Is Artificial Intelligence Possible?

Rishie Sharma
880531-0813
rishie@kth.se

2011

Abstract

In this paper the question, which is also the title of this paper, "Is artificial intelligence possible?" is examined. We investigate earlier thoughts regarding this question by famous thinkers such as Turing, Searle, Penrose and Chalmers and move on to consider what artificial intelligence is actually supposed to mean. Disregarding philosophical questions such as if machines can have a consciousness we turn our focus towards the possibility of machine intelligence matching or more interestingly, exceeding human intelligence.

Sammanfattning

I denna uppsats undersöks frågan, som också är titeln för den här uppsatsen, "Är artificiell intelligens möjlig?". Vi utforskar tidigare tankar rörande denna fråga av berömda tänkare som Turing, Searle, Penrose och Chalmers för att sedan gå vidare och begrunda vad artificiell intelligens egentligen är tänkt att betyda. Utan att fästa avseende vid filosofiska frågor som till exempel om maskiner kan ha ett medvetande, skiftar vi vår uppmärksamhet till om artificiell intelligens kan vara lika bra som eller mer intressant, överträffa mänsklig intelligens.

Contents

1 Introduction

1.1 This paper

This paper is part of of a bachelor's degree project done at the Royal Institute of Technology in Sweden during the spring semester of 2011. The course name is Degree Project in Computer Science, First Level, DD143X and it is in total 15 points. This paper is 6 of those 15 points while the remaining 9 is a computer software programming project.

1.2 Artificial Intelligence

Ever since its inception, artificial intelligence (AI) has had a wondrous connotation to it. People think of robots taking over the world, super-intelligent machines incapable of error, technological singularities and the like. While the field of artificial intelligence has progressed substantially since its birth it still has a long way to fulfil the dreams of science fiction. The question is, is it even possible to fulfil those dreams, and if so, what consequences will it have?

2 Background

In this section an overview of the most common ideas and thoughts related to the question "Is artificial intelligence possible?" is presented without any reflection or discussion.

2.1 Turing Test

Turing opted to replace the question of "Can machines think?" or in this case "Is AI possible?" with a test[1]. A test no machine since has ever passed. The machine's goal in this test is to fool a human to think that the machine is human. In the original version, the machine was to fool a human judge to think that the machine is a woman. A real woman was also to be present in the test doing her best to try to convince the human judge that she is the woman and not the machine. The idea is that if a machine can perfectly mimic human intelligence its intelligence is, at least from a materialistic point of view, the same as human intelligence. The test must be performed in such a way that only intelligence is measured and not any physical characteristics of the machine. This can effectively be achieved by limiting the communication to typed text only.

Turing had grand dreams for AI and the future and predicted that doubts against thinking machines would go away by themselves as technology advances.

> "I believe that at the end of the century the use of words and general educated opinion will have altered so much that one will be able to speak of machines thinking without expecting to be contradicted." [1]

Unfortunately 11 years has already passed since the end of that century and no one in their right mind would speak of machines thinking except in a figurative sense when your computer is being particularly slow.

The Loebner prize is an annual competition based on the Turing test that has been held since the year of 1990[2]. The competition awards prizes to the computer programs whose behaviour is considered to be the most human-like. There is also a one-time-only prize which when awarded, will end the competition.

> "In 1990 Hugh Loebner agreed with The Cambridge Center for Behavioral Studies to underwrite a contest designed to implement the Turing Test. Dr. Loebner pledged a Grand Prize of $100,000 and a Gold Medal (pictured above) for the first computer whose responses were indistinguishable from a human's." [2]

The competition is still being held as of the year of 2011.

2.2 Strong vs. Weak AI

There is a difference of opinion regarding what intelligence actually is. Some claim that even if a machine can accurately mimic human behaviour, it is not actually thinking. It is simply following a set of rules and cannot be said to possess a consciousness. Strong AI usually refers to the question if machines are able to have a consciousness, intentionality, self-awareness, understanding or other vague philosophical concepts often ascribed as human traits.

2.2.1 Chinese room

The most famous thought experiment put forth against Strong AI is that of the Chinese Room by John Searle[3]. Searle does not have any problems with the idea of Strong AI per se but rather with the idea that Strong AI is possible with the computers of today.

> "I see no reason in principle why we couldn't give a machine the capacity to understand English or Chinese, since in an important sense our bodies with our brains are precisely such machines. But I do see very strong arguments for saying that we could not give such a thing to a machine where the operation of the machine is defined solely in terms of computational processes over formally defined elements; that is, where the operation of the machine is defined as an instantiation of a computer program." [3]

The Chinese Room thought experiment is in short as follows. Let us say we can construct a computer program that takes Chinese sentences as input and outputs valid Chinese sentences in such a way that to an outsider it appears as if the machine is fluent in Chinese. Now let an English speaking person who does not understand Chinese at all enter a locked room and follow the instructions of this program. Through an opening the man receives input text in Chinese, follows the instructions of the program as to what to do with the input and through a second opening gives the produced output, which is also valid Chinese, to an outsider. To someone watching from outside it would seem like something inside the room understands Chinese, but the English speaking person has no idea of what the Chinese sentences he is receiving and producing actually mean. He is merely following rules and therefore it is wrong to say, according to Searle, that a computer program can actually understand something.

Common counter arguments include pointing out that the human in this case takes the role as the CPU or processing unit in this machine and is akin to arguing that one neuron in your brain is merely receiving and sending electrical signals and cannot be said to understand Chinese. The whole brain (of a Chinese speaking person) is however able to understand Chinese and in similar fashion the whole room with instructions, pen, paper, English speaking person etc. is able to understand Chinese. Searle's answer to this is to let the English speaking person remember all the instructions by heart and keep track of everything in his head. Now the whole system is in his head, but still he does not understand any Chinese.

Others argue that it is impossible for us to objectively judge other minds subjective experience. This is an argument directed more for Strong AI in general and is based on the philosophical problem of other minds. To you, all other humans act intelligent, but there is no way for

you to know what their subjective experience is like. You cannot know if they possess genuine understanding and consciousness but we follow the polite convention of assuming that all humans do posses Strong AI. It is argued that Searle sets higher standards for machines than we would for humans and that if a machine acts intelligent it should be assumed to be intelligent unless proven otherwise.

2.2.2 Silicon Brain

On the other hand an argument for Strong AI is the thought experiment by David Chalmers often called the Silicon Brain[4]. Picture a human brain. If anything can be said to be conscious it would be the human brain. Now imagine we could replace a neuron by an artificial neuron with the exact same chemical properties and replace neurons one by one until the entire brain is artificial, what would happen with your consciousness? There are three possible scenarios. Either your consciousness gradually disappears as you replace the neurons with artificial neurons, it disappears as soon as you replace one neuron or your consciousness remains unchanged throughout the whole process. Chalmers argues that the former two scenarios are unlikely and the only reasonable thing to happen is that your consciousness would remain all the way until your brain is entirely artificial.

2.3 Logical Weakness

K. Gödel showed in 1931 that for any consistent system capable of arithmetic there exists sentences that are true but not provable within the system[5]. This means that any theory capable of expressing elementary arithmetic cannot both be consistent and complete. This is called Gödel's incompleteness theorem and some people argue that computers' intelligence are limited by the theorem in a way that humans are not. A famous spokesperson for this view is Roger Penrose.

Roger Penrose argues that with our minds we can transcend the rules of formal systems, sidestep difficulties in reasoning and think outside of the box that constrains the computer[6]. He believes that the awareness of the mind is not something that is out of the scope of science but rather, outside of the scope of science of today. In his book Shadows of the Mind he argues for this view and talks about where these missing pieces of science might be found.

There are a number of counter arguments against this view[7]. It is argued that Gödel's theorem applies only to formal systems capable of doing (infinite) arithmetic. A Turing machine is such a system and thus limited by the incompleteness theorem. However, a physical computer is not a Turing machine because of the lack of infinite memory. In fact a computer can be described as a very large system in propositional logic and thus is not limited by the incompleteness theorem. More over some problems that are undecidable on Turing machines are in fact decidable on real computers, such as the Halting Problem.

Another counter argument is that humans are machines too and if computers have certain limitations one can only assume humans are subject to the same limitations until proven otherwise. Lastly intelligence does not have to be limited because of the inability to prove certain logical sentences. For example, Adam cannot assert that this sentence is true: "Adam cannot consistently assert that this sentence is true" but we do not think less of his intelligence for it.

3 Definition of AI

This section will have a look at different definitions of AI and will try to establish what I feel intelligence is so that it can be used in the discussion of the next section Possiblity.

3.1 Field of artificial intelligence

Those who study the field of artificial intelligence would most likely say that they are studying and designing intelligent or rational agents. With a rational agent one refers to something that acts so as to achieve the best outcome or, when there is uncertainty, the best expected outcome in any given situation[7]. In effect the field of artificial intelligence is not at all concerned with vague concepts such as consciousness, subjective experience etc. as is expected of engineers. Instead their focus is on a practical and sound definition allowing for development of practical and sound techniques of implementing this definition of AI. For the purpose of this paper however, said definition is too trivial and has a slightly too obvious answer to be of relevance.

3.2 Turing test

A machine capable of passing the Turing test would certainly qualify as AI in my book. However, there may be other kinds of AI apart from one mimicking human intelligence. Imagine an intelligent robot capable of communicating with humans. Even though he is intelligent it might be obvious to a human that he is a robot and would thus fail the Turing test. It might be capable of reasoning, abstract thinking and excel at many activities and yet incapable of fooling a human that it is a human. Robots mimicking human behaviour like those which would be able to pass the Turing test would be desirable when working with customer service but such behaviour might be unnecessary when doing certain scientific research or other activities that do not require much casual interaction with humans. An AI who can help us prove theorems or find ways of creating new materials by use of intelligent reasoning should be excused from not being able to pass the Turing test. Thus the definition of AI where only machines capable of passing the Turing test are said to be intelligent seems too strong.

3.3 Discussion about intelligence

Intelligence is not something that is either present or not present. If adult humans possess strong intelligence then there are many weak types of intelligence. A dog obviously possesses some sort of intelligence. What about a human baby or a 5-year old? Most people would say that gorillas are intelligent. To find out what intelligence might be let us instead consider things that are obviously not intelligent.

A stone is not intelligent, not aware, has no understanding and is completely, entirely dead. A star is not intelligent. A house is not intelligent (although lately the term intelligent house has become popular). Likewise many would argue that a computer is not intelligent. But with the advance of computer science surely it is in our right to ascribe some intelligence to advanced systems. Google figures out what you want to search for before you even type it out and even if you type it incorrectly. Is that not some sort of intelligence? Through the use of GPS and other positioning techniques your phone knows where you are and can tell you how to go where you want. Is that not more intelligent than anything a dog can do? With advanced sound analysis techniques a machine can listen to you humming a song and tell you which song it is[8].

People would argue that these systems do merely what we tell them to do. They have no real intelligence, they have no real understanding. But are not chemical and physical reactions the same? They do what the creator of the universe told them to do, and what are our brains if not a bunch of chemical and physical reactions? Quantum physical reactions, some would argue. In quantum mechanics a particle/wave is described as a probability of finding it in a certain position at a certain time. Certain experiments even imply that one particle/wave has been in two places at once. This certainly challenges the classical deterministic view of the universe. Whether or not this has any importance to the question of if artificial intelligence is possible is unfortunately beyond the scope of this paper.

If one believes there is something else to intelligence than mere cause and effect then the answer to the question "What is intelligence?" is simple in a sense. Intelligence is that something else, whatever it may be. On the other hand if there is nothing else to intelligence than mere cause and effect then there would be no essential difference between a stone falling when dropped, a dog running off to fetch a ball when thrown and a person solving a mathematical problem when given one. It is a difficult idea to accept; the idea that maybe there is nothing special about our intelligence.

3.3.1 Postulates

Even if intelligence is completely deterministic there is a difference between a stone falling when dropped and a student solving a mathematical problem. The student is using the student's intelligence which is something I would like to describe as including the following abilities:

1. The ability to learn by example.
2. The ability to improve the efficiency of performing a certain task by analysing your own behaviour and modifying it appropriately.
3. The ability to generalise solutions and so that they can be applied to varying situations, while realising that using a generalised solution in a specific situation rarely is as efficient as using a solution optimised for that specific situation.

However this is merely one part of intelligence. There is another part to intelligence which is the part allowing us to interact with each other.

4. The ability to communicate the solution to a problem to another intelligent being in such a way that the other intelligent being can accurately solve the problem.
5. The ability to learn the solution of a problem by communicating with another intelligent being.
6. The ability to work together with another intelligent being in order to accomplish a task that was previously impossible (or simply difficult) to accomplish alone.

These abilities are not meant to be *the* definition of intelligence but rather some fundamental abilities I believe to be essential for an intelligent being. I will try to justify all of them briefly. Take for example the act of playing badminton. A normal human would be able to follow the instructions and examples of a teacher to learn the rules of the game using abilities 1 and 5. Likewise the teacher is using ability 4 and simply performing the task to teach the newcomer. While not explicitly learning what to do in every specific situation that can possible occur within a game of badminton, you are still able to generalise what you have learned with ability 3 and use it to play badminton according to the rules. Now to become a good badminton player you would have to keep practicing and keep learning by example to learn more about what to do in certain situations that often occur within the game. Using ability 2 you would have to perfect

your swing to be able to smash as hard as possible while also perfecting your accuracy to be able to place the shuttlecock exactly where you want it (where your opponent is not).

In badminton there are also *doubles* which are games where two pairs battle against each other and the rules slightly differ. In the case of doubles well thought out team play is required to win over your two opponents. You would have to use your sixth ability to carefully decide who takes which shuttlecocks rather than both of you running towards the shuttlecock, running in to each other and hitting each other with your rackets (more common than you would think among rookies). While I do realise that this sixth ability is very broadly defined and would probably benefit from being defined into several more detailed abilities (like abilities involving devising tactics, communicating them to the individuals involved in the tactic etc.), the ability is defined well enough for the purpose of this paper.

However, the list is lacking something important still. It is something that humans may not do on a daily basis but certainly have the ability for, *to step outside of the box*, so to speak. With the abilities I have mentioned so far there is obviously some sort of intelligence involved but with only these abilities you would be stuck inside the box. While the earlier abilities can be worded in a precise and mechanical manner, this last ability I found difficult to express without the use of vague or ambiguous words.

Consider the most ingenious inventions or acts of mankind. With only the aforementioned abilities, how would a population invent the wheel for example? At first thought it may seem impossible because of the mechanical nature of the abilities. On second thought however, it is merely a matter of one individual watching something rolling down a hill, using the first ability to learn by example and adapt the example with the third ability to its problem of moving something. But what about the ability to make fire? First off you would have to see fire in a natural fire caused by for example hot and dry weather or a lightning strike. In this case however, how can you learn from example? Learning to make a fire seems to be a more complicated procedure requiring you to realise for example that fire is created in hot and dry conditions and then remember where you have felt hot and dry conditions before and how those were created. It seems to require some sort of fundamental understanding of cause and effect that the previously mentioned abilities lack. You could also come to the conclusion that fire is caused by lightning and by chance happen to see lightning sparks created when striking certain objects together and optimally use both insights to create fire in hot and dry conditions with the aid of lightning sparks.

7. The ability to deduce properties about objects and events by the use of observation and experiment and use these properties to solve problems.

This seventh ability has by itself limited usefulness. It lacks a generalising ability needed to draw general conclusions about certain classes of objects. I have come to think of this ability as *the racist ability* because it in essence encompasses prejudice of behaviour of objects, including other intelligent beings. While the third ability could be reformulated to allow this sort of ability it I could not find a satisfying reformulation and decided to formulate a new one. Even if this ability has led to many horrible acts throughout human history, it is a very important part of our intelligence.

Picture an intelligent being. It sees a stone fall and by the seventh ability asserts that the stone falls. It sees a leaf fall and uses the seventh ability to assert that the leaf falls. It sees water flow downwards when flowing over a sloped surface and asserts that the water flows downwards. But with only the seventh ability it would be surprised (assuming it is capable of having such a feeling) each time a stone or a leaf falls. It would have no capability of assuming that similar objects would behave in a similar way.

8. The ability to generalise previously observed properties about objects to hold for similar

objects you have not previously seen to have this property.

Even if I have not been able to encapsulate Einstein with these eight abilities, a being capable of all these abilities would, I believe, possess some intelligence.

4 Possibility

In this section we examine each of the abilities mentioned in the previous section one by one. If there are machines today already more or less capable of the ability then those are mentioned otherwise some speculation is carried out of how the future might look like. Lastly some discussion about the possibility of one machine doing all of these things intelligently is carried out.

The ability to learn by example and adapt the example procedure properly to fit the problem you are facing.

Learning by example is essentially what supervised learning with artificial neural networks (ANN) is[7]. You give the ANN problems with given solutions and it will form a model more or less matching the given problems and solutions and then you can give it a new problem without a solution and it will give you a solution. ANNs are computational models inspired by biological neural networks. They are a good alternative when solving problems that are difficult to solve with classical algorithmic techniques. Such problems include handwritten text recognition and face identification among others.

The ability to improve the efficiency of performing a certain task by analysing your own behaviour and modifying it appropriately.

This ability is essentially the problem of optimising. There are many optimisation techniques used within mathematics to optimise functions but to use those the intelligent being would have to have a mathematical model of the problem it is facing. A technique used within the field of artificial intelligence that requires no mathematical model is Q-learning. Q-learning is reinforcement learning technique which tries to maximize *rewards*. These rewards are for example winning a game or when learning to walk any forward movement is a reward. Basically doing well at the task you are performing is a reward and a Q-learning algorithm tries to maximise these rewards and thus in turn maximise performance. Of course the Q-learning algorithm is just one algorithm among many and has pros and cons in different situations, but the point is that there are today computer programs capable of this ability to at least some extent.

The ability to generalise solutions so that they can be applied to varying situations, while realising that using a generalised solution in a specific situation rarely is as efficient as using a solution optimised for that specific situation.

This ability would be useful in machines capable of performing varying tasks. If you know how to throw for example a dodgeball you can use similar movement when throwing a baseball or a handball. Because of the different sizes, compositions and weights of the different balls however, the optimal movement for each one is going to be different. But in general all three cases involve the act of throwing a round object. While throwing one of the balls cannot be done optimally for the first time by just knowing how to throw one of the other balls, the machine does not have to relearn the entire process of throwing a ball.

There are not many AI systems today aimed towards solving many different kinds of problems with one system that I am aware of. That is mostly because of the fact that just solving one specific problem with an AI is a plenty hard enough problem. However, this ability in its simplest form does not require much effort and is perhaps not even worthy of being in the list. In the ball example, if you know how to throw a dodgeball it is just a matter of trying to throw the

baseball like you throw a dodgeball, see what happens and proceed from there in your quest of learning how to throw a baseball rather than starting from scratch.

The ability to communicate the solution to a problem to another intelligent being in such a way that the other intelligent being can accurately solve the problem.

There are undoubtedly problems where a computer can tell you the solution, even guide you through the solution step by step. An instance of the Japanese game of Sudoku can be solved by a computer and it can tell you how to solve it by telling you what number to write in which square repeatedly until they entire grid is filled. While the ability actually tries to encompass the ability of communicating solutions in natural language, it must be admitted that even standardised language is a sort of communication.

The ability to learn the solution of a problem by communicating with another intelligent being.

There are certain intelligent beings capable of talking to machines. These are often called nerds, geeks or computer programmers. Computer programmers are indeed capable of teaching machines how to solve problems. It is precisely what they do. An AI capable of learning something by understanding spoken natural language would of course be desirable and would be a lot more intelligent than a machine only able to understand programming languages.

The ability to work together with another intelligent being in order to accomplish a task that was previously impossible (or simply difficult) to accomplish alone.

Ants are not very intelligent per se but do indeed work well together. Flocking algorithms are usually simple but give rise to quite complex cooperative behaviour. One might wonder if working together really is intelligent behaviour at all. To show what kind of behaviour is aimed at with this ability I shall give some examples. When you are building a house and you need to lift a stone that is too heavy for you to lift by yourself, you need to ask someone to help you lift the stone. Another example is when hunting down an animal, you would rather do it in a group than by yourself. Running a multibillion dollar international company is not something you would to do by yourself either.

There are many different types of organisational or cooperative behaviour. Hierarchical structures, favour for favour, game theory related strategies where each person tries to maximise their own success etc. There are also cultural differences regarding how *helpful* or willing to do favours for others people are and difference regarding when it is expected of someone to do a favour and so on. The essence of this ability is to be able to devise a plan that involves other intelligent beings, communicate your plan to everyone involved and argue about the conditions by which they are willing to be part of your plan. I believe this requires intelligence, and very sophisticated intelligence at that. It would require complex planning abilities which just in itself is very hard, natural language skills if hoping to involve a human in the plan which also is a very hard problem and lastly the ability to understand emotions and other complex social behaviours to be able to convince the people involved in the plan to agree to the plan. The day a machine is capable of this is the day we should start worrying about machines taking over the world.

The ability to deduce properties about objects and events by the use of observation and experiment and use these properties to solve problems.

Honda's advanced robot ASIMO is able to recognise stationary objects and avoid them while walking[9]. It is also capable of recognising faces and addressing people by name. This little Japanese robot can follow people around and can not only avoid stationary objects while walking but can even avoid moving objects. While not that intelligent yet the robot does have a basic understanding of objects. This seventh ability is certainly possible and what needs to be done

is to improve the understanding of ASIMO to come closer to that of a human's understanding of objects.

The ability to generalise previously observed properties about objects to hold for similar objects you have not previously seen to have this property.

Newer versions of ASIMO have the ability to when shown a chair, deduce general properties about it and then when shown a second, different chair assert that this is indeed also a chair[10]. This is done based on physical characteristics of the chair. This only a basic understanding of physical characteristics of objects and lacks the ability to generalise physical properties about objects. For example realising that most objects fall when dropped or most mammals cry out when in pain is not encompassed with ASIMO's ability. This kind of deduction requires more advanced intelligence that is not present in any machines today that I am aware of. This could be of the same level of difficulty as smoothly merging any specialised AIs into one, thinking AI, which will be the topic of the next section.

4.1 Combining powers

Research about the human brain indicate that different areas of the brain are responsible for different functions[11]. There are many different very complex functions of our brain that we have no control over which are done *subconsciously*. Likewise there are things we do control actively. The occipital lobe for example is responsible for visual processing and is located in the rear part of the brain. We do not have to actively decode the electrical signals sent from our eyes in to meaningful pictures and we do not have to actively recognise people or objects. This is all done by our brain subconsciously and how these things actually work is hidden to us. What you can do is choose to focus on a specific object in your field of vision and notice details about its structure, colour and patterns. You can also compare the object to other objects from your memory and thus use different parts of your brain in unison. This type of control of different parts of the brain is believed to be done in an area of the brain called the prefrontal cortex.

Many of these subconscious functions are being studied in the field of artificial intelligence today. For example face recognition, optical character recognition, walking, running, speech synthesis, speech recognition and so on. When we manage to have all these functions working satisfyingly by themselves it is this active part that connect all the different functions together, akin to the prefrontal cortex, that needs to be built if artificial intelligence shall ever have a hope of matching human intelligence.

I do believe it is possible to create an artificial intelligence and the more I read and hear about work done in artificial intelligence, brain research, behavioural studies and the like, the less credit I feel the intuitive concepts of intelligence, consciousness, understanding, self, awareness, intention and the like have.

5 Consequences

The ultimate dream concering AI for any science fiction fan is of course the AI that has surpassed human intelligence. No longer bound by the limits of our flesh the robots create a new era of artificial world domination! Many people have speculated about this type of scenario which have led to interesting ideas such as technological singularity.

5.1 Technological singularity

If we can create an intelligent being more intelligent than ourselves, then could not that intelligent being create an intelligent being even more intelligent still[12]? This is the idea behind the concept of a technological singularity. With Moore's law as inspiration one could imagine this evolution of machines creating smarter and smarter versions of themselves exploding exponentially to a point beyond imagination. This is the technological singularity, the pinnacle of technological advance. It is more wishful thinking than anything else, I believe, but a thrilling thought nonetheless.

5.2 Superintelligence

A superintelligent AI is an AI whose intelligent capabilities are in all aspects beyond that of any human. If such an AI can be created would not that AI be able to govern a country, or even the world, better than any human could? Would it not be able to judge any crime more objectively than any human judge could? There is an interesting joke about superintelligence that you might have already heard.

> A group of computer scientists build the world's most powerful computer HyperThought. HyperThought is massively parallel, it contains neural networks, it has teraflop speed and is simply amazing. The computer scientists give HyperThought a shakedown run. It easily computes billions of decimals to pi and factors a 100 digit number, simultaneously. The scientists try to find a difficult question that may stump it. Finally, one scientist exclaims: "I know!" "HyperThought," she asks "is there a God?" "There is now," replies the computer.

Most people when reading this joke assumes that the computer is asserting its own divinity, but there is a twist to this joke. If God is our creator then from the computer's perspective God would be us humans. Humans as Gods.

A superintelligent AI would per definition be able to do everything we can do, but better. The question is would we want robots to do everything for us? Then what would we do? Questions like this while interesting do not have any real importance since we have no chance of being able to predict the future. The technological development might take a completely different turn. Perhaps we genetically enhance humans and increase our intelligence that way instead while AIs remain a long forgotten dream of the past. There is no way to tell and thus questions like those will only be of importance when the first superintelligent AI is created.

> "It amuses me to imagine a day in the distant future when humans have become extinct, surpassed by our creations, robots, who roam the universe. I like to think that these robots may have a memory of us humans, perhaps as semi-mythic fractious demigods from the distant past who created them. And, just possibly, they will remember me." - Hugh Loebner

References

[1] A. Turing, "Computing machinery and intelligence," *Mind*, vol. 59, pp. 433–460, 1950.

[2] H. G. Loebner, "Home page of the loebner prize in artificial intelligence." `http://www.loebner.net/Prizef/loebner-prize.html`, 2011. [Online; accessed 13-April-2011].

[3] J. R. Searle, "Minds, Brains, and Programs," *Behavioral and Brain Sciences*, vol. 3, pp. 417–457, 1980.

[4] D. Chalmers, *The Conscious Mind: In Search of a Fundamental Theory*. Oxford: Oxford University Press, 1996.

[5] K. Gödel, "Über formal unentscheidbare Sätze der Principia Mathematica und verwandter Systeme," *Monatshefte für Mathematik und Physik*, vol. 38, no. 1, pp. 173–198, 1931.

[6] R. Penrose, *Shadows of the mind*. Oxford University Press, 1994.

[7] S. Russell and P. Norvig, *Artificial intelligence – A modern approach*. Prentice Hall, third ed., 2010.

[8] SoundHound, "About us." `http://www.soundhound.com/index.php?action=s.about`, 2011. [Online; accessed 13-April-2011].

[9] Honda, "ASIMO Technical Information." `http://asimo.honda.com/downloads/pdf/asimo-technical-information.pdf`, 2007. [Online; accessed 13-April-2011].

[10] BBC, "James May's Big Ideas – Man-Machine," 2008. Documentary.

[11] D. Rock, *Your Brain at Work*. HarperBusiness, 2009.

[12] Singularity Institute for Artificial Intelligence, Inc., "What is the Singularity?." `http://singinst.org/overview/whatisthesingularity`, 2011. [Online; accessed 31-May-2011].