

Is Artificial Intelligence Possible?

CHRISTOFFER LARSSON



**KTH Computer Science
and Communication**

Is Artificial Intelligence Possible?

C H R I S T O F F E R L A R S S O N

Bachelor's Thesis in Computer Science (15 ECTS credits)
at the School of Computer Science and Engineering
Royal Institute of Technology year 2011
Supervisor at CSC was Johan Boye
Examiner was Mads Dam

URL: [www.csc.kth.se/utbildning/kandidatexjobb/datateknik/2011/
larsson_christoffer_K11094.pdf](http://www.csc.kth.se/utbildning/kandidatexjobb/datateknik/2011/larsson_christoffer_K11094.pdf)

Royal Institute of Technology
School of Computer Science and Communication

KTH CSC
SE-100 44 Stockholm, Sweden

URL: www.kth.se/csc

Abstract

A presentation of the most common philosophical theories concerning Artificial Intelligence. Furthermore it discusses intelligence, humanity and ethics concerning Artificial Intelligence.

Contents

| | | |
|----------|---|-----------|
| 1 | Summary | 3 |
| 2 | Introduction | 3 |
| 3 | Background | 3 |
| 3.1 | Presentation of various ideologies | 3 |
| 3.2 | Presentation of various terms | 4 |
| 3.3 | Mind and Consciousness | 5 |
| 3.3.1 | Science fiction | 5 |
| 4 | Intelligence | 6 |
| 4.1 | Definition | 6 |
| 4.2 | Intelligence for computers and double standards | 8 |
| 5 | Artificial Intelligence Theories | 9 |
| 5.1 | Turing Test | 9 |
| 5.1.1 | Presentation | 9 |
| 5.1.2 | Objections | 10 |
| 5.2 | The Chinese room | 13 |
| 5.2.1 | Presentation | 13 |
| 5.2.2 | Objections | 14 |
| 5.3 | The Brain prosthesis experiment | 16 |
| 5.3.1 | Presentation | 16 |
| 5.3.2 | Objections | 17 |
| 6 | Discussion | 17 |
| 6.1 | The Evolution of the Mind | 17 |
| 6.2 | Multithreading the Brain's processes | 18 |
| 6.3 | Understanding the Brain | 18 |
| 6.4 | Feelings and Personality | 19 |
| 6.5 | Artificial intelligence | 20 |
| 6.5.1 | Rights | 20 |
| 6.5.2 | Benefits | 20 |
| 6.5.3 | Disadvantages | 21 |
| 6.5.4 | Cyborgs | 22 |
| 7 | Conclusion | 23 |

1 Summary

This document gives a brief presentation of various ideologies that are connected to the debate of the possibility of artificial intelligence as well as the human mind. It speculates about the definition of intelligence in general and even what it means to be human. It also presents the most common and most quoted philosophical theories, the Turing test and Chinese room among some others.

The document summarizes the theories and objections against these. It also discusses the human body and speculates if there is anything one should try to mimic in order to achieve artificial intelligence.

The document proceeds to speculate about robot rights and ethics if artificial intelligence is possible. Do we have a right to do what we want with intelligence we have created? Furthermore, discuss benefits and disadvantages of artificial intelligence and what is truly essential to be considered artificial intelligence.

The document is then finished with conclusions.

2 Introduction

Is artificial intelligence possible? The question whether or not a machine can be intelligent has been intensely debated since the inception of the term “artificial intelligence” (and indeed well before that). Not just intelligent but also debating if machines are capable of having a conscious mind. This forces a reconsideration of the very meaning of the word intelligent. Some of the most outspoken opponents have been the American philosopher John Searle, known for his “Chinese Room” thought experiment, and the British mathematician Roger Penrose.

Scientists were much more positive in the beginning when the computer was first created. The computer was much better at calculating and there seemed to be no end to what was possible with a computer. People envisioned the world of tomorrow with flying cars and robots in every home. These goals were all extremely exaggerated and we are still far off from achieving even half of what people thought would be possible before the 21st century. These miscalculated predictions and disappointments made people give up on artificial intelligence, for a long time, and even now scientists are reluctant to speculate about the future. [10]

This essay will not try to speculate when artificial intelligence will be possible, it will focus solely on the possibility of artificial intelligence, admittedly with a very materialistic and minimalistic approach. However, it also discusses ethics and rights concerning artificial intelligence.

3 Background

3.1 Presentation of various ideologies

Not all of the following terms are mentioned in document but they will give you some insight in to the ideologies surrounding not only artificial intelligence but also man and mind. These will also help you get in to a appropriate state of mind before dwelling further into the philosophy and theories about mind, man and machines.

Functionalism is a theory which core idea is that mental states (beliefs, desires, being in pain, etc.) are constituted solely by their functional role - that is, they are causal relations to other mental states, sensory inputs, and behavioral outputs. Since mental states are identified by a functional role, they are thought to be able to be manifested in various systems, even perhaps computers, so long as the system performs the appropriate functions. Computers are physical devices with electronic substrate that perform computations on inputs to give outputs, similar to brains which are physical devices with neural substrate that perform computations on inputs which produce behaviours.

Mechanism is the belief that living things are like machines, composed of parts lacking any intrinsic relationship to each other. Thus, the source of a thing's activities is not the whole itself, but its parts or an external influence on the parts.

Dualism is a set of views about the relationship between mind and matter, which begins with the claim that mental phenomena are, in some respects, non-physical.

Materialism is the theory of materialism holds that the only thing that exists is matter; that all things are composed of material and all phenomena (including consciousness) are the result of material interactions.

Biological naturalism is a theory about the relationship between mind and body (i.e. brain), and hence an approach to the mind-body problem. It was first proposed by the philosopher John Searle in 1980 and is defined by two main theses: 1) all mental phenomenon from pains, tickles, and itches to the most abstruse thoughts are caused by lower-level neurobiological processes in the brain; and 2) mental phenomena are higher-level features of the brain.

Computationalism, or also referred to as computational theory of mind, is the view that the human mind ought to be conceived as an information processing system and that thought is a form of computation.

Singularitarianism is a technocentric ideology and social movement that is defined by the belief that a technological singularity - the creation of a superintelligence - is a likely possibility within the not too distant future, and that deliberate action ought to be taken to ensure that the "Singularity" occurs in a way beneficial to humans.

3.2 Presentation of various terms

Searle, in conjunction with the introduction of "the Chinese room" also introduced definitions of what he called "Weak AI" and "Strong AI" in order to be able to focus on the questions he found relevant. Searle proposed [16]:

Weak AI

The philosophical position that machines can demonstrate intelligence, but do not necessarily have a mind, mental states or consciousness.

Strong AI

The appropriately programmed computer with the right inputs and outputs would thereby have a mind in exactly the same sense human beings have minds.

In this document, the terms "real AI" or "perfect AI" may be used to mean the same as strong AI.

3.3 Mind and Consciousness

The mind has been a standard topic in philosophy from ancient times to present day. Great philosophers like Plato, Aristoteles and Descartes among others have pondered the nature of the mind. The ancient Greeks thought the mind was an epiphenomenon. The mental world exists as a derivative parallel to the physical world, affected by the physical world but not able to have an effect on the physical world.

One of the most famous philosophers of all time, Descartes, might be most famous for the quote "Cognito ergo sum"; I think therefore I am. But he certainly made it easy for himself; he did not define what it means to "think". You may think it is trivial; you understand what it means "to think". However it gets much more complicated when you try to describe it. How can we claim computers think, or to not be able to think, unless we have something to judge it by?

3.3.1 Science fiction

Man has for a long time been intrigued by the thought of creating artificial intelligence or life. It is constantly a reoccurring theme in media and entertainment. In Mary Shelley's "Frankenstein", the monster can be considered to be a created intelligence, a created man. In Alchemy there existed an idea of Homunculus, an artificially created man. [5]

Yet at the same time people seem to fear the thought of an artificial intelligence, with movies and books constantly depicting robots that will come to the conclusion that man is obsolete and hence should be eradicated. One of the earliest movies displaying this animosity towards artificial intelligence is Fritz Lang's Metropolis from 1927.

Others include, I, Robot, loosely built around the concept of Asimov's three laws of robotics, and two of the most famous films of all time, Matrix and Terminator. Among these Matrix is arguably the more philosophical of them and in fact the Matrix is very similar to one of the favourite devices of philosophers: the brain in the vat.

Brain in the vat [15]

Imagine that your brain was removed from your body at birth and placed in a vat. The vat sustains your brain, allowing it to grow and develop. At the same time, electronic signals are fed to your brain from a computer simulation of an entirely fictitious world, and motor signals from your brain are intercepted and used to modify the simulation as appropriate. The brain in the vat can be in an identical state to the brain of someone eating a hamburger; yet in one case, the hamburger exists; in the other, it does not. Even in the former case, then,

the belief only refers to the actual hamburger in the eyes of a third party who has independent access to the internals of the brain and to the external world containing the hamburger. The brain by itself does not refer to the hamburger.

The similarities with the Matrix, where most of the humans are only living within their own brain without even realizing it, are apparent.

4 Intelligence

4.1 Definition

Different areas of science or just in everyday life have different definitions of intelligence. All of them suited for their field. In psychology you try to measure a person's intelligence or if they are of proper mind as to not be a danger to themselves or society, while biologist might simply be concerned with if the organism is sentient. There are various different standards on intelligence, and there is also a clear need for a definition when it comes to artificial intelligence. Here is a few definitions of intelligence according to a few dictionaries and encyclopedias.

Definition of INTELLIGENCE [19]

1. a (1) : the ability to learn or understand or to deal with new or trying situations : reason; also : the skilled use of reason (2): the ability to apply knowledge to manipulate one's environment or to think abstractly as measured by objective criteria (as tests)
b Christian Science : the basic eternal quality of divine Mind
c : mental acuteness : shrewdness
2. a : an intelligent entity; especially : angel
b : intelligent minds or mind <cosmic intelligence>
3. : the act of understanding : comprehension
4. a : information, news
b : information concerning an enemy or possible enemy or an area;
: an agency engaged in obtaining such information
5. : the ability to perform computer functions

Let us examine these listings in more detail:

The first listing, is divided by "or". A person seemingly only has to learn something but not understand it is accepted as intelligent by this definition. However, let us say that a person has to be able to learn **and** understand **and** to deal with new or trying situations. One could claim that this is already possible, in the field of machine learning. You can teach a computer to learn to play backgammon having no knowledge of the game before, and even teach it to play it better than most humans. This seem to fulfil the requirements of

intelligence, furthermore seeing as how it plays it even better than most humans, is it therefore more intelligent than a human?

The second part of the first listing; the ability to apply knowledge to manipulate it's environment. You can make a robot learn to walk, it's body being the environment, or teach it to handle objects. Does this make it intelligent?

Furthermore, "To think abstractly" is too unspecified to serve as guidance to what constitutes intelligence. Moreover intelligence quotient tests have received a lot of criticism partly because IQ tests only examine particular areas embodied by the broadest notion of "intelligence".

The second uses the word "intelligence" to explain "intelligence" which is a metonym. The first listing contains the third so I'll leave it out of consideration.

The fourth is just another way of saying information and thus does not apply in this situation.

If we take it to mean as number five on the list then computers are already intelligent. When you start to question "intelligence" you involuntarily have to questions such as "understanding" as well. What does it mean to understand? To think?

In 1983 Howard Gardner proposed a the theory of multiple intelligences to analyse and better describe the concept of intelligence. For example a person can be considered to be an idiot yet still be a genius on the soccer field or with a violin in their hands. One proposed intelligence was mathematical intelligence, an area in which computers outperform any man. [7] A savant exhibits extraordinary ability in a highly specialized area, such as mathematics or music, but can be considered idiots or even be retarded in other areas. Is a general intelligence needed or can we consider savants intelligent, even if they are only geniuses in one area?

This debate has yet to be resolved and thus arguments for or against the possibility of artificial intelligence has to be preceded by a definition of what the writer means with "intelligence".

Here follows two definitions of intelligence from the Encyclopedia Britannica and the Swedish nationalencycolpedin:

According to the Encyclopedia Britannica:

Intelligence, in education, is the ability to learn or understand or to deal with new or challenging situations.

In psychology the term may more specifically denote the ability to apply knowledge to manipulate one's environment or to think abstractly as measured by objective criteria (such as the IQ test). Intelligence is usually thought of as deriving from a combination of inherited characteristics and environmental (developmental and social) factors. General intelligence is often said to comprise various specific abilities (verbal ability, ability to apply logic in solving problems, etc.), but critics contend that such compartments fail to reflect the nature of cognition and that other models, perhaps based on information processing, are needed. [3]

Nationalencycolpedin [12]

Intelligence, in everyday language, the same as wisdom, ingenuity, talent or intellect. Suggestions for definitions of intelligence have often emphasized abstract thinking, relational thinking, learning, adapting to new situations and effective utilization of experience. Intelligence testing has also been criticized for inadequate theoretical basis.

Observations show that people with good performance in one area also tend to exhibit good performance in other areas, which can be interpreted as expressions of a general intelligence. Covariation between measures of performance in specific areas, such as linguistic and mathematical domains, is even higher, indicating that even special talent dimensions are significant.

One popular way of measuring intelligence is with the intelligence quotient test, the so called IQ-tests. When IQ tests are initially standardized, using a standardization sample, by convention the average result is set to 100 and the standard deviation of the results is set to 15 points. When IQ tests are revised they are again standardized using a new standardization sample and the average result set to 100. However, if the new sample is tested using older tests in almost every case they score substantially above 100. This substantial increase in average scores on intelligence tests all over the world is called the **Flynn effect**, after J.R. Flynn who first documented this perpetual trend. However, Flynn believes that the increase is actually an increase in abstract problem solving rather than intelligence. Otherwise, someone who scored among the best 10% a hundred years ago, would today be among the 5% weakest. In short, someone who would be considered very bright a century ago, should now be considered a moron or even be retarded. [6]

4.2 Intelligence for computers and double standards

There seem to be much stricter standards of intelligence for computers. A lot of times we base our perception of intelligence on ourselves, what we can do. That is why a robot or computer often is expected to understand speech and be able to talk. However, a deaf person that can not even register speech is not considered unintelligent. Neither is a mute person who is unable to talk, or a blind person who can not read. These disabilities do not make them any less intelligent than the average man. Even a person that embodies both does not have to be of any lesser intelligence than any one else. The story of Helen Keller, a famous deaf-blind, yet she still managed to become an accomplished author, political activist and lecturer is a famous one and just goes to show that speech and sight does not contribute to intelligence. Being blind as well as deaf hindered her from learning to read, sign language or even to talk through conventional ways. However she overcame this handicap by spelling out the words in her hand. This is fairly similar to a computer that reads texts, reading it as string of characters one after another. This should be enough, if you strip away the unnecessary things this is all you need. A computer can be intelligent even without the ability to understand human speech or speak of its own, it does not have to “see” or even be able to move and yet it can still possibly be intelligent. The AI research might want to aim for more minimalistic goals.

5 Artificial Intelligence Theories

Here is presented two of the most well known thought experiments, the so called Turing test and "The Chinese room", as well as one referred to as "the Brain Prosthesis experiment". Each thought experiment is presented shortly, and for most part followed by objections against their hypothesis, which in turn is succeeded by comments on the viability of such objections or statements.

5.1 Turing Test

5.1.1 Presentation

Alan Turing suggested in 1950 in his article "Computing machinery and Intelligence", that the question "Can machines think" be replaced with a game that he named "The imitation game". A computer that succeeds at "winning" the imitation game is said to pass, what has come to be called the Turing test, and as such can be considered intelligent. It is thus, by Searle's definition, an example of weak artificial intelligence.

The imitation game is devised such that there is one interrogator and two participants, one of which is a computer. The interrogator will ask both of the participants various questions, all conversation taking place through text so the judge can neither hear or see the participants, and then tries to determine which one of them is human. If the interrogator mistakenly claims the computer to be the human, the computer pass the Turing test.

Turing also examined a wide variety of possible objections to the possibility of intelligent machines, including virtually all of those that have been raised since his paper appeared. [15] We will discuss those under the headline "Objections".

The Turing test is basically a behavioral test for computers. A modern version of the Turing's experimental design would probably use an online chat room. However there have been various adaptations of the Turing test. One of the most recent ones is a new competition that was introduced in 2008: 2K BotPrize. This contest is an adaptation of the Turing test for a first person shooter video game called Unreal Tournament 2004. During the competition, both human and artificial players enter a game anonymously. Human players are able to observe and judge other opponents, tagging them as humans or bots. If a bot is wrongly identified as human on more than 50 % of the judgements, it is considered to pass the Turing Test (no bot has passed this Turing test to date).

The winner of the 2010 edition of the BotPrize competition achieved a humanness ratio of roughly 32 %, however it was close to the "least human" human player, who scored 35 % humanness. [2]

There have also been other adaptations, such as a Turing test for computer generated music. In the experiment people were to judge which music was created by a human and which was created by the program. [10]

Turing notes that no one (except philosophers) ever asks the question "can people think?" He writes "instead of arguing continually over this point, it is usual to have a polite convention that everyone thinks." [18]. Turing's test

extends this polite convention to machines: If a machine acts as intelligently as human being, then it is as intelligent as a human being.

5.1.2 Objections

Some objections are more current than others. Objections like the theological, "Heads in the sand" or Extra sensory perception objections are deemed unnecessary for discussion or debate. For people who are interested they are recommended to read Turing's responses to these objections themselves.

The Mathematical Argument

It is well known, partly through the work of Turing himself as well as that of Godel, that certain questions cannot be answered correctly by any formal system. One example is the halting problem. Another example is Godel's incompleteness theorem.

The first incompleteness theorem states that no consistent system of axioms whose theorems can be listed by an "effective procedure" (essentially, a computer program) is capable of proving all facts about the natural numbers. For any such system, there will always be statements about the natural numbers that are true, but that are unprovable within the system. The second incompleteness theorem shows that if such a system is also capable of proving certain basic facts about the natural numbers, then one particular arithmetic truth the system cannot prove is the consistency of the system itself. Effectively Godel shows that to find a complete and consistent set of axioms for all of mathematics is impossible.

The philosopher J.R.Lucas bases his arguments on Godel's incompleteness theorem. Briefly, for any non-trivial formal system F (a formal language, and a set of axioms and inference rules), it is possible to construct a so-called "Godel sentence" $G(F)$ with the following properties:

- $G(F)$ is a sentence of F , but cannot be proved within F .
- If F is consistent, then $G(F)$ is true.

Lucas claims that because a computer is a formal system, then there are sentences whose truth it cannot establish - namely its own Godel sentences. But if we accept that the brain is a deterministic physical device operating according to normal physical laws, then it is just as much a formal system as the computer and thus has the same limitations as computers. [15]

Turing himself stated that

"Although it is established that there are limitations to the powers of any particular machine, it has only been stated, without sort of proof, that no such limitations apply to the human intellect." [18]

Furthermore, these "Godel - statements" always refer to the system itself, similar to "this statement is false" or "I am lying". This was first pointed out by C. H. Whitely, when he proposed the sentence "Lucas cannot consistently assert this sentence." If you think about it, you'll see that (1) it is true, and yet (2) Lucas cannot consistently assert it. So Lucas is also "incomplete" with respect to truths about the world. This shows that Lucas himself is subject to the same

limits that he describes for machines, as are all people, and so Lucas' argument is pointless. [9]

One should also note that Godel's argument only applies to what can theoretically be proved, given an infinite amount of memory and time. In practice, real machines (including humans) have finite resources and will have difficulty proving many theorems. It is not necessary to prove everything in order to be intelligent. [15]

Arguments from consciousness, various disabilities and informality of behaviour.

These arguments are fairly similar in that they all deny the machines of being capable of something. The arguments take the form of claims to the effect that "a machine can never do X." As examples of X, Turing lists the following: Be kind, resourceful, beautiful, friendly, have initiative, have a sense of humour, tell right from wrong, make mistakes, fall in love, enjoy strawberries and cream, make someone fall in love with it, learn from experience, use words properly, be the subject of its own thought, have as much diversity of behaviour as a man, do something really new. [18]

I would go as far as to say that a lot of those characteristics are not even present in some humans. Not everyone is kind or friendly. Not everyone is beautiful or can make people fall in love with them. Furthermore, there is no research that would support the theory of correlation between appearance and intelligence. Beauty has no influence on intelligence and thus is not a viable argument. Turing suggests that this scepticism of this nature arises from experience of seeing machines lacking these qualities. Naturally you conclude that having an ugly appearance, limited purpose, little variety of behaviours etcetera are necessary properties of machines in general. He draws similarities to the 1940s, when the general population found it difficult even to believe that machines could find the solutions of equations or predict ballistic trajectories. [18]

One of the most influential and persistent criticisms of AI is the claim that human behaviour is far too complex to be captured by any simple set of rules; and because computers can do no more than follow a set of rules, they cannot generate behaviour as intelligent as that of human. However claims that "X is impossible for computers" (e.g., X might be beating a chess master) tend to be overtaken by actual events. [10]

Lady Lovelace's objection

Lady Lovelace, commenting on Babbage's Analytical Engine, said,

"It has no pretensions to originate anything. It can do whatever we know how to order it to perform." [18]

This objection, that computers can only do what they are told to do and are therefore not capable of creativity is refuted simply by noting that one of the things we can tell them to do is to learn from their experience. [15] page 823) For example you can create a program that by using machine learning can learn to play chess so well that in fact it can beat its creator.

Argument from continuity in the Nervous System

”The nervous system is certainly not a discrete-state machine. A small error in the information about the size of a nervous impulse impinging on a neuron, may make a large difference to the size of the outgoing impulse. It may be argued that, this being so, one cannot expect to be able to mimic the behaviour of the nervous system with a discrete-state system.”

It is true that a discrete-state machine must be different from a continuous machine. But if we adhere to the conditions of the imitation game, the interrogator will not be able to take any advantage of this difference. For instance, if a differential analyser asked to give the value of pi (actually about 31416) it would be reasonable to choose at random between the values 312, 313, 314, 315, 316 with the probabilities of 005, 015, 055, 019, 006 (say). Under these circumstances it would be very difficult for the interrogator to distinguish the differential analyser from the digital computer. [18]

Weaknesses of the test

The Turing test is based on the assumption that human beings can judge a machine’s intelligence by comparing its behaviour with human behaviour. Every element of this assumption has been questioned: the human’s judgement, the value of comparing only behaviour and the value of comparing against a human. Because of these and other considerations, some AI researchers have questioned the usefulness of the test. In practice, the test’s results can easily be dominated not by the computer’s intelligence, but by the attitudes, skill or naivet of the questioner. One criticism of the Turing test is that it is explicitly anthropomorphic. If our ultimate goal is to create machines that are more intelligent than people, why should we insist that our machines must closely resemble people?

The Turing test does not directly test whether the computer behaves intelligently - it tests only whether the computer behaves like a human being. Since human behavior and intelligent behavior are not exactly the same thing, the test can fail to accurately measure intelligence in two ways:

Some human behavior is unintelligent

The Turing test requires that the machine be able to execute all human behaviors, regardless of whether they are intelligent. It even tests for behaviors that we may not consider intelligent at all, such as the susceptibility to insults, the temptation to lie or, simply, a high frequency of typing mistakes. If a machine cannot imitate human behavior in detail, it fails the test.

Some intelligent behavior is inhuman

The Turing test does not test for highly intelligent behaviors, such as the ability to solve difficult problems or come up with original insights. In fact, it specifically requires deception on the part of the machine: if the machine is more intelligent than a human being it must deliberately avoid appearing too intelligent. If it were to solve a computational problem that is impossible for any human to solve, then the interrogator would know the program is not human, and the machine would fail the test.

Because it cannot measure intelligence that is beyond the ability of humans, the test cannot be used in order to build or evaluate systems that are more intelligent than humans. Because of this, several test alternatives that would be able to evaluate superintelligent systems have been proposed.

In a recent experiment run at the Laboratory of Intelligent Systems in the Ecole Polytechnique Federale of Lausanne, Switzerland, robots developed the ability to lie with despite not haven been programmed to, and even going against their original programming. The experiment was constructed that 1000 small robots would search for good resources, and turn on a light when they had to let other robots know about the resource so they would also get some. However, resources are scarce so after a few generations of robots some stopped turning on their light when they found a good resource in order to keep it all to itself. [1] Is this intelligent behaviour? Is this human behaviour? It is easy to conceive that humans would act the same, but this might be considered inhuman, and even though a human might act in the same way, it might still not pass the Turing test. Moral questions make it difficult to decide intelligent behaviour.

5.2 The Chinese room

5.2.1 Presentation

“The Chinese room” thought experiment was first presented by John Searle in 1980. With “The Chinese room” concept he argues that even if a computer would pass the Turing test and is mistaken for human, does it truly understand the questions it is asked or the answers it gives? He also created the definition of weak and strong artificial intelligence to distinguish it from the AI in the Turing test as to be able to debate, not about whether it is intelligent or not, but if it has a mind. Searle clearly separates the functions of the brains from the “higher-level” mental phenomena, showcasing the theory of biological naturalism.

The system consists of a human, who understands only English, equipped with a rule book, written in English and various stacks of paper, some blank, some with indecipherable inscriptions. (The human therefore plays the role of the CPU, the rule book is the program, and the stacks of paper are the storage device.) The system is inside a room with a small opening to the outside. Through the opening appear slips of paper with indecipherable symbols. The human finds matching symbols in the rule book, and follows the instructions. The instructions may include writing symbols on new slips of paper, finding symbols in the stacks, rearranging the stacks, and so on. Eventually, the instructions will cause one or more symbols to be transcribed onto a piece of paper that is handed through the opening to the outside world.

From the outside, we see a system that is taking input in the form of Chinese sentences and generating answers in Chinese that are as obviously “intelligent” as those in the conversation imagined by Turing. Searle then argues as follows: the person in the room does not understand Chinese; the rule book and the stacks of paper, being just pieces of paper, do not understand Chinese; therefore there is no understanding of Chinese going on. Hence, according to Searle, running the right program does not necessarily generate understanding. [15]

5.2.2 Objections

Here follows a summary of objections to “The Chinese room” and John Searle’s rebuttal to these objections:

The Systems Reply

“While it is true that the individual person who is locked in the room does not understand the story, the fact is that he is merely part of a whole system, and the system does understand the story. Understanding is not being ascribed to the mere individual; rather it is being ascribed to this whole system of which he is a part.” [16]

Searle’s response is to reiterate the point that the understanding is not in the human, and cannot be in the paper, the man in the room is the only thing which could possibly “have a mind” or “understand”, so there cannot be any understanding. He further suggests that one could imagine the human memorizing the rule book and contents of all the stacks of paper, so that there would be nothing to have understanding except the human; and, when one asks the human whether the it understands Chinese, the reply will be negative. [15]

Comments

The objection is that although one can ask if the human in the room understands Chinese, this analogous to asking if the CPU can take cube roots. In both cases, the answer is no, and in both cases according to the Systems reply, the entire system does have the capacity in question. Certainly, if one asks the Chinese room whether it understands Chinese, the answer would be affirmative. The shift from paper to memorization is simply physical instantiations of a running program. The real claim made by Searle has the following form:

1. Certain kinds of objects are incapable of conscious understanding
2. human, paper, and rule book are objects of this kind
3. If each of a set of objects is incapable of conscious understanding, then any system constructed from the objects is incapable of conscious understanding
4. Therefore there is no conscious understanding in the Chinese room.

While the first two steps are on firm ground, the third is not. Searle just assumes it is true without giving any support for it. But if you do believe it, and if you believe that humans are composed of molecules, it follows that you must also believe that individual molecules are capable of thinking. [15]

Some argue that it is possible for there to be two minds in the same physical place, similar to the way a computer can simultaneously “be” two machines at once: one physical (like a Macintosh) and one “virtual” (like a word processor).

The Robot Reply

”The robot reply is that although the symbols manipulated by the Chinese room may not have real meaning to the room itself, a fully equipped robot would not be subject to the same limitations. Its internal symbols would have meaning to it by virtue of its direct experience of the world. Suppose we put a computer inside a robot and it would actually operate the robot. The robot would have a sensors attached to it that enabled it to see, it would have arms and legs that enabled it to “act”, and all of this would be controlled by its computer ”brain”.” [16]

Searle’s reply is that the addition of such “perceptual” and “motor” capacities adds nothing by way of understanding. If you put the Chinese room inside the robot’s “head”: the sensors are redesigned to generate Chinese symbols instead of streams of bits, and the effectors redesigned to accept Chinese symbols as control inputs. Then we are back were we started. Conversely, the outputs of human sensors, for example, along the optic nerve or the auditory nerve, might as well be in Chinese. It is hard to argue that connecting artificial sensors to these nerves would remove consciousness from the brain involved. [15]

The Brain Simulator Reply

Suppose we design a program that simulates the actual sequence of neuron firings at the synapses of the brain of a native Chinese speaker. The machine takes in Chinese stories and questions about them as input, it simulates the formal structure of actual Chinese brains in processing these stories, and it gives out Chinese answers as outputs. We can even imagine that the machine operates in the manner that actual human brains presumably operate when they process natural language. Now if we refuse to say that the machine understood the stories, wouldn’t we also have to deny that native Chinese speakers understood the stories?

Searle proposes the idea to let the man inside the room to operate water pipes instead of shuffling cards, with each water connection corresponding to to a synapse in the Chinese brain. After doing following the instructions and operating the water pipes in the correct way the Chinese answers is outputted at the end of the water pipes. Searle means that even though the system takes Chinese as input, it simulates the formal structure of the synapses of the Chinese brain, and it gives Chinese as output, the man does still not understand Chinese, nor do the water pipes. [16]

Comments

This is a variation on the “systems reply” that appears more plausible because “the system” now clearly operates like a human brain, which strengthens the intuition that there is something besides the man in the room that could understand Chinese. [15]

The Combination Reply and The Other Minds Reply

If you take all three previous replies together they are collectively more convincing and even decisive. Searle’s respond to this is to amalgamate the corresponding objections he raised against the previous arguments.

The combination of the three arguments would mean a robot with a brain-shaped computer lodged in its cranial cavity, the computer programmed with

all the synapses of a human brain, the whole behavior of the robot is indistinguishable from human behavior, and the whole thing as a unified system and not just as a computer with inputs and outputs.

Assuming the robot body was an imitation of the human body, it would be, for all intents and purposes, indistinguishable from a human. This makes it more apparent that something is awry with Searle's argument. It brings to question whether other people understand Chinese or anything else.

Several people have noted that Searle's argument is just a version of the problem of other minds, applied to machines. Since it is difficult to decide if people are "actually" thinking, we should not be surprised that it is difficult to answer the same question about machines. [15]

The Many Mansions Reply

The many mansions argument basically says that Searle bases his arguments on the present state of technology of analog and digital computers and that eventually we might have a technology that can recreate the causal process, which Searle claims to be essential for consciousness. To which Searle replies that he has no objections with this reply except that it redefines artificial intelligence, that mental processes are computational processes over formally defined elements. Searle says it is that thesis that he challenged and if the claim changes his objections do no longer reply because there would no longer be a testable hypothesis for them to apply to. [16]

However, as the previous comments about his rebuttals should show, the Chinese room parable has many flaws and does not prove that mental processes are not computational processes. Hence artificial intelligence may be achievable even with analogue and digital computers. Several critics have also pointed out that the man in the room would probably take millions of years to respond to a simple question, and would require "filing cabinets" of astronomical proportions. This brings the clarity of Searle's intuition into doubt. (wikipedia)

5.3 The Brain prosthesis experiment

5.3.1 Presentation

Suppose that we can build microscopic electronic devices that mimic the input/output behaviour and connectivity of all the neurons in the brain and can be smoothly interfaced to neural tissue. Lastly, suppose that we can replace individual neurons with the corresponding electronic devices without interrupting the operation of the brain as a whole. The experiment consists of gradually replacing all the neurons with electronic devices, and then reversing the process to return the subject to his or her normal biological state. The electrical devices that can only perform computations must therefore be equivalent to the neurons that are replaced. This shows that the experiment is based on the idea of computationalism and materialism.

By the definition of the experiment, the subject's external behaviour must remain unchanged compared to what would be observed if the operation were not carried out. Although the presence or absence of consciousness cannot easily be ascertained by a third party, the subject ought at least to be able to record any changes in his or her own conscious experience. [15]

5.3.2 Objections

There are no substantial objections against this argument. One person believed that if a person underwent the operation they might get trapped inside their own body, yet their body still operates and even answers like he would but his mind would be separated from his actions and words, while his mind slowly fades away. However, if this was the case then it should have already occurred when changed a very small part of his brain. If we check with the patient after every change of a neuron and he states he is fine, then it is highly unlikely that there will be a sudden change from changing one neuron. The brain loses neurons all the time, and it does not change with the disappearance of a few. Thus it is not plausible that the patient's mind would end up being separated while exchanging his brain cells.

6 Discussion

6.1 The Evolution of the Mind

The idea to have nature and man as a model when constructing machines is not a new one. Already Leonardo Da Vinci made sketches for inventions clearly inspired by models in nature. [17] In the beginning of artificial intelligence scientists glanced at nature, in particularly the human brain to see if they could steal any ideas, that could prove useful in an attempt to create intelligent machines. Particularly the neurons of the brain. There are approximately 10 billion neurons in the human brain. [9] The average neuron has as many as 1,000 connections [10], but it is possible for a neuron to have 200,000 separate entry ports. [9] However it is difficult to mimic neurons so it is usually custom to use a slightly simplified computer model of a neural network. [11] Yet, efforts to provide a detailed electronic models of real neurons appear to show that while real neurons are more complicated than typical computer models, the difference in complexity is modest. [10] This brings to question if we really need to aim at resembling our own brain as much as possible. There might yet be a model that exceeds the human brain that has yet to be thought of.

A neuron by another name is commonly referred to as a brain cell. In other words a neuron is just one single cell, and our brain is built up of billions that creates a whole. There are is a plethora of one single-celled organisms but they are generally not considered to have any sort of intelligence or mind and their behaviour depends solely on their physical mechanisms. Single-celled organisms do not have a mind, yet we do. Other higher lifeforms are also arguably intelligent, which points at that the mind is a product of evolution. Life on this planet began as single-celled organisms, without any intelligence. As we evolved we developed a body and organs that were more focused on specific tasks, and the brain that supervises and controls the body. Remembering danger, and location of food proved essential and thus the memory evolved. Is the mind just the next step in evolution? If the mind is a byproduct of our brain, which is only made up of live tissue and electricity, then surely someday we will be able to recreate it which would result in also creating a mind.

6.2 Multithreading the Brain's processes

A recent discovery has shown that besides the brain there is a mass of nerves called the enteric nervous system in the bowel region. This cluster of nerves have a higher amount of nerves than anywhere else in the body and though it has no higher level of mental capacity, has been shown to function separate from the brain. [4] Furthermore, we are said to not use our entire brain. However this of course refers to conscious thought. The rest of the brain is occupied with necessary information processing and filtering as well as various other autonomous functions. Besides this we are said to have one conscious mind as well as a unconscious one.

The differences in performance between machine and man is largely because of the way the brain is constructed and how it operates compared to the equivalent in today's technical systems. All information processing in the brain is performed in parallel in our highly connected biological neural network. On the other hand, large parts of the information processing in today's computers still occur sequentially with one instruction at a time. Everything that happens in a computer has to go through a few processors. The method can be compared to pushing tons of sands through the hole of a tiny hourglass. [17] Perhaps this division of tasks is what enables the brain to have conscious thoughts. Like a computer would use multi-threading or perhaps multicore processors to free up some computational time for higher functions. Each level in the brain is within itself a processor or even several processors and the mind is made up of several processors working together, with absolutely essential processes working in the background as to give more leeway for higher-level features of the brain which creates the mind.

Like the aforementioned case where a computer can be both a physical machine and a virtual machine. The mind would be something akin to a virtual machine.

In order to achieve conscious thought within a computer we might have to separate operations and aim for parallel systems.

6.3 Understanding the Brain

The human brain has occasionally been called the most complicated structure in the universe. One thing that has been studied is the functions of the brain's larger parts. Another thing that has been studied is the function of the brain's smallest parts, the so called brain cells and the synapses that connect the different brain cells with each other. However between these two extremes there are several levels and structures that has yet to be explored. [17]

Can the brain's top level - the mind - be understood without understanding the lower levels on which it both depends or does not depend. [9] If we have to understand the lower levels in order to understand the higher levels, then we have to understand the further lower levels of the lower levels in order to understand them and so on ad infinitum. It would be an endless recursion, going further and further down to the even lower levels. Do we really have to understand molecules in order to understand the neurons? Yet, do we have to understand the atoms in order to understand the molecules? Do we have to understand quarks and leptons in order to understand atoms? Usually you get by with a certain level of abstraction. You do not have to know how the

molecules work in order to know how a car works. In fact, most people only want to know enough in order to drive, not what propels it.

We do not need to know the lower levels, we can understand an organ without needing to know what molecules it is made of, but at what point do we have enough information to accurately explain the brain's activities? It might be in the levels between the larger parts and the neurons, or we might need to understand the intermediate levels as well as the neurons before we understand the nature of the mind.

6.4 Feelings and Personality

Things we closely relate with a mind are personality and feelings. Some objections against artificial intelligence are that a machine can never be moved by a song, or fall in love or like music etcetera but these objections fall short of objecting against artificial intelligence. Personality is a very abstract term. As long as someone shows consistency with their preferences and attitude it can be called their personality. Some people can be talked about as being boring and not having a personality, yet you might as well say that that just happens to be their personality. It is not the case of them lacking a personality, "not having a personality" is also a personality. Personality is not something you can have or be without, so it is irrelevant for artificial intelligence. Regardless of what the artificial intelligence does, or not do, it will have a personality.

Furthermore, personality is something you develop as you live, it is your experiences that shapes who you are. Any machine would probably be the same. Hence like a newborn child it might not have a personality when it is first created but will form a personality in due time.

The term "feelings" does not refer to the act of feeling (i.e. sensing), but instead a conscious subjective experience of emotion. Emotion is the complex psychosocial experience of an individual's state of mind as interacting with biochemical (internal) and environmental (external) influences. Basically emotions are the result of the body reacting to various situations. Primal emotions such as fear are present in all higher life forms. Animals that fear for their survival and self-preservation do not risk their safety and as a result tend to live longer and thus have a higher chance of procreating. There has also been research in inducing feelings with the aid of chemicals. For example, there is a substance in chocolate that when ingested gives the person a feeling similar to infatuation.

The fact that feelings can be induced by chemicals shows that they are mostly mechanical functions of the body and as such there is no reason why they should not be able to be replicated by a machine. However, perception of the physical world does not necessarily result in a universal reaction among receivers, but varies depending on one's tendency to handle the situation, how the situation relates to the receiver's past experiences, and any number of other factors. Thus, similar with personality, artificial intelligence must first gain experience before it would learn emotions.

Yet, emotions are not a requirement for artificial intelligence. Emotions are not even a requirement for high intelligence in humans. Psychopathy, though not the correct term anymore, is the commonly known term for a personality disorder characterized by and abnormal lack of empathy combined with strongly amoral conduct but masked by an ability to appear outwardly normal. The common stereotype of a psychopath is someone who lacks emotions, is highly

intelligent and uses it to manipulate people in his or her surrounding. How ever questionable the persons personality might be it does not change the fact that the person is highly intelligent and undoubtedly human. There are highly intelligent humans lacking certain emotions so why would artificial intelligence without emotions be considered less intelligent or to not have a mind? Emotions are not enough to argue the intelligence or the existence of a mind.

6.5 Artificial intelligence

Will artificial intelligence help us uncover the mystery that is the mind, or do we have to solve that in order to create artificial intelligence? Assuming we manage to create artificial intelligence, we will have to change the definition of what constitutes a person. Should human rights no longer be restricted to humans but be applicable to artificial intelligence as well?

6.5.1 Rights

What purpose will AI serve in our society? There is certainly benefits in making robots that are intelligent to an extent; rudimentary or dangerous work can be left for the much sturdier robots without the need for supervision. However is there really a need for strong artificial intelligence, or would it be better with less intelligent systems that only have enough intelligence for the required task and no more? Strong artificial intelligence might choose not to perform the task handed to them considering the risk to themselves, or if they are similar to humans, that they just might not want to do it or they might find it boring. If we consider artificial intelligence as people and extend human rights to include them, then we have no right to force them into labour. Even though it might be scientifically and financially in the our best interest to refuse them human rights and be able to use them as test subjects or force them to do hazardous and dangerous labour.

If we lack a way of determining the presence of a mind, or if the debate of what constitutes a mind has yet to be resolved before the arrival of artificial intelligence there is a possibility that they might be considered to not possess a mind. Apartheid towards artificial intelligence might certainly be a potential scenario as a result.

6.5.2 Benefits

An interesting question is how far the development will go when it comes to constructing technical systems that takes over tasks that man previously performed. Will machines be able manage all tasks that man does today? Will they perform better than us? Will we make robots that cleans and picks up after us in our homes? Will there be robots in human form play tennis and soccer just as well as today's players? Will there be robots that can correctly identify and name all people it has previously met and talked to?

A corollary question is what really separates man from machine when it comes to performing advance tasks that has to do controlling processes? What advantages does the human have with the implementation of different tasks? What advantages does the machine have? If man has some special advantages we also have to take into consideration if it is plausible that the machine will

eventually reduce this “lead” or if it has other advantages that outweigh the disadvantages. It is highly unlikely that man will change. Evolution is a much slower process compared to the technical developments. While machines and robots get more intelligent with each passing year, man will barely change in a thousand years. Man is still more intelligent, more mobile and more versatile than all robots, but that my change in not too distant future. [17]

The benefits of artificial intelligence are plenty. Computer networks are much faster than our neural networks [11], and we are not able to have store much more than the brain is capable of. All this means that artificial intelligence could potentially be smarter than humans. A computer that is able to give an intelligent diagnosis in medicine, finance and other areas could potentially be more accurate than their human counterparts. If suited with a much sturdier body than they can do heavier and more dangerous labour than man, their bodies could potentially be tailored for a certain task.

A report from Stanford Research Institute calculates that the market potential for intelligent or semi-intelligent robots will be as large as for cars in a few years. [17]

Robots could take over many dangerous, boring, unhealthy and monotonous jobs. It can be physically demanding environments: in mines, at the bottom of the ocean, in the heavy industry, in nuclear power plants, certain surveillance assignments, assembly work etcetera. It can also include welding, painting, cleaning, cashier work, gardening and so on. The benefits are apparent. Robots can work continuously without food or sleep and they do not require any extensive safety equipment. Some estimations indicate that 20-40% of today’s work belongs to the category dangerous, unhealthy and monotonous work, but of course it all depends on how you calculate. [17]

On the other hand, there is a possibility that robots will assume new tasks that today is not performed because they are not economically viable. Besides this, one can imagine “humanoids” that are only for private uses. Robots that mow the lawn, shovel snow, robots as tennis and conversation partners, robots as language teachers, robots that do housecleaning, paints the house etcetera. In the future perhaps everyone will have their own loyal robot, a constant companion that helps with every possible thing and that even can be pleasant company. [17]

6.5.3 Disadvantages

If intelligence also breeds personality there might be AI with bad or inappropriate personalities for the job they were tailored for. Even if you could change their personality within the program it would be highly unethical or maybe even illegal if they were to be considered to have the same rights as humans. You also get programs or robots who are unwilling to perform their tasks. It would also be unethical to knowingly have them perform life threatening work.

For example if we put artificial intelligence in every thing, even in video games. Would it be ethically justifiable to have real AI in video games? Real AI would certainly bring life and a sense of realism unprecedented in the gaming world but the implications would be considerable. If every non playable character was an AI, would it be defensible to kill in the game? Games are already intensely debated whether they are too violent or not. However, could this even be considered to be homicide? For all effects and purposes the AI is a person,

and killing it for entertainment is not to be encouraged. However, some might argue that this would prove to be a deterrent or help them vent so they do not kill “real people”. What kind of impact would that have on how people view other people? How would that affect people?

People can come to prefer the company of “humanoid” robots instead of other people. This is not that unlikely; there might be robots that offer more intellectual stimulation than what other people can provide. There might be robots that are more pleasant to talk to because they are more articulate, possess more knowledge and are better at expressing themselves. Already today there are a lot of people that constantly “hangs out with” computers, in the form of computer games and role-playing games. (This “socialization” gives both knowledge, entertainment and is intellectually stimulating). The step from this to spending time with even an more advanced technical system is perhaps not that far. There is a risk that robots can run amok and become dangerous to people or that they can go berserk and destroy property. [17] It is these concerns that have given rise to the ideology known as Singularitarianism.

6.5.4 Cyborgs

The mind might have to be properly understood and mapped before this can be a possibility but if we are capable of creating artificial intelligence and transfer it to an artificial body then it is plausible that we would also be able to take our own intelligence and transfer it to an artificial body or even a computer. Would people who transfer their mind to a robot body still be regarded as human? What if they just transplanted their brain to an artificial body?

Consider the thought that a man is injured in an accident and loses his hand, but receives a robot hand in replacement of his lost one. People with prosthetics are common and it is doubtful you would find any one who considers them less human because of it. Further consider if the man instead of a hand lost an arm, and had that replaced with a robot prosthetic. People would undoubtedly still consider him the same person. Still, consider if the person lost his entire body in an accident or to a terminal illness and they were successful in salvaging his brain and transplant it into an artificial body. Would you still regard him as a person? Of course this would be much easier if the artificial body was indistinguishable from his previous body, to the degree that if you did not inform his family they would have never have known he had changed bodies. Is this person human? Will everyone consider him human? His behaviour and way of thinking is unaffected, in the eyes of everyone and himself, he is, as you would say, “the same person”.

If we decide that the man is in fact still human then consider the following, the setup is opposite to the previous example, if we are able to save our “mind” on a computer or even manage to press it all into a chip and replace the only the brain of the human body, so that the person still has its original body but his brain replaced with a chip; is he still human? The brain could be replaced with artificial neurons as in the brain prosthesis experiment, they are both equivalent in theory. This time we changed the brain and not the body, with his heart and everything intact and unaffected by the change, as far as the body can tell, it is still controlled by the same brain. As previous, his behaviour or way of thinking has not changed even the slightest. Now if we pose the question again; is he human? This thought experiment really questions what it means to be human.

If we accept them both as human we have to accept that being human is not any physical feature, we can not ascribe “human” to anything.

7 Conclusion

Artificial intelligence has yet to be conceived but already it seems to be constrained and subject to prejudice. We seem to be suffering from a superiority complex. We have convinced ourselves that we are special, superior to other species and we are unwilling to accept anything that argues against that notion. This is probably most prevalent in religion, however this idea permeates society and even nonreligious people seem to share this point of view.

Hence, even though we strive to create an artificial intelligence we want it to be confined to what is possible with human intelligence. We still hold ourselves as the most intelligent beings and as such we judge other intelligence by comparing it to ours and effectively think that intelligent beings have to be as similar to us as possible. The more similar they are to us, the more intelligent they are considered to be.

We are unable to be objective, our bodies and brains are so intertwined it hinders us from making clear and objective observations. We are convinced that creativity is born from emotions, yet a computer could potentially simulate every possible combination and find something among all those possibilities that works and thus creating something truly new. Our brain is overflowing with hormones, meaning our thoughts and emotions are strongly linked to each other, making us unable to imagine intelligence without emotions. The demands of artificial intelligence also change over time. It seems that artificial intelligence would be what computers can not do today.

All dystopian views of artificial intelligence reflects our own insecurities and demerits. We are simply afraid that the artificial intelligence would do to us what we would have done to them. However, if we are intent on creating artificial intelligence then we will surely succeed at it one day. Humans consist of natural substances, born out of evolution and nature, and as such we should at some point be able to comprehend or mimic those substances. There is no reason why we should not be able to replicate or even improve upon it. The thought of strong AI is indeed exciting, yet it might be more preferable to have AI that are designed for specific tasks with limited intelligence that does not refuse orders. Then AI may come in all shapes and forms and not the just the predefined stereotypical views we have about it.

The philosophy of artificial intelligence such as the Turing test or the Chinese room are in fact fairly useless to artificial intelligence researchers, and are mostly ignored. Science is based on trial and error. You do not simply debate if something is possible but instead you prove it with experiments. Artificial intelligence is no different; artificial intelligence research will surely continue until we finally succeed or real physical undoubtable proof that it is in fact impossible is presented.

References

- [1] Abrams, Michael. (2008). Robots Evolve and Learn How to Lie. *Discover magazine*. January Edition. <http://discovermagazine.com/2008/jan/robots-evolve-and-learn-how-to-lie>.
- [2] Arrabales, Raul; Munos, Jorge. (2011). "Conscious-like Bot wins the 2K BotPrize". *ERCIM (European Research Consortium for Informatics and Mathematics) NEWS: Special theme: Intelligent and Cognitive Systems*. Number 84. 34-35.
- [3] *Britannica Online Encyclopedia*. (2011). <http://www.britannica.com/EBchecked/topic/289766/human-intelligence>.
- [4] Brown, Harriet. (2005). "The other brain also deals with many woes". *The New York Times*. (August 23 Edition). <http://www.nytimes.com/2005/08/23/health/23gut.html>.
- [5] Campbell, Marry Baine. (2010). "Artificial Men: Alchemy, Transubstantiation, and the Homunculus". *Republics of Letters: A Journal for the Study of Knowledge, Politics, and the Arts* 1. no 2. <http://rofl.stanford.edu/node/61>.
- [6] Flynn, J R. (1994). "IQ gains over time". *Encyclopedia of human intelligence*. In R. J. Sternberg (Ed.). 617-623. New York: Macmillan.
- [7] Gardner, Howard. (1983). *Frames of Mind: The Theory of Multiple Intelligences*. New York. Basic Books.
- [8] Hadhazy, Adam. (February 12, 2010). "Think Twice: How the Gut's 'Second Brain' Influences Mood and Well-Being". *The Scientific American*. <http://www.scientificamerican.com/article.cfm?id=gut-second-brain>.
- [9] Hofstadter, Douglas R. (1979). *Godel, Escher, Bach: an Eternal Golden Braid*. Basic Books. ISBN: 0-14-017997-6.
- [10] Kurzweil, Ray. (1999). *The age of Spiritual Machines: When computers exceed human intelligence*. Viking Penguin. ISBN: 0-670-88217-8.
- [11] Marsland, Stephen. (2009). *Machine Learning: An Algorithmic Perspective*. Chapman & Hall/CRC machine learning and pattern recognition series. Taylor & Francis Group. ISBN: 978-1-4200-6718-7.
- [12] *Nationalencyklopedin*. (2011). <http://www.ne.se/lang/intelligens>
- [13] Penrose, Roger. (1989). *The Emperor's New Mind: Concerning Computers, Mind, and The Laws of Physics*. Oxford University Press. ISBN: 0-14-014534-6.
- [14] Penrose, Roger. (1994). *Shadows of the Mind: A Search for the Missing Science of Consciousness*. Oxford University Press. ISBN: 0-19-853978-9.
- [15] Russel, Stuart J.; Norvig, Peter. (2003) *Artificial Intelligence: A Modern Approach*. 2nd edition. Upper Saddle River. New Jersey: Prentice Hall. ISBN: 0-13-790395-2.

- [16] Searle, John R. "Minds, Brains, and Programs". *The Behavioral and Brain Sciences*. vol. 3. 417-457.
- [17] Thomas, Bertil. (2003). *Framtidens Intelligens*. Lund. Studentlitteratur. ISBN: 91-44-03058-4.
- [18] Turing, A M. (1950). "Computing machinery and intelligence". *Mind: A quarterly review of psychology and philosophy*. Vol. LIX. No. 236. 433-460.
- [19] *Mirriam-Webster's online dictionary*. (2011). <http://www.merriam-webster.com/dictionary/intelligence>.

