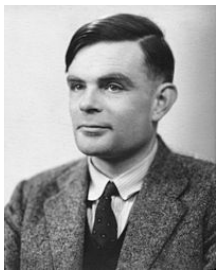

Turing's Imitation Game
Computing Machinery and Intelligence

Can Machines Think?

1) The Question: Can machines think?
    i) If defined naturally by what we deem "machine" and "think" to mean, then it could be determined by a survey such as a Gallup poll.
2) The Imitation Game
    i) Three players: Man (A), Woman (B), Interrogator (C)
    ii) Interrogator is in a separate room and may only ask questions (in written form or through an intermediary) of the Man and the Woman.
        (a) The interrogator only knows them as X and Y
        (b) Ex: Will X please tell me the length of his or her hair?
            1. If X is the man (A), then he must answer, but his objective is to cause the interrogator to make the wrong identification
            2. He may answer "My hair is shingled, and the longest strands are about nine inches long."
        (c) The third player's objective (B) is to help the interrogator
            1. These hints may not be helpful, since the other player (A) may also make similar remarks
            2. Ex: "I am the woman, don't listen to him!"
    iii) At the end of the game, the interrogator must determine which of X and Y are the man and woman.
3) The Question (revised): What will happen when a machine takes the part of A in the imitation game?
    i) Will the interrogator decide wrongly as often when trying to decide between man and machine as between man and woman?
    ii) Discussion: Is this question equivalent to the original? Why? Why not?
        (a) Turing's critique:
            1. Draws a line between physical and intellectual capacities of humans (we're just now—2010s—able to 3-D print skin)
            2. Little point in dressing up a "thinking machine" like a human in artificial flesh (no beauty competitions)

3. Just keep the interrogator from physical interactions with the players in the game
4. Q&A format can bring up any facets of the "human endeavor" such as intellectual discourse, poetry, game-playing, etc.
5. Witnesses can brag about strengths or charms, but cannot demonstrate them for the interrogator

iii) Man vs Machine seems unfair
  (a) Man would make a poor imitator of machine
    1. Slow
    2. Inaccurate in arithmetic
  (b) Machine could "fake" it
    1. Machine could do something which should be described as thinking, but yet is very different from what the man does.
    2. Pretend to be slow
    3. "…if, nevertheless, a machine can be constructed to play the imitation game satisfactorily, we need not be troubled by this objection." (is this a fair assumption?)

iv) Turing ignores the notion that a machine's best strategy is to do something other than imitate true thinking of a man, and assumes that the best strategy is to "try to provide answers that would naturally be given by a man.
  (a) Is this a good assumption?
  (b) Are current Turing Tests following this assumption?

4) What is a machine?
  i) Only permit digital computers in the Imitation Game
    (a) No "men born in the usual manner"
    (b) No test tube babies or clones
      1. "it is probably possible to rear a complete individual from a single cell of the skin (say) of a man. To do so would be a feat of biological technique deserving of the very best praise…"
  ii) Are there imaginable computers that would do well in this game?
    (a) Turing does not suggest trying experiment
  iii) What is a digital computer?
    (a) Replacement for a human computer:
      1. Follows fixed rules from a rule book
      2. Not allowed to deviate from those rules

3. Has an unlimited supply of paper (and his memory)
4. May use a 'desk machine,' aka calculator

(b) Computer in 3 parts
1. Store
    i. The unlimited supply of paper
    ii. The book of rules stored as a table of instructions
    iii. Memory
2. Executive Unit
    i. Performs the calculations
3. Control
    i. Ensures that the instructions in the table are followed correctly and in the right order
    ii. Instructions may be "add the number stored in position 6809 to that in 4302 and put result back into position 4302"
    iii. Instructions may be encoded to fit a specific packet (memory) size, ex: 6809430217
    iv. Instructions may allow for loops and conditionals, ex: 'Now obey the instruction stored in position 5606, and continue from there'
    v. Fun domestic analogy: see page 438

(c) Aside: Programming
1. Human computers don't have rule books, they remember what they must do
2. Finding out what they do and translating that into a set of instructions for a machine to perform is programming

(d) Aside: Digital Computers with Randomness
1. Allowed to cast a die for results
2. Some considered these machines to have free will, but Turing disagrees
3. It is impossible to determine observing one of these machines whether it has a random element or not
4. Infinite capacity (store) computers are allowed (TMs with infinite tape)
5. Electrical machines are allowed
6. Debunk the idea that nervous system is electrical and so is the electrical digital computers

7. Babbage's Analytical Engine was a mechanical computer, so using electricity is not of theoretical importance

iv) Digital Computers are Discrete State Machines
   (1) Discrete states with sudden jumps between them
   (2) Wheel machine example
   (3) Given initial state and input signals, it is always possible to predict all future states
      (a) Laplace's view: given complete state of the universe, it is possible to predict all future states
      (b) Butterfly effect? (p440)
   (4) Manchester Machine

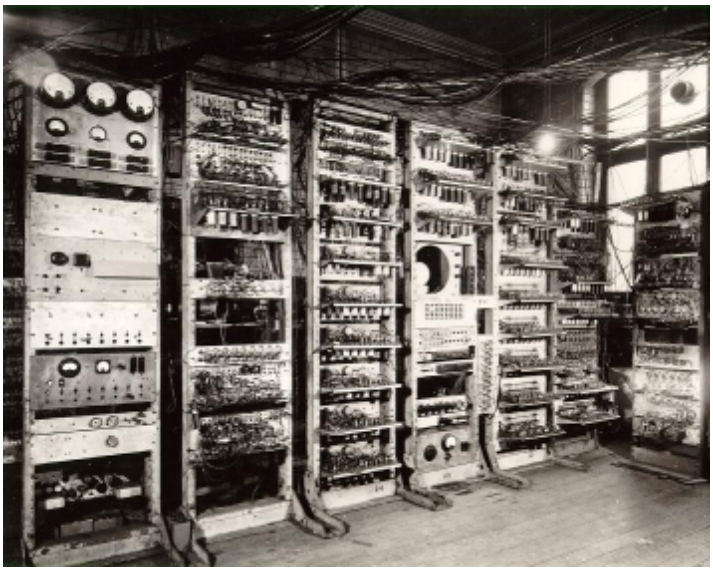

      (a) 4,050 valves and had a power consumption of 25 kilowatts
      (b) Word = 1x40-bit number or 2x20-bit instructions
      (c) Main store: 2 Williams tubes holding 32x40-bit words each (page)
      (d) Magnetic Drum store: 128 pages
      (e) According to Turing: $2^{165000} = 10^{50000}$ states
      (f) Storage capacity is about 165,000
         (i) 100 sheets of paper, 50 lines each, 30 digits per line, writing digits 0-9, the number of states is $10^{100*50*30} = 10^{150000}$
   (5) Universality
      (a) Digital computer can mimic behavior of any discrete state machine given the instruction table
         (i) DSM vs simulated DSM would be indistinguishable

5) The Question: “Is it true that by modifying this [digital] computer to have an adequate storage, suitably increasing its speed of action, and providing it with an appropriate programme, [it] can be made to play satisfactorily the part of A in the imitation game, the part of B being taken by a man?”
6) Discussion and Contrary Views
    a) Turing believes:
        i) “in 50 years it will be possible to programme computers, with a storage capacity of about 10^9, 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 5 minutes of questioning”
        ii) “at the end of the century the use of words and general educated opinion will have altered so much that one will be able to speak of machines thinking without expecting to be contradicted.”
    b) Theological Objection
        i) Do machines have souls? Can God confer a soul into machinery? An elephant?
        ii) Would need capabilities to minister to that soul.
        iii) Likens the argument to the solar system debate
    c) Heads in the Sand Objection
        i) Consequences of machines thinking would be horrible
        ii) Belief that man is superior (likely intellectuals value their own thinking too much)
    d) Mathematical Objection
        i) Logically, there are certain things a machine (infinite tape TM) cannot do
            (1) Godel’s theorem, and other results
            (2) TM cannot solve the halting problem
        ii) Disabilities that machines have, which humans do not
        iii) Humans also “too often” give wrong answers
        iv) “There might be men cleverer than any given machine, but then again there might be other machines cleverer again, and so on.”
    e) Argument from Consciousness
        i) Need for thoughts and feelings to influence the writing/output
        ii) Not the “monkeys with typewriters” could eventually produce Shakespeare’s works argument
        iii) Only way to be sure a machine thinks here is to be the machine and feel yourself think. The only way to know another man thinks is to be him and feel yourself think.

iv) Viva voce, similar to imitation game

f) Arguments from Various Disabilities
   i) You'll never make machines do X
   ii) Machines cannot make mistakes (machine could be programmed to introduce mistakes in arithmetic, not to give the right answer)
   iii) Abstract machines may not make errors of function, but may commit errors of conclusion to throw off the interrogator
   iv) Future machines (today's research) on self-modifying and self-repairing code
   v)

g) Lady Lovelace's Objection
   i) Computers can do whatever we know how to order them to do, but not originate anything on its own
   ii) Cannot take us by surprise

h) Argument from Continuity in the Nervous System
   i) Nervous system is not a discrete state machine, so we can't model human intelligence by a DSM

i) Argument from Informality of Behavior
   i) Behaviors of a man in every possible situation cannot be easily written down
   ii) Ex: stoplight that is malfunctioning with red and green both illuminated. Should you stop or go?
   iii) Writing down all rules appears impossible

j) Argument from Extra-Sensory Perception
   i) With E.S.P., anything is possible.
   ii) Interrogator may know who is who by telepathy.

7) Learning Machines
   a) Can a machine be made to learn more than the program it was given?
   b) An idea is injected into the machine
      i) Will it be subcritical and die away after being introduced (never grow)
      ii) Will it be supercritical instead of subcritical and grow into a new theory or line of thinking
   c) Turing estimates storage capacity of human brain: 10^{10}-10^{15} binary digits (Do you agree?)
      i) 10^9 should be satisfactory to play imitation game successfully
      ii) $2x10^9$ = content of Encyclopedia Britannica
      iii) 60 workers, working for 50 years, producing 1000 digits a day will accomplish this amount of work.
         (1) Will this work??

(2) Is his reasoning sound?

iv) REMEMBER

(1) 10^{15} binary digits = 125 TB

(2) 10^{9} binary digits = 125 MB

(a) MS Office '11 MAC: 180 MB

d) Argues for looking at a child and producing a child's brain, which is capable of learning

i) Faster than evolution, since experimenter can speed up process using his intelligence

(1) Can find proper mutations to fix any weakness in the process

8) Since 1950

a) 1991 Hugh Loebner started the annual Loebner prize contest. $100,000 prize to author of first program to pass an unrestricted Turing Test

i) http://www.loebner.net/Prizef/loebner-prize.html

ii) write a chat bot that will best mimic the responses of a human and trick the judge

iii) "Each year an annual cash prize and a bronze medal is awarded to the most human-like computer."

iv) 2013 Contest

(1) Judge, IBM Thomas J. Watson Research Center (Hawthorne, NY, USA)
Judge Professor Mike McTear (Emeritus), University of Ulster (Belfast, N. Ireland & Granada, Spain)
Judge Dr. Roger Schank, Professor Emeritus, Northwestern University & CEO, Socratic Arts, (Florida, USA)
Judge Professor Noel Sharkey, University of Sheffield (England)

v) 2012 Contest

(1) No bots fooled the judges

vi) 2011 Contest

(1) "None of the AI systems fooled the judges, therefore the Turing Test has not been passed and the Silver Medal was not awarded. The bronze medal and cash prizes were therefore awarded based on the ranks awarded by the judges"

vii) http://www.loebner.net/Prizef/2005_Contest/Transcripts.html

b) 100-year anniversary of Turing's birth Competition

i) Located where Turing cracked the Enigma code

(1) 150 Turing Tests completed

(2) Eugene won: http://www.princetonai.com/bot/bot.jsp

c) Cleverbot: http://www.cleverbot.com

d) ALICE: http://alice.pandorabots.com
e) The Long Bet Project is a wager of $20,000 between Mitch Kapor (pessimist) and Kurzweil (optimist) about whether a computer will pass a Turing Test by the year 2029

9) Discussion Questions
   a) What is needed to pass a Turing Test?
      i) Natural language processing
      ii) Reason
      iii) Knowledge
      iv) Machine learning?
      v) Vision?
      vi) Robotics?
   b) Should we consider passing the test?
   c) Do you think that these competitions and strategies to aid a machine winning the Turing test encompass the true nature of the test, or are they missing the bigger picture?
   d) With today's computing resources, will the Turing test actually measure intelligence of machines?
   e) How do machines such as IBM's Watson perform? Is it a fair comparison to distinguish a person from a machine that can search the entire Internet for responses?