 5, 2012)

Edward Feigenbaum,
Professor at Stanford University, USA
(1994 Turing Award Winner)

Extreme Ture-Ing In The Spectrums Of AI

Abstract

"I became obsessed with the Turing machine… I had no access to a scientific library. I looked up in my English-Greek dictionary the verb "to ture"… In the end, I managed to piece together the puzzle, how a man named Alan Turing crafted his machine in order to answer the paramount question of his time: what can be computed, and what cannot."

Christos H. Papadimitriou
Communications of the ACM
Vol. 55 No. 9, (August 2012) Page 42

In this talk, we "ture" some abstractions (or "spectrums") of the work of AI scientists in the first Turing Century, looking for points that have been little explored, but might have great impact and value in the second Turing Century.

The landscape to be "tured" contains these three spectral dimensions: Cognitive -> Perceptual; Knowledge -> Search; and What -> How. On our "ture" we will glimpse ranges of Applications, peaks of Creativity, and unexpectedly large watersheds of Learning and rivers of Expertise. At the end of this "ture", we should have more insight into, and power over, how we and our robots compute "what can be computed".

Q3: Any perspectives on future AI? [L]

A3: Spectrums of AI [F]

Herbert A. Simon

Previous (Herb Brooks)

Herbert Alexander Simon (June 15, 1916 – February 9, 2001) was an American researcher in the fields of cognitive psychology, computer science, public administration, economic sociology, and philosophy (sometimes described as a "polymath"). In 1978, he received the Nobel Prize in Economics for his pioneering research into the decision-making process within economic organizations. His later work involved artificial intelligence, developing computer simulations of problem-solving. Simon was able to develop more complex models of economic decision-making by bringing psychological concepts into play, thus leading to models that more closely resembled human social behavior.

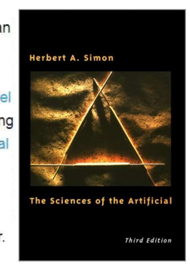

Contents

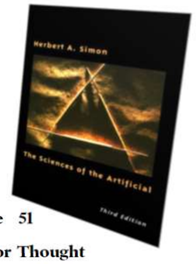

THE ARCHITECTURE OF COMPLEXITY

HERBERT A. SIMON*

Professor of Administration, Carnegie Institute of Technology

(Read April 26, 1962)

A NUMBER of proposals have been advanced in recent years for the development of "general systems theory" which, abstracting from properties peculiar to physical, biological, or social systems, would be applicable to all of them.[1] We might well feel that, while the goal is laudable, systems of such diverse kinds could hardly be expected to have any nontrivial properties in common. Metaphor and analogy can be helpful, or they can be misleading. All depends on whether the similarities the metaphor captures are significant or superficial.

and to analyze adaptiveness in terms of the theory of selective information.[3] The ideas of feedback and information provide a frame of reference for viewing a wide range of situations, just as do the ideas of evolution, of relativism, of axiomatic method, and of operationalism.

In this paper I should like to report on some things we have been learning about particular kinds of complex systems encountered in the behavioral sciences. The developments I shall discuss arose in the context of specific phenomena, but the theoretical formulations themselves make

STATE DESCRIPTIONS AND PROCESS DESCRIPTIONS

"A circle is the locus of all points equidistant from a given point." "To construct a circle, rotate a compass with one arm fixed until the other arm has returned to its starting point." It is implicit in Euclid that if you carry out the process specified in the second sentence, you will produce an object that satisfies the definition of the first. The first sentence is a state description of a circle, the second a process description.

These two modes of apprehending structure are the warp and weft of our experience. Pictures, blueprints, most diagrams, chemical structural formulae are state descriptions. Recipes, differential equations, equations for chemical reactions are process descriptions. The former characterize the world as sensed; they provide the criteria for identifying objects, often by modeling the objects themselves. The latter characterize the world as acted upon; they provide the means for producing or generating objects having the desired characteristics.

The distinction between the world as sensed and the world as acted upon defines the basic condition for the survival of adaptive organisms. The organism must develop correlations between goals in the sensed world and actions in the world of process. When they are made conscious and verbalized, these correlations correspond to what we usually call means-end analysis. Given a desired state of affairs and an existing state of affairs, the task of an adaptive organism is to find the differ-

ceeded. The solution is genuinely new to us—and we do not need Plato's theory of remembering to explain how we recognize it.

There is now a growing body of evidence that the activity called human problem solving is basically a form of means-end analysis that aims at discovering a process description of the path that leads to a desired goal. The general paradigm is: given a blueprint, to find the corresponding recipe. Much of the activity of science is an application of that paradigm: given the description of some natural phenomena, to find the differential equations for processes that will produce the phenomena.

THE DESCRIPTION OF COMPLEXITY IN SELF-REPRODUCING SYSTEMS

The problem of finding relatively simple descriptions for complex systems is of interest not only for an understanding of human knowledge of the world but also for an explanation of how a complex system can reproduce itself. In my discussion of the evolution of complex systems, I touched only briefly on the role of self-reproduction.

Atoms of high atomic weight and complex inorganic molecules are witnesses to the fact that the evolution of complexity does not imply self-reproduction. If evolution of complexity from simplicity is sufficiently probable, it will occur repeatedly; the statistical equilibrium of the system will find a large fraction of the elementary particles participating in complex systems.

- The **dynamic properties** of hierarchically organized systems, which can be ***decomposed*** into subsystems in order to analyze their behavior.
- The **relation** between complex systems and their ***descriptions*** (states vs. process).

On *Evolution of AI Research*

Evolving AI

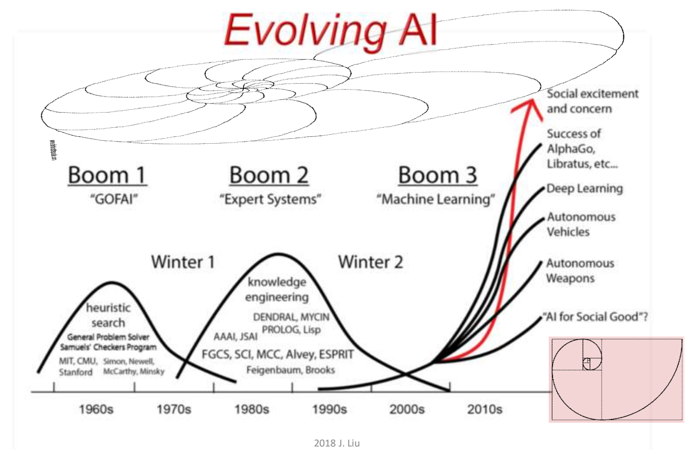

http://www.technologystories.org/ai-evolution/ 2018 J. Liu

On *Evolution of AI Research*: The Beginning of a NEW CHAPTER

- "**AI** is thought to be an *impossible* dream by many. But ***not*** to us in AI. …the future promises to generate orders of magnitude greater impact than its progress to date."

Raj Reddy
(1994 Turing Award)

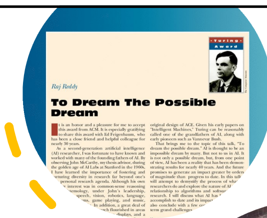
Raj Reddy
To Dream The Possible Dream

Enjoy Your **AI** Journey !

Prof. Jiming Liu *jiming@comp.hkbu.edu.hk*