

IS IT POSSIBLE TO DESIGN A MACHINE WITH MENTAL CAPACITIES?

By Gareth Oystyk

Introduction

Is it possible to design a machine with mental capacities? Some believe it is. Others believe it is either very unlikely or impossible.

Alan Turing believes machines that exhibit human behaviour have mental capacities. I shall describe his argument in detail to support the possibility that we can design a machine with mental capacities. I will describe John Searle's Chinese Room argument as a critique of Turing's position.

Hubert Dreyfus is pessimistic about the prospects of designing a machine with mental capacities. He believes that AI's failure to produce machines with mental capacities is due to false assumptions about the nature of human intelligence. Dreyfus believes that it is not possible to design a digital computer with mental capacities.

At the end of the essay, I turn to a related question, "Why do we bother attempting to design machines with mental capacities?"

This essay is divided into three parts. The first part contains arguments for the possibility of designing a machine with mental capacities. The second part contains arguments against that possibility. The third part contains a brief exploration of the question posed in the previous paragraph. Please note that throughout this essay I will substitute the phrase "artificial intelligence" with the acronym "AI".

Part One – It Is Possible

Those who believe that it is possible to design a machine with mental capacities are numerous. Computer scientists, philosophers, psychologists, and numbers of other professionals have demonstrated an interest in developing these machines. AI is a new field of research – it emerged in the late 1950s. Since then, the number of arguments in

support of AI has grown. But one person in particular originally inspired those in the AI community, and that person was Alan Turing.

In his influential paper, "Computing Machinery and Intelligence" (published 1950),ⁱ Turing clearly states that he believes that it is possible to design a machine with mental capacities. At the beginning of his paper, Turing proposes to address the question "Can machines think?" but he realizes that the question itself is ambiguous, particularly because of the ambiguity surrounding the words "machine" and "think".ⁱⁱ So he proposes another question, summarized in my own words, "Can we design a machine that will successfully play the imitation game?" This new question needs some background information. Turing proposes the following test for the attribution of mental capacities. This test is known as the imitation game. Three people play the game. The first person is a man, the second a woman and the third an interrogator of either gender (the gender is unimportant). All three individuals are put in a separate room. The interrogator can communicate with the man and the woman by using a teleprinter or some other intermediary device that will not run the risk of indicating the gender of either of the other players. The goal of the game is for the interrogator to successfully identify the woman. The woman is asked to convince the interrogator that she is a woman; however, the man is asked to do the same. The interrogator is allowed to ask any question he or she wants, such as "How do you treat a yeast infection?" or "How do you feel when you are having sex?" or even "What is your bra size?" It seems probable that the interrogator will identify the woman correctly in more than half of the cases; however, what happens when the man is replaced by a machine that has the ability to answer the same questions? Alternatively, what would we think if the interrogator only identified the woman correctly half of the time?ⁱⁱⁱ Turing believes that a machine that passes the test (i.e. causes the interrogator to guess correctly only some of the time) has mental capacities. Therefore, Turing believes that sophisticated language use is a sufficient condition for the attribution of mental capacities.

This position is plausible for many reasons. First, the one feature that seems to separate us from other animals is our ability to use sophisticated language. Although attempts have been made to teach some animals, such as primates, to use sophisticated language, it seems that these attempts have failed.^{iv} Secondly, upon reflection many people will notice that their thoughts involve language. We think by running a series of thoughts through our minds, and in doing so, we come to notice that most of these thoughts involve language, at least at the level of sophisticated, argumentative.

Next, Turing explores the prospects of designing a machine that is capable of using sophisticated language and interacting with people. Turing believes that human beings are machines and that the human brain is a certain type of biological computer; however, for the sake of argument, Turing excludes human beings from his definition of machine. Turing is also not concerned with finding an extant machine capable of success in the imitation game; he wishes to discuss whether or not there are imaginable computers that would be successful in the imitation game. Turing identifies one particular type of machine that may be capable of success, the digital computer. Turing resolves to "only permit digital computers to take part in our game."^v

What is a computer? Analog computers have three parts: the store, which stores information such as a table of instructions, the executive unit, which executes or processes the instructions, and the control, which ensures that the instructions are followed correctly. The information that is used in the digital computer is broken into packets. Digital computers ultimately encode their information in a binary fashion, using only zeros and ones. In addition, the digital computer also has the ability to receive information from outside its environment (inputs) and in turn manipulate its environment (outputs).^{vi}

How is a digital computer similar to a human computer? Turing assumes that all human actions follow fixed rules that can be articulated and translated into an intelligible procedure. Turing also assumes that human beings do not have the authority to deviate from these rules. So in order to make a computer mimic human behavior, all we must do is

articulate those procedures that guide human actions and translate them into instructions that can be followed by the digital computer.^{vii} There may be other similarities. Normally, digital computers use electricity. Our nervous system exhibits electrical activity. Perhaps this is why digital computers may have mental capacities. Turing disagrees. He cites the work of Babbage and his planned analytical engine – a digital computer, planned in the 1880's, which would have used only mechanical processes. Here, it is clear that Turing is a functionalist. Functionalists do not believe that the material used to build a computer is of any theoretical importance. Instead, they believe, as Turing states, that we should look for “mathematical analogies of function.”^{viii}

To sum up, Turing believes that all human actions can be articulated as a set of procedures that can be programmed into a digital computer, giving it the ability to mimic human behavior. If the machine successfully mimics human behavior, then it has mental capacities. So, in place of the question, “Can machines think?” Turing asks “Are there imaginable digital computers which would do well in the imitation game?”^{ix}

Turing's argument can be summarized as follows:^x

1. Behavioral similarity with human beings is a sufficient condition for a thing to have mental capacities attributed to it.
2. The behaviour of human beings can be articulated as a set of general procedures.
3. It is possible to design digital computers that can simulate the behaviour articulated in any set of general procedures.
4. Digital computers are machines.
5. Therefore, it is possible to design machines that can simulate the behavior of human beings.
6. Therefore, it is possible to design a machine that has mental capacities attributed to it.

Please note my use of “attribute” in the above argument. It is not possible to know for certain that a machine has mental capacities, just as it is not possible to know for certain that another human being has

mental capacities. That is why it is more accurate to say that we can “attribute” mental capacities to a machine.

Turing’s position is somewhat compelling; however, one argument in particular has threatened Turing’s position seriously. This is John Searle’s Chinese Room Argument. It can be found in his paper “Minds, Brains and Programs.”^{xj} Like Turing, Searle concedes that human beings are machines of a biological sort. But he wonders whether or not Turing’s functionalism is a good theory. Is it really the case that the nature of the computer’s materiality is of no theoretical importance? Furthermore, he wonders if we would be justified to conclude that a machine has mental capacities because it simulates human behaviour. Searle summarizes Turing’s position in the following way: computer programs that simulate human cognition are not models of minds, but actual minds.^{xii}

Within the AI community, Searle identifies two distinct groups: Strong AI and Weak AI. Those associated with strong AI believe that programs are minds and that thinking is merely the manipulation of symbols (syntax). Those associated with weak AI believe that computer programs are only models of human minds but not actual minds. In his paper, Searle only critiques strong AI. Searle does not believe that computer programs alone are a sufficient condition for mental capacities. He also doubts that digital computers will ever be able to exhibit mind-like qualities. He demonstrates this in his Chinese Room Argument.^{xiii}

Suppose you are placed in a room with two baskets. Each basket contains a number of Chinese symbols. Two Chinese speakers wait separately outside the room and pass symbols to you through slots in the wall. The two Chinese speakers wish to use you to communicate with each other. You are provided with a rulebook, written in English, which instructs you to match certain Chinese symbols in one basket with others in the other basket. The rulebook, however, does not explain what these symbols mean in English. A symbol is passed through to you through one slot. You match it up with the proper symbol and pass the latter back out the slot. The baskets of symbols are

analogous to a database or the store in a digital computer. The rulebook is the program. Those who wrote the book are the programmers. The symbols that come in from the slots are the inputs; those you pass out are the outputs. Suppose the rulebook was so well written and suppose you were so quick matching symbols that your answers were indistinguishable from those of a Chinese speaker. According to Turing a machine that can successfully play the imitation game has mental capacities. If a machine has mental capacities, it should understand the language that it is using. But if you were in this Chinese Room, do you think you would understand Chinese? I assume that you would not, for you are only manipulating meaningless symbols based on a set of rules. Because the essence of digital computers is to manipulate symbols according to exact rules, it seems that the digital computer cannot understand language, and more generally, cannot possess mental capacities. The difference between digital computers and human computers is that digital computers use only syntax, whereas human computers use both syntax and semantics.^{xiv}

Searle lists several axioms or premises before he reaches a number of conclusions:

1. Computer programs are formal (syntactic).
2. Human minds have mental contents (semantics).
3. Syntax by itself is neither constitutive of nor sufficient for semantics.
4. Brains cause minds.

Conclusion 1: Computers are neither constitutive of nor sufficient for minds.

Conclusion 2: Any other system capable of causing minds would have to have causal powers (at least) equivalent to those of brains.

Conclusion 3: Any artifact that produced mental phenomena, any artificial brain, would have to duplicate the specific causal powers of brains, and it could not do that just by running a formal program.

Conclusion 4: The way that human brains actually produce mental phenomena cannot be solely by virtue of running a computer program.^{xv}

These axioms and conclusions indicate two areas where Searle differs from Turing. Searle believes that material composition of the computer is of great theoretical importance (brains cause minds), but that we could, nevertheless, “come to be able to create thinking systems artificially.”^{xvi} Secondly, Searle notes that the simulation of mental capacities is no different than simulating digestion. When a computer simulates digestion, there is not any real food being digested anywhere.^{xvii} However, like Turing, Searle is somewhat optimistic that in the distant future we may produce artifacts with mental capacities. But he is very critical of the importance placed on the digital computer as the instrument that will reach that goal.

How can Turing reply to Searle? Consider the following question, “How do we know that other human beings, beside ourselves, have mental capacities?” There are no instruments that can measure or capture mental capacities (thoughts, emotions, and sensations); we can only measure physiological processes in our brains. For convenience’s sake, we always infer that other human beings have mental capacities like us, primarily because they act like us and have bodies like us. This is a weak solution to what is known as the Problem of Other Minds, but it is the only one we have. So Turing is right to identify his test as a mental capacity detector. It would be a form of speciesism to attribute mental capacities to only human beings and not machines if a computer that exhibited human behaviour and appearance existed.

Suppose we got our hands on a fresh human cadaver that was missing its brain. We could build a digital computer sophisticated enough to control that body in such a way that it behaved as any other human being. We could place this computer in the skull cavity of the cadaver and connect it up to the central nervous system. This would create a cybernetic organism – a cyborg. We turn the computer on and the cyborg comes to life. A moment later, due to our excitement, we have heart attacks and die. The cyborg knocks over a Bunsen burner on

his way out of the laboratory and the laboratory catches fire, destroying all evidence of the cyborg's history. This being proceeds to live among human beings for its entire life and is never detected. It makes friends and even a family (it still possesses reproductive organs). The cyborg's lover and friends would never believe it to be something other than a normal human being. Then a mysterious visitor arrives who explains that the cyborg is really just a complicated computer. His friends and family would passionately argue that he has mental capacities. One could hold that failing to attribute mental capacities to a machine that behaves like us is merely arbitrary because we attribute mental capacities to other human beings that behave like us every day. And furthermore, the goal of AI seems to be to produce machines that behave like we do. Whether or not they really possess mental capacities, I believe, is irrelevant.

Part Two – It Is Not Possible

Many find the previous arguments convincing. It is only a matter of time, some say, until we design a machine with mental capacities. But Hubert Dreyfus believes otherwise. He is particularly critical of AI research involving digital computers, where in fact, most researchers and scientists in the AI community use digital computers. Dreyfus takes a phenomenological stance, similar to that of the French philosopher, Merleau-Ponty, who believes that embodiment is a necessary condition for mental capacities. Dreyfus also places an emphasis on the importance of intuition in skill development, and on the holistic nature of sense perception. Digital computers seem to lack these attributes.

In his article "Misrepresenting Human Intelligence,"^{xviii} Dreyfus critiques the assumptions made by those in the AI community. He believes that AI makes "overly ambitious goals and predictions."^{xix} He identifies the basic project of AI as the production of a machine that possesses intelligence equal to or better than human beings. AI has become a "degenerating research program" because it rests on shaky assumptions.^{xx}

What sort of assumptions does AI rest upon and why are they problematic? Turing, for example, assumes that human knowledge is constructed out of meaningless bits of sense data and then rearranged or organized according to rule-governed operations. But Dreyfus argues that human knowledge does not consist of such an arrangement. He believes that any system built upon such an assumption will never display intelligent behavior. Dreyfus does not identify the lack of speed or storage capacity as the source of the current problems in AI, but the concepts around which the machines are being built. These concepts include academically rejected yet popularly assumed theories of human mentality. Two general assumptions are “mental processes are sequences of rule-governed operations” and “these operations are carried out on determinate data which represent facts or features in the world.”^{xxi}

In order to show the shortcomings of digital computers, Dreyfus raises several examples. Suppose someone utters the phrase “the book is in the pen.” This phrase is ambiguous – it has several meanings. The only way to find out exactly what this sentence means is to place it in a context. If I am a farmer and I utter this phrase, I may be indicating that I dropped my diary in the pigpen. If I am a new parent and I utter this phrase, I may be indicating that I left a storybook in my infant’s playpen. There is no way to discover what the phrase means unless it is uttered in a context – a situation.^{xxii} It is always the case that we are in a situation in so far as we have bodies.^{xxiii} We cannot completely cut ourselves off from the rest of the world. The permanence of our bodies ensures our location in a situation. It limits the range of possibilities available for making decisions, which makes the decision process faster and reduces the amount of time it takes to identify relevant features in the environment. Dreyfus indicates that our situation is not determined by a set of context-free facts about the environment or the set of beliefs that we bring to that situation. Every situation grows out of previous situations. Because we are embodied, we are always situated and every situation blends backwards into a previous situation in the style of an infinite regress.^{xxiv} We cannot conceive ourselves as being without a

situation; we cannot exist independently from our bodies, and therefore, we cannot exist before our bodies do, before we are in a situation. Digital computers lack a situation similar to the situation of a human being. If they lack a situation, they must lack the sort of mental capacities enjoyed by human beings.

Many in the AI community have identified practical uses for thinking machines. There are plans to design machines as experts. These computers can contain all available knowledge on the expert task they are designed to perform, so it seems that they may be better at identifying problems because they will not overlook any information. But can a computer ever become an expert in the popular sense of the term? Dreyfus thinks it is unlikely.^{xxv} In order to show this, he outlines the process through which beginners become experts.

Stage One: The Novice.^{xxvi} The novice is given the rules or relevant facts required as knowledge for that particular skill. However, they lack a sense of the overall goal. In this stage, a computer can easily be constructed that surpasses the skill of the novice, for it may be possible to design a computer that mimics any articulated type of behavior.

Stage Two: The Advanced Beginner. The advanced beginner gains experience by applying skills in certain situations. The advanced beginner uses previous success or failures to recognize relevant features in new situations. The advanced beginner does not approach every situation anew in such a way that all rules are applied in the way they were originally articulated. Dreyfus believes that this is easy to do for human beings, but impossible to do for computers. It does not seem to be the case that computers bring their past experience to a new situation. They simply treat every situation as the first situation and execute the same programs or procedures they originally did. In this sense, it makes it difficult for computers to adapt to new situations.

Stage Three: Competent. The competent person has a lot of experience and begins to order the important or relevant features of the situation according to a plan. When the competent person decides on a plan and executes it, he or she feels responsible for the outcomes of the

plan. Success brings pleasure to the competent person; failure is unpleasant and generally not forgotten. The results of that personal choice are brought into the next situation as limits to those new possibilities. I believe that it is unlikely that a computer actually experiences the feeling success or failure because computers must always follow instructions. The following pseudo-code may illuminate my idea:

- If asked, "What sandwich would you like?" utter the phrase "I would like a tuna sandwich"
- If you receive a tuna sandwich, move robotic mouth to smiling position.
- If you do not receive a tuna sandwich, but receive another sandwich, utter the phrase "I do not like that sandwich. In fact, I would rather eat a dog."
- If you do not receive a sandwich at all, utter the phrase "Garçon, what kind of service is this anyway?"

If the robot receives a tuna sandwich, it will execute the proper instructions. If it does not receive a tuna sandwich or any sandwich at all, it will also execute the proper instructions. A digital computer must always follow instructions. It will never fail to follow instructions. It will always succeed, in so far as it always follows instructions. I find it difficult to conceive of myself living without both successes and failures. Success and failure imply each other. I need to experience the pain of failure in order to experience the pleasure of success. A computer could never experience the pleasure of success because it is always succeeding. It would always feel the same way; there would be nothing to compare the feeling of success with.

Stage Four: Proficiency. The proficient person is very involved in a given situation. They often rely on intuition. But there are times when the proficient person will reflect on the situation due to attempts to justify his or her actions.

Stage Five: Expertise. The expert often acts purely from intuition. For example, people that determine the gender of baby chickens cannot see the chicken's sex organs, they simply rely on intuition to perform

their task, yet experts at sexing chicks are very effective.^{xxvii} How can this be explained? Skills become part of the expert. Merleau-Ponty is useful here. He calls this process "sedimentation."^{xxviii} Skills become part of our body, part of the habitual body. The habitual body is that aspect of our existence that absorbs repetitive tasks in such a way that we can divert our consciousness to more important things.^{xxix} When learning to ride a bicycle, people must first concentrate on the task at hand, the position of their limbs, the distribution their weight, the texture of the road, the speed they are moving at. But after years of experience, little attention is required when riding a bicycle. In fact, many people can have a conversation with themselves while riding a bicycle and not pay any attention to the complicated task they are performing. Try to teach a machine to ride a bicycle while talking to itself. Merleau-Ponty and Dreyfus both mention that the tools used by the expert become enveloped by the body. Chess pieces become extensions of the chess master's body; the bicycle becomes an extension of the person on the bicycle. But I must clarify this odd statement. The body as experienced, according to Merleau-Ponty, is not material or spatial; it is intentional.^{xxx} We exercise our intentions through our situated bodies. Instruments are not experienced as additions to the material body, but become incorporated into our bodies in so far as they exhibit significance in relation to a particular goal. The body envelops the bicycle because it signifies a mode of motility in respect to the intention of the person on the bicycle. Chess pieces are enveloped because of their significance as subversive elements in the field of force relations that exists between pieces on a chessboard.^{xxxi} Skills belong to the habitual body or as Dreyfus calls it, the intuition. Because the habitual body frees up our consciousness (our reflective being), the habitual body is necessarily pre-reflective, or unconscious. And the nature of pre-reflective being is such that it cannot be reflected upon. In other words, reflecting upon pre-reflective experience makes it reflective.^{xxxii} It is impossible to grasp the processes that go on in the pre-reflective aspect of our existence. Therefore, it is impossible to articulate the rules, if they exist, which govern those pre-reflective

activities. One of the supposed virtues of the digital computer is that it can simulate any process that can be articulated. So if those in the AI community wish to design machines that are capable of being experts, they will certainly fail, because we cannot articulate all the rules that govern the actions of human beings.

Dreyfus believes we should attempt to design computers that do not rely on analytic, rule-governed behavior, but on concrete, lived experience and intuition. The assumption that human minds are digital is very problematic and in many ways contradictory to our lived experience. Dreyfus does not think that it is possible to design digital computers with mental capacities. No amount of speed or memory will suffice because "our performance is entirely different in kind from that of a digital computer."^{xxxiii}

The following is a summary of Dreyfus' argument:^{xxxiv}

1. Behavioral similarity with human beings is a necessary condition for the attribution of mental capacities.
2. Some of the behaviour of human beings cannot be articulated as a set of procedures.
3. Digital computers can only produce behaviour that is articulated in a set of procedures.
4. Digital computers are machines.
5. Therefore, it is not possible to design a machine that can simulate all the behaviour of human beings.
6. Therefore, it is not possible to design a machine that has mental capacities attributed to it.

I find Dreyfus' arguments very compelling, primarily because of the emphasis he places on the non-digital and intuitive nature of our experience. The methodology of phenomenological philosophers is such that they describe experience and prompt you to ask yourself if their description matches yours. I am more inclined to agree with Dreyfus and Merleau-Ponty that my mental experience contains intuition and the permanence of my body rather than digital processes or database queries. However, there is a problem with the phenomenology Dreyfus adopts. He relies on descriptions of concrete,

lived experience. But at the same time, like Merleau-Ponty, he supposes that these concrete examples are universal in character. Although Dreyfus describes our experience in such a way that we agree with him, it is impossible to certify that all experience occurs in that fashion. Machines may in fact experience the world, but in a way radically different way than we do. However, the goal of AI is to design machines that think and behave as we do. So even if machines did experience the world, albeit in a way completely foreign to us, their experience is of no concern for those in the AI community.

Part Three – Why Bother?

Why should we bother designing machines with mental capacities? This part of my essay is a very brief exploration of this important question.

It is apparent that digital computers have affected our world in important ways. Information is more readily available than ever. Medical technology has benefited tremendously from computers. Computer modeling allows us to study a wide variety of natural phenomena. Computers are useful, so I do not argue that we should stop using and improving upon them, but why should we bother designing computers that think?

Suppose an inventor spent his entire life isolated in his laboratory working on a top-secret project. Days before he dies, he decides to present his invention. Journalists and important figures in the scientific community come to witness this event. Now suppose the inventor unveils his achievement and to the surprise of many, the inventor presents a stone wheel. Most people would balk at his so-called achievement. After all, the wheel was invented thousands of years ago and since then, we have improved upon its design by encasing it in rubber and making out of different material. Now suppose a team of AI experts discovers a way to produce a thinking machine. How is their achievement any different that the inventor who re-invented the wheel? There are already six billion human beings (six billion thinking machines) on this planet, so why would one expect that the invention of

a non-biological machine with mental capacities would be so momentous? My point is this: why bother inventing something that already exists?

It is now evident that the prospects of designing a machine with mental capacities are grim. It will take a great deal of time, money and resources to design a machine with mental capacities, if it is even possible. I feel that our resources would go to better use by improving human mental capacities. This is already being pursued in several ways. Education, psychological conditioning, chemicals and drugs have been used to enhance our mental capacities. Currently, the possibility of changing the structure of the brain through genetic engineering is also realizable.

But what would happen if the AI community succeeded? If the goal of AI was to create thinking machines for use in military combat, for dangerous ocean or space exploration, for sexual pleasure, or for hard labour, I believe that inevitably, many would find these practices highly unethical and analogous to the treatment of human slaves. No one would bother re-inventing the wheel, why bother re-inventing thinking things?

Conclusion

None of the philosophers I studied completely ruled out the possibility that we may be able to design a machine with mental capacities. This is a difficult assertion to make. However, both Searle and Dreyfus show the limits of digital computers. We will have to approach the problem very differently, I believe, should we choose to undertake this task any further.

Bibliography

Dreyfus, Hubert. "Misrepresenting Human Intelligence." *Artificial Intelligence: The Case Against*. Ed. Rainer Born. New York: St. Martin's Press, 1987. 41-54.

Dreyfus, Hubert and Stuart Dreyfus. *Mind Over Machine*. New York: The Free Press, 1986.

Dreyfus, Hubert and Stuart Dreyfus. "Putting Computers in Their Place." *Social Research*, Spring 1986. 57-76.

Merleau-Ponty, M. *Phenomenology of Perception*. Trans. Colin Smith. New York: Routledge, 1999.

Searle, John. "Is the Brain's Mind a Computer Program?" *Reason at Work: Introductory Readings in Philosophy*. Eds. Steven M. Cahn et al. 3rd ed. Hardcourt Brace & Company, 1996. 753-764.

Searle, John. "Minds, Brains and Programs." *Artificial Intelligence: The Case Against*. Ed. Rainer Born. New York: St. Martin's Press, 1987. 18-40.

Turing, Alan "Computing Machinery and Intelligence." *Minds and Machines*. Ed. Alan Ross Anderson. Englewood Cliffs: Prentice-Hall, 1964. pp. 4-30.

ⁱ A.M. Turing, "Computing Machinery and Intelligence," *Minds and Machines*, Ed. Alan Ross Anderson, (Englewood Cliffs: Prentice-Hall, 1964), 4-30.

ⁱⁱ *Ibid.*, 4.

ⁱⁱⁱ *Ibid.*, 5.

^{iv} See: E.S. Savage-Rumbaugh, et al. "Symbolic Communication Between Two Chimpanzees," *Science*, vol. 201 (18 August, 1978), 641-644.

^v *Ibid.*, 7.

vi Ibid., 8-10.

vii Ibid., 8.

viii Ibid., 10.

ix Ibid., 13.

x This argument is a slight modification of the one outlined in "Notes for Phil 304A" written by Dr. C. Morgan for PHIL 342A, F01, 2000.

xi John Searle, "Minds, Brains and Programs," *Artificial Intelligence: The Case Against*, Ed. Rainer Born, (New York: St. Martin's Press, 1987), 18-40.

xii John Searle, "Is the Brain's Mind a Computer Program?," *Reason at Work: Introductory Readings in Philosophy*, Eds. Steven M. Cahn et al., 3rd ed., (Harcourt Brace & Company, 1996), 754.

xiii Ibid., 754.

xiv Ibid., 755.

xv Ibid., 755-759.

xvi Ibid., 757.

xvii Ibid., 759.

xviii Hubert Dreyfus, "Misrepresenting Human Intelligence," *Artificial Intelligence: The Case Against*, Ed. Rainer Born, (New York: St. Martin's Press, 1987), 41-54.

xix Ibid., 41.

xx Hubert Dreyfus and Stuart Dreyfus, "Putting Computers in Their Place," *Social Research*, Spring 1986, 57-76.

xxi Dreyfus, "Misrepresenting Human Intelligence," 41-43.

xxii Ibid., 43.

xxiii Maurice Merleau-Ponty, *Phenomenology of Perception*, Trans. Colin Smith, (New York: Routledge, 1999), 82.

xxiv Ibid., 44.

xxv Ibid., 47-48.

xxvi The five stages of skill development can be found on pages 48-52 in: Dreyfus, "Misrepresenting Human Intelligence," and pages 66-72 in: Dreyfus, "Putting Computers in Their Place."

xxvii Hubert Dreyfus and Stuart Dreyfus, *Mind Over Machine*, (New York: The Free Press, 1986), 197.

xxviii Merleau-Ponty, *Phenomenology of Perception*, 130.

xxix Ibid., 143-144.

xxx Ibid., 136.

xxxi Dreyfus, "Misrepresenting Human Intelligence," 51.

xxxii Merleau-Ponty, *Phenomenology of Perception*, 79-81.

xxxiii Dreyfus, "Misrepresenting Human Intelligence," 53.

xxxiv This argument is a slight modification of the one outlined in "Notes for Phil 304A" written by Dr. C. Morgan for PHIL 342A, F01, 2000.