

Subcognition and the Limits of the Turing Test

Robert M. French

Psychology Department, University of Liège, Liège, BE

email: rfrench@ulg.ac.be

Originally published in: *Mind* (1990) 99(393), 53-65

Introduction

Alan Turing, in his original article¹ about an imitation-game test of intelligence, seems to be making two separate claims. The first, the philosophical claim, is that if a machine could pass the Turing Test, it would necessarily be intelligent. This claim I believe to be correct². His second point, the pragmatic claim, is that in the not-too-distant future it would in fact be possible to actually build such a machine. Turing clearly felt that it was important to establish both claims. He realized, in particular, that if one could rigorously show that *no* machine could ever pass his test, his philosophical point, while still true, would lose a great deal of its significance. He thus devoted considerable effort to establishing not only the philosophical claim but also the pragmatic claim.

Ever since his article appeared most philosophers have concentrated almost exclusively on attacking or defending the philosophical claim. There are those who believe that passing the Turing Test constitutes a sufficient condition for intelligence and those who do not. The philosophical importance of this first claim is that it provided a clean and novel test for intelligence that neatly sidestepped the vast philosophical quagmire of the mind-body problem. The philosophical claim translates elegantly into an operational definition of intelligence: Whatever *acts* sufficiently intelligent *is* intelligent.

However, in this paper I will take issue with Turing's pragmatic claim, arguing that the Turing Test's very capacity to probe the deepest, most essential areas of human cognition makes it virtually useless as a real test for intelligence. I strongly disagree with Hubert Dreyfus' claim, for example, that "as a goal for those actually trying to construct thinking machines, and as a criterion for critics to use in evaluating their work, Turing's test was just what was needed."³ We will see that the Turing Test could be passed only by things that have experienced the world as we have experienced it, and this leads to the central point of the present paper, namely, that the Test provides a guarantee not of intelligence but of culturally-oriented *human* intelligence.

I establish this consequence of the Turing Test by first proposing a first set of "subcognitive" questions that are explicitly designed to reveal low-level cognitive structure. Critics might object that there is something unfair about this type of question and suggest that it be disallowed. This leads to another important claim of this paper, which is that in fact, there is no way to distinguish questions that are subcognitive from those that are not. To support this claim, I present another class of questions that seem at first glance to be "cognitive", but in reality, prove to be every bit as dependent on unconscious mechanisms as the initial class of questions. Close examination of some of Turing's original questions reveals that they, too, are subcognitive. In like manner, any sufficiently broad set of questions making up a Turing Test would necessarily contain questions that rely on subcognitive associations for their answers. I will show that it is impossible to tease apart "subcognitive" questions from ones that are not. From this it follows that the cognitive and subcognitive levels are inextricably intertwined.

It is this essential inseparability of the subcognitive and cognitive levels -- and, for that matter, even the physical and cognitive levels -- that makes the Turing Test a test for *human* intelligence, not intelligence in general. This fact, while admittedly interesting, is not particularly useful if our goal is to gain insight into intelligence in general. But if we cannot use the Turing Test to this end, it may turn out that the best (or possibly only) way of discussing general intelligence will be in terms of categorization abilities, the capacity to learn new concepts, to adapt old concepts to a new

environment, and so on. Perhaps what philosophers in the field of artificial intelligence need is not simply a *test* for intelligence but rather a *theory* of intelligence. The precise elements of this theory are, as they were in 1950 when Turing proposed his imitation-game test, still the subject of much controversy.

On Nordic Seagulls

Consider the following parable: It so happens that the only flying animals known to the inhabitants of a large Nordic island are seagulls. Everyone on the island acknowledges, of course, that seagulls can fly. One day the two resident philosophers on the island are overheard trying to pin down what "flying" is really all about.

Says the first philosopher, "The essence of flying is to move through the air."

"But you would hardly call this flying, would you?" replies the second, tossing a pebble from the beach out into the ocean.

"Well then, perhaps it means to remain aloft for a certain amount of time."

"But clouds and smoke and children's balloons remain aloft for a very long time. And I can certainly keep a kite in the air as long as I want on a windy day. It seems to me that there must be more to flying than merely staying aloft."

"Maybe it involves having wings and feathers."

"Penguins have both, and we all know how well they fly . . ."

And so on. Finally, they decide to settle the question by, in effect, avoiding it. They do this by first agreeing that the only examples of objects that they are absolutely certain can fly are the seagulls that populate their island. They do, however, agree that flight has something to do with being airborne and that physical features such as feathers, beaks, and hollow bones probably are superficial aspects of flight. On the basis of these assumptions and their knowledge of Alan Turing's famous article about a test for intelligence, they hit upon the Seagull Test for flight. The Seagull Test is meant to be a very rigorous sufficient condition for flight. Henceforth, if someone says, "I have invented a machine that can fly," instead of attempting to apply any set of flight-defining criteria to the inventor's machine, they will put it to the Seagull Test. The *only* things that they will certify with absolute confidence as being able to fly are those that can pass the Seagull Test. On the other hand, they agree that if something fails the Test, they will not pass judgment; maybe it can fly, maybe it can't.

The Seagull Test works much like the Turing Test. Our philosophers have two three-dimensional radar screens, one of which tracks a real seagull; the other will track the putative flying machine. They may run any imaginable experiment on the two objects in an attempt to determine which is the seagull and which is the machine, but they may watch them only on their radar screens. The machine will be said to have passed the Seagull Test for flight if both philosophers are indefinitely unable to distinguish the seagull from the machine.

An objection might be raised that some of their tests (for example, testing for the ability to dip in flight) might have nothing to do with flying. The philosophers would reply: "So what? We are looking for a sufficient condition for flight, not a *minimal* sufficient condition. Furthermore, we understand that ours is a very hard test to pass, but rest assured, inventors of flying machines, failing the Test proves nothing. We will not claim that your machine *cannot* fly if it fails the Seagull Test; it may very well be able to. However we, as philosophers, want to be absolutely certain we have a true case of flight, and the only way we can be sure of this is if your machine passes the Seagull Test."

Now, of course, the Seagull Test will rightly take bullets, soap bubbles, and snowballs out of the running. This is certainly as it should be. But helicopters and jet airplanes -- which *do* fly -- would also never pass it. Nor, for that matter, would bats or beetles, albatrosses or hummingbirds. In fact, under close scrutiny, probably only seagulls would pass the Seagull Test, and maybe only seagulls from the philosophers' Nordic island, at that. What we have is thus not a test for flight at

all, but rather a test for flight as practiced by a Nordic seagull.

For the Turing Test, the implications of this metaphor are clear: an entity could conceivably be extremely intelligent but, if it did not respond to the interrogator's questions in a thoroughly human way, it would not pass the Test. The *only* way, I believe, that it would have been able to respond to the questions in a perfectly human-like manner is to have experienced the world as humans have. What we have is thus not a test for intelligence at all, but rather a test for intelligence as practiced by a human being.

Furthermore, the Turing Test admits of no degrees in its sufficient determination of intelligence, in spite of the fact that the intuitive human notion of intelligence clearly does. Spiders, for example, have little intelligence, sparrows have more but not as much as dogs, monkeys have still more but not as much as eight-year-old humans, who in turn have less than adults. If we agree that the underlying neural mechanisms are essentially the same across species, then we ought to treat intelligence as a continuum and not just as something that only humans have. It seems reasonable to ask a good test for intelligence to reflect, if only approximately, those differences in degree. It is especially important in the study of artificial intelligence that researchers not treat intelligence as an all-or-nothing phenomenon.

Subcognitive Questions

Before beginning the discussion of subcognitive questions, I wish to make a few assumptions that I feel certain Turing would have accepted. First, I will allow the interrogator to poll humans for the answers to some of the questions prior to posing them during the imitation game itself. (I will call the humans who are polled the "interviewees".) I also want to make explicit an assumption that is tacit in Turing's article, namely that the human candidate and the interrogator (and, in this case, the interviewees) are all from the same culture and that the computer will be attempting to pass as an individual from that culture. Thus, if ever the computer replies, "I don't speak English" or something of the sort, the interrogator will immediately deduce, rightly, that the other candidate is the human being. Finally, while I believe that it is *theoretically* possible to build a machine capable of experiencing the world in a manner indistinguishable from a human being -- a machine that can fall off bicycles, be scratched by thorns on roses, smell sewage, and taste strawberries --, I will assume that no computer is now, or will in the foreseeable future be, in a position to do so.

I will designate as *subcognitive* any question capable of providing a window on low-level (i.e., unconscious) cognitive structure. By "low-level cognitive structure", I am referring, in particular, to the subconscious associative network in human minds that consists of highly overlapping activatable representations of experience. This is the level currently being explored by new approaches to cognitive modelling.⁴

The first class of questions is explicitly designed to reveal low-level cognitive structure (and I think everyone will agree that they do so). I will respond to the anticipated objection that these explicitly subcognitive questions are unfair by following up with another set of questions that seem, at first glance, to be at a higher cognitive level than the first set. These questions will turn out, under closer examination, to be subcognitive also. I will conclude with a final set of questions that seem uncontestedly to be innocent high-level cognitive questions but that will be just as hard as the others were for the computer to answer in the way a human would.

Associative Priming

This first set of questions is based on current research on associative priming, often called semantic facilitation. The idea is the following: Humans, over the course of their lives, develop certain associations of varying strength among concepts. By means of the so-called lexical decision task it has been established⁵ that it requires less time to decide that a given item is a word when that item is preceded by an associated word. If, for example, the item "butter" is preceded by the word

"bread", it would take significantly less time to recognize that "butter" was a word than had an unassociated word like "dog" or a nonsense word preceded it.

The Turing Test interrogator makes use of this phenomenon as follows: The day before the Test, she selects a set of words (and non-words), runs the lexical decision task on the interviewees and records average recognition times. She then comes to the Test armed with the results of this initial test, asks both candidates to perform the same task she ran the day before, and records the results. Once this has been done, she identifies as the human being the candidate whose results more closely resemble the average results produced by her sample population of interviewees.

The machine would invariably fail this type of test because there is no *a priori* way of determining associative strengths (i.e., a measure of how easy it is for one concept to activate another) between *all* possible concepts. Virtually the only way a machine could determine, even on average, all of the associative strengths between human concepts is to have experienced the world as the human candidate and the interviewees had.

A further example might help to illustrate the enormous problem of establishing the associative weights between concepts in an *a priori* manner. Certain groups of concepts, say, the steps in baking a cake, are profoundly sequential in nature. The associative strengths between sequentially related concepts involved in baking a cake (opening the flour bin, breaking the eggs, mixing the flour and eggs, putting the mixture in the oven, setting the oven temperature, removing a baked cake) are profoundly dependent on the human experience of cake-baking. Even if we made the assumption that concepts like "removing a cake from an oven", "breaking eggs", "setting oven temperature", and so on could be explicitly programmed into our computer, the associative strengths among these concepts would have to reflect the temporal order in which they normally occurred in human experience if the machine were to pass the Turing Test. We would have to be able to set these strengths in an *a priori* manner, not only for category sequences associated with cake-baking, but also between the concepts of *all* the concept sequences experienced by humans. While this may be theoretically possible, it would certainly seem to be very implausible.

Now, suppose a critic claims that these explicitly subcognitive questions are unfair because – ostensibly, at least – they have nothing to do with intelligence; they probe, the critic says, a cognitive level well below that necessary for intelligence and therefore they should be disallowed. Suppose, then, that we obligingly disallow such questions and propose in their stead a new set of questions that seem, at first glance, to be at a higher cognitive level.

Rating Games

Neologisms will form the basis of the next set of questions, which we might call the Neologism Rating Game. Our impressions involving made-up words provide particularly impressive examples of the "unbelievable number of forces and factors that interact in our unconscious processing of even . . . words and names only a few letters long".⁶ Consider the following set of questions, all having a totally high-level cognitive appearance:

On a scale of 0 (completely implausible) to 10 (completely plausible), please rate:

- 'Flugblogs' as a name Kellogg's would give to a new breakfast cereal.
- 'Flugblogs' as the name of a new computer company.
- 'Flugblogs' as the name of big, air-filled bags worn on the feet and used to walk on water.
- 'Flugly' as the name a child might give its favorite teddy bear.
- 'Flugly' as the surname of a bank accountant in a W.C. Fields movie.
- 'Flugly' as the surname of a glamorous female movie star.

The interrogator will give, say, between fifty and one hundred questions of this sort to her interviewees⁷, who will answer them. Then, as before, she will give the same set of questions to

the two candidates and compare their results to her interviewees' averaged answers. The candidate whose results most closely resemble the answers given by the polled group will almost certainly be the human.

Let us examine a little more closely why a computer that had not acquired our full set of cultural associations would fail this test. Consider "Flugbogs" as the name of a breakfast cereal. It is unquestionably pretty awful. The initial syllable "flug" phonetically activates (unconsciously, of course) such things as "flub", "thug", "ugly", or "ugh!", each with its own aura of semantic connotations. "Blogs", the second syllable, activates "blob", "bog", and other words, which in turn activate a halo of other semantic connotations. The sum total of this spreading activation determines how we react, at a conscious level, to the word. And while there will be no precise set of associated connotations for all individuals across a culture, on the whole there is enough overlap to provoke *similar* reactions to given words and phrases. In this case, the emergent result of these activations is undeniable: "Flugbogs" would be a lousy name for a cereal (unless, of course, the explicit *intent* of the manufacturer is to come up with a perverse-sounding cereal name!)

What about "Flugly" as a name a child might give its favorite teddy bear? Now *that* certainly sounds plausible. In fact, it's kind of cute. But, on the surface at least, "Flugbogs" and "Flugly" seem to have quite a bit in common; if nothing else, both words have a common first syllable. But "Flugly", unlike "Flugbogs", almost certainly activates "snugly" and "cuddly", which would bring to mind feelings of coziness, warmth, and friendship. It certainly also activates "ugly", which might normally provoke a rather negative feeling, but, in this case, there are competing positive associations of vulnerability and endearment activated by the notion of children and things that children like. To see this, we need look no further than the tale of the Ugly Duckling. In the end, the positive associations seem to dominate the unpleasant sense of "ugly". The outcome of this subcognitive competition means that "Flugly" is perceived by us as being a cute, quite plausible name for a child's teddy bear. And yet, different patterns of activations rule out "Flugly" as a plausible name for a glamorous female movie star.

Imagine, for an instant, what it would take for a computer to pass this test. To begin with, there is no way it could look up words like "flugly" and "flugbogs": they don't exist. To judge the appropriateness of any given word (or, in this case, nonsense words) in a particular context requires taking unconscious account of a vast number of culturally-acquired, competing associations triggered initially by phonetic resemblances. And, even though one might succeed in giving a program a certain number of these associations (for example, by asking subjects questions similar to the ones above and then programming the results into the machine), the space of neologisms is virtually infinite. The human candidate's reaction to such made-up words is an emergent result of myriad subcognitive pressures, and unless the machine had a set of associations similar to those of humans both in degree and in kind, its performance in the Rating Game would necessarily differ more from the interviewees' averaged performance than would the human candidate's. Once again, a machine that had not experienced the world as we have would be unmasked by the Rating Game, even though the questions comprising it seemed, at least at the outset, so cognitively high-level in nature.

If, for some reason, the critics were still unhappy with the Neologism Rating Game using made-up words, we could consider a variation on the game, the Category Rating Game,⁸ in which all of the questions would have the form: "Rate **X**s as **Y**s" (0 = "could be no worse", 10 = "could be no better") where **X** and **Y** are any two categories. Such questions give every appearance of being high-level cognitive questions: they are simple in the extreme and rely not on neologisms but on everyday words. For example, we might have, "Rate **dry leaves** as **hiding places**". Now, clearly no definition of "dry leaves" will ever include the fact that piles of dry autumn leaves are wonderful places for children to hide and, yet, few among us would not make that association upon seeing the juxtaposition of those two concepts. There is therefore some overlap, however implausible this might seem a priori, between the categories of "dry leaves" and "hiding places". We might give dry leaves a rating of, say, 4 on a 10-point scale. Or, another example, "Rate

radios as musical instruments." As in the previous example, people do not usually think of radios as musical instruments, but they do indeed have some things in common with musical instruments: both make sounds; both are designed to be listened to; John Cage once wrote a piece in which radios were manipulated by performers; etc. Once again, there is therefore some overlapping of these two categories; as a musical instrument, therefore, we might give a radio a rating of 3 or even 4 on a 10-point scale.

The answer to any particular rating question is necessarily based on how we view the two categories involved, each with its full panoply of associations, acquired through experience, with other categories. A list of such questions might include:

- "Rate **banana splits** as **medicine**",
- "Rate **grand pianos** as **wheelbarrows**",
- "Rate **purses** as **weapons**",
- "Rate **pens** as **weapons**",
- "Rate **jackets** as **blankets**",
- "Rate **pine boughs** as **mattresses**"

Just as before, it would be virtually impossible to explicitly program into the machine all the various types and degrees of associations necessary to answer these questions like a human.

Other variations on the Rating Game could be invented that would have the same effect. We could, for example, have a Poetic Beauty Rating Game where we would ask for ratings of beauty of various lines of poetry.⁹ For a computer to do as well as a human on this test, it would either have to have experienced our life and language *as we had* or contain a theory of poetic beauty that included necessary and sufficient conditions for what constituted a beautiful line of poetry. Few would seriously argue that such an experience-independent theory was possible.

Or a Joke Rating Game: "On a scale of 0 to 10 rate how funny you find each of the following jokes" followed by a list of jokes. Again, capturing the necessary and sufficient conditions for humor would seem to require a grounding in all of human experience. Most jokes depend on a vast network of associative world knowledge ranging from the most ridiculous trivia, through common but little-commented-upon aspects of human experience, to the most significant information about current events. So here again is an example of where a computer, in order to appreciate humor as we did and thereby fool the Turing Test interrogator, would almost certainly have had to experience life and language as we had.

A final variation: The Advertising Rating Game. "Given the following product **X**, rate the following advertising slogan **Y** for that product." Once again, it is hard to imagine any theory that could provide necessary and sufficient conditions for catchy advertising slogans. Good advertising slogans, like good jokes and good lines of poetry, are perceived as good because of the myriad subconscious pressures and associations gathered in a lifetime of experiencing the world.

The impossibility of isolating the physical level from the cognitive level

One of the tacit assumptions on which Turing's proposed test rests is that it is possible to isolate the "mere" (and thus unimportant to the essence of cognition) physical level from the (essential) cognitive level. This is the reason, for example, that the candidates communicate with the interrogator by teletype, that the interrogator is not permitted to see them, and so on. Subcognitive questions, however, will always allow the interrogator to "peek behind the screen". The Turing Test is really probing the associative concept (and sub-concept) networks of the two candidates. These networks are the product of a lifetime of interaction with the world which *necessarily involves* human sense organs, their location on the body, their sensitivity to various stimuli, etc. Consider, for example, a being that resembled us precisely in all physical respects except that its eyes were attached to its knees. This physical difference alone would engender

enormous differences in its associative concept network compared to our own. Bicycle riding, crawling on the floor, wearing various articles of clothing (e.g., long pants) and negotiating crowded hallways would all be experienced in a vastly different way by this individual.

The result would be an associative concept network that would be significantly – and detectably by the Turing Test – different from our own. Thus, while no one would claim that the physical location of eyes had anything essential to do with intelligence, a Turing Test could certainly distinguish this individual from a normal human being. The moral of the story is that the physical level is *not* dissociable from the cognitive level. When Dreyfus says that no one expects an intelligent robot to be able to "get across a busy street. It must only compete in the more objective and disembodied areas of human behavior, so as to be able to win at Turing's game"¹⁰, he, like Turing, is tacitly accepting that such a separation of the physical and the cognitive levels is indeed possible. This may have seemed to be the case at first glance but further examination shows that the two are inextricably intertwined.

Can the Turing Test be appropriately modified?

Any reasonable set of questions in a Turing Test will necessarily contain subcognitive questions in some form or another. Ask enough of these questions and the computer will become distinguishable from the human because its associative concept network would necessarily be unlike ours. And thus the computer would fail the Turing Test.

Is it possible to modify the rules of the Turing Test in such a way that subcognitive questions are forbidden? I think not. The answers to subcognitive questions emerge from a lifetime of experience with the minutiae of existence, ranging from functionally adaptive world-knowledge to useless trivia. The sum total of this experience with its extraordinarily complex inter-relations is what defines human intelligence and this is what Turing's imitation game tests for. What we would really like is a test for (or, lacking that, a theory of) intelligence *in general*. Surely, we would not want to limit a Turing Test to questions like, "What is the capital of France?" or "How many sides does a triangle have?". If we admit that intelligence in general must have *something* to do with categorization, analogy-making, and so on, we will of course want to ask questions that test these capacities. But these are the very questions that will allow us, unfailingly, to unmask the computer.

The relevance of subcognitive factors

There remains the question of the *relevance* of these subcognitive factors that, as I believe I have shown, make it essentially impossible for a machine that has not experienced the world as we have to pass the Turing Test. Are these factors irrelevant to intelligence – just as a seagull's dipping in flight is irrelevant to flying in general – or are they a necessary substrate of intelligence? An initial part of my response is that a *human* subcognitive substrate is definitely not necessary to intelligence in general. The Turing Test tests precisely for the presence of a human subcognitive substrate and this is why it is limited as a test for general intelligence.

On the other hand, I believe that *some* subcognitive substrate is necessary to intelligence. I will not present a detailed defense of this view in this paper for two reasons: first, such a defense is beyond the scope of this paper, the goal of which has only been to discuss the limits of the Turing Test as a tool for determining intelligence and second, the necessity of a subcognitive substrate for intelligence has been compellingly argued elsewhere.¹¹ Some ideas of the defense will, however, be briefly presented below.

There is little question that intelligence relies on an extraordinarily complex network of concepts with various degrees of overlap. Philosophers from Wittgenstein¹² to Lakoff¹³ have shown that the boundaries of concepts are extraordinarily elusive things to pin down. It is probably impossible, even in principle, to describe categories in an absolute, objective manner. "Apples", for example, are almost always members of the category "food", but what about "grass", or

"shoes?" If you haven't eaten for ten days, "shoes" might well fall into your category of "food." But could something like "the Spanish Inquisition" ever be considered "food?" (Of course. Consider the following statement by a professor about to give an extraordinarily long lecture on Medieval methods of torture: "The meat of the first three hours of this lecture will be medieval torture in general. And if none of you has fallen asleep by then, we'll have the Spanish Inquisition for dessert."¹⁴) This is not a point to be taken lightly, for the associative overlap of categories essential to intelligence (and creativity) frequently occurs near the blurry boundaries of categories. And, to repeat, these boundaries are virtually impossible to define in an objective, context-independent way. Most of our thought processes are intimately tied to the associative overlap of categories. One particular example is analogy-making. Considered by many to be a *sine qua non* of intelligent behavior, it relies heavily on the ability to see two apparently unrelated situations as members, however obliquely, of the same category.

If we can view categories as being composed of many tiny (subcognitive) parts that can overlap with the subcognitive parts of other categories, we can go a long way toward explaining these associative phenomena. If, on the other hand, we deny the relevance of subcognitive factors in intelligence, we are left with the daunting, perhaps impossible, task of explicitly defining *all* of the possible attributes of each particular category in every conceivable context. It is therefore reasonable to conclude that all intelligence must have a subcognitive substrate. In particular, this implies that an intelligent computer would have to possess such a substrate, though there is no reason to believe that this substrate would be identical to our own.

Conclusion

In conclusion, the imitation game proposed by Alan Turing provides a very powerful means of probing human-like cognition. But when the Test is actually used as a real test for intelligence, as certain philosophers propose, its very strength becomes a weakness. Turing invented the imitation game only as a novel way of looking at the question "Can machines think?". But it turns out to be so powerful that it is really asking: "Can machines think exactly like human beings?". As a real test for intelligence, the latter question is significantly less interesting than the former. The Turing Test provides a sufficient condition for human intelligence but does not address the more important issue of intelligence in general.

I have tried to show that only a computer that had acquired adult human intelligence by experiencing the world as we have could pass the Turing Test. In addition, I feel that any attempt to "fix" the Turing Test so that it could test for intelligence in general and not just human intelligence is doomed to failure because of the completely interwoven and inter-dependent nature of the human physical, subcognitive and cognitive levels. To gain insight into intelligence, we will be forced to consider it in the more elusive terms of the ability to categorize, to generalize, to make analogies, to learn, and so on. It is with respect to these abilities that the computer will always be unmasked if it has not experienced the world as a human being has.

Acknowledgments

I especially wish to thank Daniel Dennett and Douglas Hofstadter for their invaluable comments on the ideas and emphasis of this paper. I would also like to thank David Chalmers, Melanie Mitchell, David Moser, and the editor of *Mind* for their remarks.

Endnotes

¹Turing, Alan M. (1950), Computing machinery and intelligence, *Mind*, Vol. 59, No. 236, pp. 433-460.

² For a particularly clear defense of this view see: Dennett, D. C. (1985). "Can Machines Think?" *How We Know*, ed. Michael Shafto. San Francisco, CA: Harper & Row.

³ Dreyfus, Hubert L. (1979), *What Computers Can't Do*, New York, NY: Harper & Row, p. 73.

⁴ Three different approaches that all address subcognitive issues can be found in:

- Feldman, J. and F. Ballard (1982). Connectionist models and their properties. *Cognitive Science*, 6(3), 205-254;
- Hofstadter, D. R., M. Mitchell, and R. M. French (1987). Fluid concepts and creative analogies: A theory and its computer implementation. CSMIL Technical Report No. 10, University of Michigan;
- Rumelhart, D. and J. McClelland (1986) (Eds.). *Parallel Distributed Processing*. Cambridge, MA: Bradford/MIT Press

⁵ A particularly relevant, succinct discussion of associative priming can be found in: Anderson, J. R. (1983) *The Architecture of Cognition*. Cambridge, MA: Harvard University Press, Chap. 3, pp. 86-125. In this chapter Anderson makes reference to the classic work on facilitation by Meyer and Schvaneveldt (Meyer, D. E. and R. W. Schvaneveldt (1971). Facilitation in recognizing pairs of words: evidence of a dependence between retrieval operations. *Journal of Experimental Psychology* 90, 227-234).

⁶ Hofstadter, D. R. (1985). "On the seeming paradox of mechanizing creativity". In *Metamagical Themes* pp. 526-546. New York, NY: Basic Books, Inc.

⁷ Even though Turing did not impose a time constraint in his original formulation of the imitation game, he did claim that "...in fifty years' time [i.e., by the year 2000] it will be possible to programme computers . . . to make them play the imitation game so well that an average interrogator will not have more than 70 per cent. chance of making the right identification after five minutes of questioning" [p. 442]. In current discussions of the Turing Test, the duration of the questioning period is largely ignored. In my opinion, one reasonable extension of the Turing Test would include the length of the questioning period as one of its parameters. In keeping with the spirit of the original claim involving a five-minute questioning period, I have tried to keep the number of questions short although it was by no means necessary to have done so.

⁸ This variation of the Rating Game was suggested to me by Douglas Hofstadter.

⁹ In fact, the interrogator in Turing's original article does indeed conduct a line of questioning about a particular turn of phrase in one of Shakespeare's sonnets.

¹⁰ Dreyfus, *op. cit.*, p. 78.

¹¹ Hofstadter, *op. cit.* "Waking Up from the Boolean Dream or, Subcognition as Computation" pp. 631-665.

¹² Wittgenstein, Ludwig. (1958) *Philosophical Investigations*. New York, NY: Macmillian Publishing Co.

¹³ Lakoff, George. (1987) *Women, Fire and Dangerous Things*. Chicago, IL: The University of Chicago Press.

¹⁴ This example is due to Peter Suber.