SFU SURJ SIMON FRASER UNIVERSITY SCIENCE UNDERGRADUATE RESEARCH JOURNAL Vol. 4 (2018-19)

Software Daemons, Their History and Use in Modern Computing Science as an Answer to the Turing Test

Rohan Ben Joseph[1]*

[1]Simon Fraser University, *Department of Computing Science*

Abstract

Any sort of Technology can be divided into categories based on the amount of control a human has over it- Open Loop, which works only when a command is given to it, like a washing machine; Closed Loop, which uses feedback, such as a thermostat that regulates temperature within a house based on the temperature it detects on the outside; and Adaptive or Autonomous, where the control system 'learns' and adapts based on information it receives. For instance, Parker Atlas, Boston Dynamics' Humanoid bot that can sense and jump onto objects,[1] can be said to be an autonomous system. However, between Closed Loop control systems and Adaptive Technology, there exists a subset of sorts- Bots: A system that performs a series of pre-defined functions and has the potential to learn from its environment and perform more advanced actions within its predefined parameters. In this review, I look at the history of these pieces of code, from Socrates' Daemon, to the Child Machine; and examine the applications these bots could currently have in Computing Science in general and as an answer to the Turing Test in Specific.

Keywords — Programming, Computing Science, Daemon

1. Introduction

The British mathematician Alan Turing proposed the Turing Test as a replacement for the question "Can machines think?" in his 1950 Mind article 'Computing Machinery and Intelligence'[2]. Since then, Turing's ideas have been widely discussed, attacked, and defended over and over. At one extreme, Turing's paper has been considered to represent the "beginning" of artificial intelligence (AI) [3] and the Turing Test has been considered its ultimate goal. At the other extreme, the Test has been called useless, even harmful [4]. In between are arguments on consciousness, behaviorism, the 'other minds' problem, operational definitions of intelligence, necessary and sufficient conditions for intelligence-granting, and so on [5, 6, 7, 8].

In multitasking computer operating systems, a daemon is a computer program that runs as a background process, completing automated tasks without a need for interaction from the user. Named after Maxwell's Daemon, these computer programs could be used to perform or optimise practically any repetitive task [9].

*Corresponding Author: rohan.benjoseph@gmail.com

If carefully unpacked, the word daemon uncovers a provocative and useful dualism [10]. It is of more than academic importance that we learn to think clearly about the actual cognitive powers of computers, for they are now being introduced into a variety of sensitive social roles, where their powers will be put to the ultimate test: In a wide variety of areas, we are on the verge of making ourselves dependent upon their cognitive powers [11]. The key to these cognitive powers lies in the functions that these computers perform in the background, running calculations, deciding the priority of tasks, deciding the best way to allocate memory etc.; all of which it does through daemons and other repeating functions. It is through these daemons that computers can even hope to achieve the cognitive powers that humans have.

As machine learning aims to address larger, more complex tasks, the problem of focusing on the most relevant information in a potentially overwhelming quantity of data has become increasingly important. For instance, data mining of corporate or scientific records often involves dealing with both many features and many examples, and the internet and World Wide Web have put a huge volume of low-quality information at the easy access of a learning system [12, 13]. Similar issues arise in the personalization of filtering systems for information retrieval, electronic mail, netnews, and the like [14].

In the following review, I will attempt to provide an outline of their many applications, highlight the potential limitations to the studies, and point out future topics for exploration. I have considered several research articles and Theses associated in favour of the usage of bots and daemons in AI and some against. Thus, I will attempt to provide a balanced view of the long-term benefits and risks associated with bot use.

2. The Origin of Species

It is said that one of Socrates's friends took council of the Oracle of Delphi and inquired the name of the wisest person in the world, and the Oracle replied "Socrates" [15]. When Socrates found out about this, he tried to prove her wrong, and started on a life-long quest to search out wisdom in others by asking them questions - a method known as maieutics [16], an attempt to 'give birth' to the latent truth inside a person. Thus was Socrates' journey prompted [17].

Socrates's death came at the hands of the Athenian polity, when they charged him for asebeia (impiety) on two counts: corrupting the youth of the city (through his maieutics) and failing to acknowledge the gods of the city and introducing new gods. They gave him the choice of exile or death, and Socrates chose death by wilful consumption of a poisonous hemlock beverage [18]. It is easy to focus on the fact that he was sentenced to die because of his philosophizing, but we cannot dismiss the latter half of his sentence of impiety, that he failed to acknowledge the gods of the city and attempted to introduce other gods [19]. Why would they accuse Socrates of this? What did it mean?

"They are the envoys and interpreters that ply between heaven and earth, flying upward with our worship and our prayers, and descending with the heavenly answers and commandments, and since they are between the two estates, they weld both sides together and merge them into one great whole. It is only through the mediation of the spirit world that man can have any intercourse, whether waking or sleeping, with the gods. And the man who is versed in such matters is said to have spiritual powers, as opposed to the mechanical powers of the man who is expert in the more mundane arts."

- Plato while speaking about Socrates' Daemon in 'The Symposium' [20]

Socrates had a bot. Not in the literal sense, of course. But Socrates had a non-human helper, or so he claimed. He called this entity a daemon. Socrates' Daemon had many real, hard-coded linguistic and symbolic links with today's bots. It was Intelligent and ready to offer advice based on the situation without prompting, seemingly performing background functions without Socrates' interference.

In his 1867 thought experiment, "Maxwell's Demon" [21], James Clerk Maxwell attempted to show that thermodynamics is not strictly reducible to mechanics. Maxwellian Demons are mechanical devices that carry out measurements on a thermodynamic system, manipulate the system so as to extract work from it, and erase all records of the measurement outcomes [22]. If successful, they decrease the total entropy of the universe, thereby violating the Second Law of Thermodynamics [23]. Monitoring the speed at which the various molecules bounced around the chambers, the daemon could tell which specific molecules contained high or low energy states. Sliding a door open at intervals, he could separate the molecules into two different groups based on relative energy level [24].

Neither inward oracle nor false god, Maxwell's Daemon was hardly evil. He was merely a little otherworldly helper; wouldn't it be grand to be able to have little helpers fulfilling our wishes and doing our bidding? Maxwell's Daemon represented wishful thinking on a grand scale; until computer science.

3. The Imitation Game

The opening sentence of Turing's 1950 paper declares "I propose to consider the question, 'Can machines think?' [2] The paper provides a philosophical framework for answering this question. These 7 sections are briefly summarised below. The Imitation Game Often referred to as the "Turing test", this is a form of parlour game involving a human interrogator who alternately questions a hidden computer and a hidden person in an attempt to distinguish the identity of the respondents. The Imitation Game is aimed at providing an objective test for deciding whether machines can think.

Critique of the New Problem. Turing discusses the advantages of the game for the purposes of deciding whether machines and humans could be attributed with thinking on an equal basis using objective human judgement.

The Machines Concerned in the Game. Turing indicates that he intends digital computers to be the only kind of machine permitted to take part in the game. Digital Computers. The nature of the new digital computers, such as the Manchester machine, is explained and compared to Charles Babbage's proposals for an Analytical Engine.

Universality of Digital Computers. Turing explains how digital computers can emulate any discrete-state machine.

Contrary Views on the Main Question. Nine traditional philosophical objections to the proposition that machines can think are introduced and summarily dismissed by Turing.

The Child Machine. In the final section of the 1950 paper Turing addresses the motivation and possible approaches for such endeavours [25, 26, 27]. Turing goes on to discuss three distinct strategies which might be considered capable of achieving a thinking machine. These can be characterised as follows: 1) AI by programming, 2) AI by ab initio machine learning and 3) AI using logic, probabilities, learning and background knowledge [25, 28].

In the next sections we discuss various phases of AI research as it has been conducted over the past half century.

4. The Chinese Room Problem

The argument and thought-experiment now generally known as the Chinese Room Argument was first published in a paper in 1980 by American philosopher John Searle. It has become one of the best-known arguments in recent philosophy. It is one of the best known and widely credited counters to claims of strong artificial intelligence (AI)—that is, to claims that computers do or at least can (someday might) think. According to Searle's original presentation, the argument is based on two key claims: brains cause minds and syntax doesn't suffice for semantics [29].

Searle's Chinese Room experiment parodies the Turing test and echoes René Descartes' suggested means for distinguishing thinking souls from unthinking automata [30].

Its target is what Searle dubs "strong AI". In case of strong AI, Searle says, "the computer is not merely a tool in the study of the mind, rather the appropriately programmed computer really is a mind in the sense that computers given the right programs can be literally said to understand and have other cognitive states" [31]. Searle contrasts strong AI with "weak AI". In case of weak AI, computers just simulate thought, their seeming understanding isn't real understanding (just as-if), their seeming calculation is only as-if calculation, etc.

Searle asks us to imagine that a man is seated in a sealed room with 2 doors: one allowing input from one source outside the room (in the form of a slot) and one allowing output to the source outside the room (also in the form of a slot). The input from the outside source are Chinese squiggles that have been printed on card, but to the man in the room they are nothing more than incomprehensible gibberish (since he does not know the first thing about Chinese). The man is told that upon receiving the input squiggles, he must open a heavily-indexed reference book, wherein he must scrupulously track down the squiggle he received and find the matching squiggle of another sort. Once the man finds the matching squiggle, he must record it on an output piece of card and send it back through the output door's slot. Unknowingly the man has just performed some sort of translation that is altogether opaque to his understanding [32].

To the outside source, the Chinese room as a whole, is a sort of system and is being treated as a subject of a Turing test. The interested parties of the outside source are typing in questions in Chinese and receiving answers in Chinese. If the Chinese room is of good quality, then it should be possible to convince the interested parties that the room, or something inside it, is intelligent, thus suggested that the room, or something inside it, could pass the Turing Test. Searle suggests that this is an error, as the man in the room does not have any conscious states that exhibit and sort of understanding of the questions that he receives. To him it is all just squiggles. Certainly it might simulate intelligence impressively, but Searle suggests that this is precisely the problem, since it means only that we have an automata that is extremely good at fooling our test [33].

5. NLP and the ELIZA Effect

ELIZA was a computer program written by Joseph Weizenbaum of MIT University in the late 60s which is considered to be the first chatterbot, i.e. a program that can partially mimic a human in a conversation with a human [34]. In many ways ELIZA is has provided insights not just into what a serious NLP (Natural Language Processing) system should achieve but also has provided a lot of insight into human reactions to computer systems which look like "intelligent" systems but are not so. ELIZA was not meant to be an AI system, it was meant to be a toy or a parody system. ELIZA was first implemented in the SLIP language (Symmetric List Processor), a language incensed by Weizenbaum himself as an extension to FORTRAN but with better functionality to process doubly linked lists [35]. For its time, ELIZA was revolutionary in many aspects, as interfaces were not really common in the computers of the late 60s due to the absence of personal computing and thus the idea of interactive computation had not arisen yet or entered into popular fancy. Even though the perceived intelligence of it was an illusion (and a very bad illusion) the fact remains that it was the first genuine human machine interface (pretending to be a human – intelligent machine interface) attempting to use natural language [36].

As Weizenbaum discovered, many subjects who experimented with ELIZA got emotionally attached to it. Many did so despite Weizenbaum's informing them that there is no intelligence involved and that in fact ELIZA is not 'answering' them but only regurgitating a hardcoded script.

ELIZA was clearly just a daemon. It simply took input from a user and gave a previously stored vague response. It was a background process with no intelligence to speak of; the software equivalent of a ticket machine. However, the way Weizenbaum applied these processes into her code made her appear to be a form of Machine Intelligence. Another example of such a process would be Usenet's Serdar Argic.

6. Serdar Argic

In 1991, Usenet's culture and history discussions suffered under a flood of huge swaths of repetitive propaganda concerning the supposed Armenian murders of Turks in 1918 (history shows that the killing was the other way around), coming from a poster named Serdar Argic at a site known as zuma.UUCP [37].

Serdar responded to, seemingly, every and any Usenet post he could find that mentioned Turkey or Armenia, even in newsgroups that had nothing to do with either country. The poster was generally harangued with such phrases as "your criminal Armenian grandparents" (even if the poster happened to be of a non-Armenian ethnicity) and with over-the-top subject headings such as "The Self-Admitted Crook and Liar", "The Criminal SDPA-ASALA Grandparents of The Gum Brain", or "A mouthpiece for the fascist x-Soviet Armenian Government". This was usually followed by a lengthy essay concerning the alleged Armenian mass murders [38].

Some participants tried to argue with Argic, but that only made matters worse as he replied to each post with more harangues, along with successively more hysterical accusations concerning secret Armenian conspiracies. Some watched in amusement, and some even wrote parodies mocking the overwrought style of the posts [39]. But the amusement quickly turned to annoyance when it became apparent that the sheer volume of Serdar Argic posts was overwhelming the discussions on the hardest-hit newsgroups.

It quickly became apparent, however, that his responses didn't have much intelligence behind them. For one thing, they followed a distinct repeating pattern. For another, Argic did not appear to distinguish between the nation and the bird: posts containing references to Thanksgiving turkey were as likely to become targets as posts discussing Turkey's foreign policy [32].

Over time, a consensus built on the Usenet community: Serdar Argic was not a person, but a computer program which scanned the news articles and responded to any article that contained certain words, plugging in the name of the article's writer ("John Sugaharo's criminal Armenian grandparents") and other random phrases. Because of the robotic nature of the responses, this program was promptly dubbed "the zumabot" [40].

Serdar Argic clearly wasn't a well-developed example of Artificial Intelligence. However, shrouded in the anonymity of the net, he might well have been some crazed man incensed by the Armenian Genocide. Is this not exactly what passing the Turing Test would mean? And if so, couldn't a more advanced daemon having ELIZA like qualities give us a clear answer the the Turing Test?

7. The Future of Bots: The Child Machine

Turing closes the Mind paper with the following statement: "We can only see a short distance ahead, but we can see plenty there that needs to be done." As the present article indicates, Turing's vision was far from myopic. Indeed, he foresaw many of the key issues which dominated Artificial Intelligence research over the last fifty years [2].

In his writings on intelligence and machinery, Turing often employs analogies. One analogy he states explicitly and calls the "guiding principle" of his investigation into "possible ways in which machinery might be made to show intelligent behavior" is "the analogy with the human brain" [3, 41, 42].

The analogy may not be precise, but it is pretty clear: humans undergo education processes for a portion of their lives (which Turing estimates at about the first twenty years of their lives), and their behavior after that is very much affected by the education they have received, even though they still receive other interference – most of the time, in fact. The point is to approximate the human process of education with some analogous process suitable for machines [3]. The major points of his proposal are that, on analogy with a human's life, we plan for these three stages of a machine: first, there is the infant stage of a machine, which is a machine that has not been educated and is at least partly unorganized. It need not be a blank slate, but it is important that large amounts of its behavior are undetermined [2]. This is followed by the child-machine stage, during which the machine is educated. The first stage of education is to get the machine to a point where "it could be relied on to produce definite reactions to certain commands." Education involves a teacher who is intentionally trying to teach or modify the machine's behavior to effect some specific kinds of behavior. The machine's behavior is in flux during this time [3].

Even if the machine is given the means to educate itself using some kind of program during the child-machine stage, there is still oversight and monitoring by a teacher of sorts who checks up on its progress and intervenes if necessary. The machine that results when education is ended is supposed to behave in a way that can be predicted "in very broad outline" by someone familiar with how it has been educated — but its behavior might not, in fact probably will not, be fully predictable. Finally, there is the adult-machine, which is still capable of learning, but is also capable of quite complex behavior without additional intervention.

Thus, daemon qualities, if applied to a regular Machine Learning Process would be capable of so much more. Indeed, I believe that this would open up the door to a wide field of possibilities within the realm of Artificial Intelligence.

References

[1] Changing your ideas of what robots can do. URL https://www.bostondynamics.com/atlas.

[2] Alan M Turing. Computing machinery and intelligence. In *Parsing the Turing Test*, pages 23–65. Springer, 2009.

[3] Stephen Muggleton. Alan turing and the development of artificial intelligence. *AI communications*, 27(1):3–10, 2014.

[4] Virginia Savova and Leonid Peshkin. Is the turing test good enough? the fallacy of resource-unbounded intelligence. In *IJCAI*, pages 545–550, 2007.

[5] Katrina LaCurts. Criticisms of the turing test and why you should ignore (most of) them. *Official blog of MIT's course: Philosophy and theoretical computer science*, 2011.

[6] Robert M French. Subcognition and the limits of the turing test. *Mind*, 99(393): 53–65, 1990.

[7] Ayse Pinar Saygin, Ilyas Cicekli, and Varol Akman. Turing test: 50 years later. *Minds and machines*, 10(4):463–518, 2000.

[8] Stuart M Shieber. *The Turing test: Verbal behavior as the hallmark of intelligence*. Mit Press, 2004.

[9] A Patricia Ambler, Harry G. Barrow, Christopher M. Brown, Rod M. Burstall, and Robin J. Popplestone. A versatile system for computer-controlled assembly. *Artificial Intelligence*, 6(2):129–156, 1975.

[10] B Buchanan, E Feigenbaum, and N Sridharan. Heuristic theory formation: data interpretation and rule formation. *Machine intelligence*, 7:267–290, 1972.

[11] HG Barrow and SH Salter. Design of low-cost equipment for cognitive robot research. *Machine Intelligence*, 5:555–566, 1970.

[12] Stuart J Russell and Peter Norvig. *Artificial intelligence: a modern approach*. Malaysia; Pearson Education Limited,, 2016.

[13] Steven Pinker. *How the mind works*. Penguin UK, 2003.

[14] Susan G Sterrett. Too many instincts: Contrasting philosophical views on intelligence in humans and non-humans. *Journal of Experimental & Theoretical Artificial Intelligence*, 14(1):39–60, 2002.

[15] Richard Kraut. Socrates. URL https://www.britannica.com/biography/Socrates.

[16] Daniel Jones. *English pronouncing dictionary*. Cambridge University Press, 2006.

[17] Patricia E Easterling. *The Cambridge companion to Greek tragedy*. Cambridge University Press, 1997.

[18] Nicholas D Smith, Paul Woodruff, Paul B Woodruff, et al. *Reason and religion in socratic philosophy*. Oxford University Press on Demand, 2000.

[19] Essay on socrates: A great philosopher. URL https://www.bartleby.com/essay/Socrates-A-Great-Philosopher-PKA9ESYVJ.

[20] *Plato's Symposium: A Critical Guide*. Cambridge Critical Guides. Cambridge University Press, 2017.

[21] Andrew Rex. Maxwell's demon–a historical review. *Entropy*, 19(6):240, 2017.

[22] Charles H Bennett. Demons, engines and the second law. *Scientific American*, 257 (5):108–117, 1987.

[23] Nathanaël Cottet, Sebastien Jezouin, Landry Bretheau, Philippe Campagne-Ibarcq, Quentin Ficheux, Janet Anders, Alexia Auffèves, Rémi Azouit, Pierre Rouchon, and Benjamin Huard. Observing a quantum maxwell demon at work. *Proceedings of the National Academy of Sciences*, 114(29):7561–7564, 2017.

[24] John D Norton. The simplest exorcism of maxwell's demon: The quantum version. 2014.

[25] Donald W Davies. The bombe a remarkable logic machine. *Cryptologia*, 23(2): 108–138, 1999.

[26] Richard Langton Gregory. *Machine Intelligence*, volume 6. Halsted Press, 1967.

[27] Susan G Sterrett. Nested algorithms and"the original imitation game test": A reply to james moor. *Minds and Machines*, 12(1):131–136, 2002.

[28] Susan G Sterrett. Turing's two tests for intelligence. In *The turing test*, pages 79–97. Springer, 2003.

[29] John Searle. Can computers think. *Searle, J. Minds, Brains, and Science*, pages 28–41, 1984.

[30] John Preston and Mark JM Bishop. *Views into the Chinese room: New essays on Searle and artificial intelligence*. OUP, 2002.

[31] Ayse Pinar Saygin, Ilyas Cicekli, and Varol Akman. Turing test: 50 years later. *Minds and machines*, 10(4):463–518, 2000.

[32] John R Searle. Minds, brains, and programs. *Behavioral and brain sciences*, 3(3): 417–424, 1980.

[33] Joseph Weizenbaum. Computer power and human reason: From judgment to calculation. 1976.

[34] Joseph Weizenbaum et al. Eliza—a computer program for the study of natural language communication between man and machine. *Communications of the ACM*, 9(1):36–45, 1966.

[35] Lucy A Suchman. *Plans and situated actions: The problem of human-machine communication*. Cambridge university press, 1987.

[36] Wendy Grossman. *Net. wars*. NYU Press, 1997.

[37] Serdar argic in toronto's eye weekly. URL http://www.jaedworks.com/shoebox/serdar-eye.html.

[38] The zumabot's tale. URL http://www.jaedworks.com/shoebox/zumabot.html.

[39] URL http://www.columbia.edu/~sss31/Turkiye/lists.html.

[40] Politics in the middle east. URL http://www.cs.cmu.edu/afs/cs/project/theo-11/www/naive-bayes/20_newsgroup/talk.politics.mideast/77226.

[41] Donald Michie. Trial and error. *Science Survey, Part*, 2:129–145, 1961.

[42] James Moor. *The Turing test: the elusive standard of artificial intelligence*, volume 30. Springer Science & Business Media, 2003.