

# Updating the Turing Test. Wittgenstein, Turing and Symbol Manipulation\*

Carlo Penco

University of Genoa, Genoa, Italy  
Email: [penco@unige.it](mailto:penco@unige.it)

Received April 12<sup>th</sup>, 2012; revised May 14<sup>th</sup>, 2012; accepted May 30<sup>th</sup>, 2012

In this paper I present an argument against the feasibility of the Imitation Game as a test for thinking or language understanding. The argument is different from the five objections presented by Turing in his original paper, although it tries to maintain his original intention. I therefore call it “the Sixth Argument” or “the Argument from Context”. I show that—although the argument works against the original version of the imitation game—it may suggest a new version of the Turing Test, still coherent with the idea of thinking and understanding as symbol manipulation. In a new form, the main idea which lies behind the original Imitation Game remains untouched by the criticism of Searle’s Chinese room argument and suggest a possible implementation which avoids some of the shortcomings of the original Turing Test.

*Keywords:* Philosophy of Logic; Meaning; Context; Imitation Game; Turing; Wittgenstein

## Introduction

Wittgenstein’s *Philosophical Investigations* was published just a few years after his pupil’s famous paper on *Mind*, “Computing Machinery and Intelligence” (1950). In this paper Turing, who had attended Wittgenstein’s lectures in 1939, was relying on a vision of language as an essential feature of intelligence. He was following the intuition of his teacher, according to whom “thinking is essentially operating with signs”<sup>1</sup>. Wittgenstein’s main novelty in his book posthumously published in 1953 was a vision of language which partly challenged the Turing Test. I don’t mean that Wittgenstein was against the idea that machines can think; he made a few remarks on the topic, such as the following:

Could a machine think?—Could it be in pain?—Well, is the human body to be called such a machine? It surely comes as close as possible to being such a machine. But a machine surely cannot think!—Is that an empirical statement? No. We only say of a human being and what is like one that it thinks. We also say it of dolls and no doubt of spirits too. Look at the word “to think” as a tool (Wittgenstein, 1953: pp. 359-360).

In this quotation we have a very general remark on the use of the verb “to think”: if something is sufficiently “like” a human being it is reasonable to attribute to it the property of thinking. And being sufficiently “like” a human being implies—among other things—showing correct linguistic behavior. As Harnad 2000 (p. 429) remarks, “‘likeness’ can take two forms: likeness in structure and likeness in function”; Turing chooses likeness in function, particularly concerning the function tested by linguistic behavior. Harnad criticizes the limitation of the original Turing Test using only linguistic behavior, and proposes a hi-

\*A previous version of this paper was presented at the IIT (Italian Institute of Technology); I thank Giulio Sandini and all the participants for their criticism and their patience in front of a philosophical discussion a bit far from the direct commitments of the researchers. Thanks also to Marcello Frixione for comments on a previous version of the paper.

<sup>1</sup>A well known recurrent motto used in Wittgenstein 1958.

erarchy of Turing tests, where the fundamental one is the 3<sup>rd</sup> level of test, where robotics is central; this step is motivated by the fact that “things that human beings can do go beyond mere verbalising” and only with robotics we may map these intelligent abilities. The idea had already been presented in Harnad 1991 with the TTT (Total Turing Test), which is exactly like the Turing test, but requires machines to answer all kinds of input, not just verbal<sup>2</sup>. My proposal is slightly different, although it shares with Harnad the idea that the meaning of words cannot develop in isolation from action (“it is hard to imagine how our words could have the meanings they have if they were not first grounded in these nonverbal interactions with the world”)<sup>3</sup>. The central idea of the present paper is connected with a better and more adequate definition of what is meant by “linguistic behavior”. With his conception of “language game” Wittgenstein insisted that no linguistic expression has any meaning unless considered inside a context of actions and goals. Certainly the imitation game is a “language game” with clear actions and goals, a game to which Wittgenstein could have given his approval and interest.

In what follows I want to make two opposite remarks on the Turing Test from the point of view of Wittgenstein’s extreme contextualism. On the one hand Wittgenstein helps us in clarifying the limits of the imitation game with respect to what we may call the core of our language use, the connection between linguistic expressions and actions. On the other hand, some of Wittgenstein’s ideas may respond to the concerns raised by John Searle’s famous mental experiment of the Chinese room, and offer a possible alternative to a renewed Turing Test to help us better understand what we could be meant by saying that machines could think.

<sup>2</sup>This idea has provoked the search for stronger and stronger kinds of Turing Tests, like the proposal by Schweizer, 1998 for a Truly Total Turing Test, requiring machines able to develop languages and new ideas. Contemporary robotics is not so far from this conception. But Turing himself suggested putting sense organs in a machine and teaching it to learn a language. A completely opposite view is the view of “restricted” Turing Tests, devoted to single topics or abilities (see for instance Rajaraman, 1997).

<sup>3</sup>Harnad, 2000: p. 429.

## The Sixth Argument

In Loebner Prize, a competition based on the Turing Test, the prize is given to the computer whose responses are indistinguishable from a human's. After 10 years of competition for the Loebner Prize, the Turing test has never been completely overcome, and no programs have had much more success than the original Eliza by Weizenbaum<sup>4</sup>. Although there are still many defences of the Turing test (e.g. Copeland, 2000; Moor, 2001), and Turing himself suggested different variations<sup>5</sup>, according to Luger, 2005, the main contemporary reactions to the Turing Test are critical: 1) it deals with purely symbolic problem solving without any connection with perceptual skills or actions in real world, given that the dialog is conducted by means of a keyboard in a separate room; 2) it needlessly constrains machine intelligence to fit the mould of human intelligence. Partly following this second attitude, Hayes & Ford, 1995 claim that the test is a "distraction" from real artificial intelligence research<sup>6</sup>. French 2000 suggested that the real problem is not how to pass the Turing Test, but "why we can't pass it". Both reactions are reasonable, but the Turing Test still has some appeal, and I want to explore reaction a) from a particular point of view: let us keep the Test as dealing with "symbol manipulation"; even working inside the concept of symbol manipulation we are forced to challenge the original setting of the test. And, as we will see, a more difficult test could prove to be a more tractable one.

When Turing devised the original setting for the test he anticipated many possible objections: 1) the theological objection; 2) the "head in the sand" objection; 3) the mathematical objection; 4) the argument from consciousness 5) the argument from various disabilities. If we stick to the original setting and the original idea that thinking is symbol manipulation, can we be content with the answer Turing gave to these objections? I don't think so, and I think that Wittgenstein's insistence on the context dependence of meaning constitutes a sixth objection which is apparently very difficult to overcome in principle with the imitation game: I will call it "the Context Argument".

The Context argument might be stated in a very simple definition, which the standard Turing test could not overcome:

CA: We cannot have language understanding unless there is a processing of the context of utterance and the cogni-

<sup>4</sup>The point is that Turing Test could have been considered "successful" even with Eliza, if a very "naive" person might believe Eliza is a real person; but this is based on a "trick" and should be rejected (on the Turing test not being a trick see also Harnad, 2000; on Eliza being too naive to enter the Loebner Prize, see Saygin, Cicekli, & Akman, 2000). The Loebner Prize judges are a bit more sophisticated (even if sometimes ignorant). But the basic chat boxes are still made of tricks, and unable even to fulfil the basic anaphoric links between two sentences. On the debates on the Turing Test see the collection edited by Moor, 2003, but it is also possible to find something interesting in documents linked in web pages like: <http://aaai.org/AITopics/TuringTest>.

<sup>5</sup>Turing himself on the possibility of having an "unrestricted version" of his test answered in a BBC broadcast: "Oh yes, at least 100 years, I should say" This shows that also Turing could see a future for his test (See Moor, 2001: p. 91).

<sup>6</sup>See also Hayes & Ford, 1997 reviewing a book of interviews on AI, with opposite attitudes towards Turing Test. Minsky, not interviewed in the book had an attitude not dissimilar from the one held by Hayes, and even offered 100 dollars to whoever could bring an end to the Loebner Prize. Given that the Loebner Prize is supposed to stop when somebody succeeds in passing the Turing Test, Loebner called his prize "Minsky-Loebner Prize", assuming that, if there is a winner—and the Prize stops—Minsky would enrich the average prize by 100 dollars.

tive context in which a sentence is used.

The *context of utterance*, as defined by Kaplan and Lewis, is given basically by speaker, time and location of the utterance, and requires a direct connection with action and perception in order to interpret the proper use of indexicals and demonstratives (the referential use of "I", "you", "now", "here", "today" and so on). No amount of given information about the world, no universal encyclopedia, like for instance CYC, can replace this basic *ability* in language understanding: an enormous knowledge base may connect a huge number of inferences about a great deal of general and particular information about the world, but cannot help in understanding "I have to meet her today" or "I left all the beers in the fridge yesterday". The *cognitive context* includes general rules of conversation such as those inspired by the work of Grice dealing with conversational implicatures, and with presuppositions from the "common ground" of a conversation as suggested by Stalnaker. Moreover, language understanding cannot rely only on semantic networks defining inferences running over the lexicon of a language: linguistic competence deals not only with inferential aspects of the lexicon, but also with referential aspects, the ability to tell a cat from a dog, or to recognize the person now in front of you as the same person you met yesterday. Proper language understanding requires the ability to *use* symbols in context, non only to manage the inferential relations among symbols (which is just *one* aspect of linguistic competence, as stressed by Marconi, 1997). Being organized in a closed room and devoted only to linguistic exchange and reasoning without any kind of interaction with the real world, the standard Turing test seems unable to face the challenge of the Context Argument.

In the next paragraph I want to exemplify some particular ways in which the topic of context dependence has been developed in philosophy and semantics. I claim that these developments may help us to suggest another form of the Turing test, which could overcome not only the criticism raised by the Context Argument, but also the radical challenge posed by Searle's Chinese Room argument (Searle, 1980).

## The Context Argument: Details

The problem of contextual dependence is not only a problem for the Turing test, but for semantics in general and for formal semantics in particular. Standard model theoretical semantics treats the meaning of a sentence as its truth conditions: understanding a sentence is knowing when the thought expressed by the sentence is true. This is the fundamental step of traditional philosophy of language, the step according to which we speak of the "truth-conditional content" of a sentence. I understand the meaning of " $2 + 2 = 4$ " if I know that the sentence is true only if 2 and 2 is 4; or, I know the meaning of "the Eiffel Tower is located in Paris" only if I know to what "Eiffel Tower", "Paris", "in" and "located" refer to, and I know that the sentence is true only if the Eiffel Tower is located in Paris. Since the beginning of the analytic tradition in philosophy of language it has been apparent that there are problems for this general project. Formal semantics works quite well with expressions where time and place are clearly expressed and we know the meanings and referents of all the expressions composing the sentence. Problems arise with indexicals and demonstrative expressions, that is with context dependent expressions like "I", "you", "he", "now", "today", "this", "that" and

so on. As Perry 1997 makes evident, these expressions do not represent a syntactic set (they have different syntactic roles, from pronouns to adverbs), but they are a semantic set, characterized by their dependence on the context of utterance, defined as a limited set of parameters: “speaker, time, location”. Apparently if I say:

(1) “On April, 11, 2012 Carlo is trying to finish a paper”

not only does everybody understand the sentence, but it is easy to get the truth conditions: the sentence is true only if on April, 11, 2012 Carlo is trying to finish a paper. But things become more problematic when I say:

(2) “Today I am trying to finish a paper”

Every English speaker understands the sentence, and the sentence uttered at the time and place of the utterance has clear truth conditions. But if I find sentence (2) written on paper, there is no way to give a semantic evaluation (to evaluate at what conditions it is true) because the sentence can be understood only in the context of its utterance. Certainly there is *some kind* of truth condition: the sentence is true if the person speaking on the day to which “today” refers was trying to finish a paper during that day. But this is only a schema of truth conditions, it is just a set of rules waiting to be applied (Perry would speak of “reflexive truth conditions”).

Therefore, while (1) has a determined meaning, represents a set of procedures whose application permit us to check the truth of the sentence, in (2)—unless we know the value of the parameters of the context of utterance (time, location and speaker)—there is no way to evaluate the truth of the matter.

The general rules attached to the indexicals (“today” refers to the day of the utterance; “I” refers to the speaker of the utterance...) are procedures which need to be activated and filled with the appropriate contents in order to evaluate the sentence. But we need knowledge of the context to fill the gap. Probably the traditional Turing test would have no problem with the term “I” which is interpreted as “the individual writing on the screen at the moment”. But problems arise when we deal with expressions like “here”, “there”, “he”, “she” or “this” and “that” (when used demonstratively and not anaphorically). The standard treatment of this kind of context dependent expressions—indexicals and demonstratives—has been given by David Kaplan. Kaplan 1989 distinguishes what he calls “character” from what he calls “content”. The character is the linguistic rule attached to the indexical, and the content is the referent of the indexical. In the last couple of decades there has been much discussion on the logical and semantical treatment of “indexicals”. Different classifications of indexicals have been given, relying on the presence/absence of intentions or gestures (according to which only “now” and “I” are “pure indexicals”, given that the other expressions seem to have a common, at least implicit, demonstrative or intentional aspect). Philosophers and logicians have devoted much time and effort to find a suitable way to treat such context dependent expressions semantically.

Now a problem arises: are these the only expressions whose referents (and meanings) depend on the context? Some authors (e.g. Cappelen & Lepore, 2004) assert that indexicals are the only strictly context dependent expressions, and all other expressions contribute a “minimal content” to the thought expressed by an utterance. For instance to give a semantic evaluation of “he is tall” we just need to apply the rule for “he” in the context of utterance and we assert that the property of being tall

applies to the person referred to as “he”. Indexicals are the “basic set” of context dependent expressions<sup>7</sup>. But not everybody agrees. Significant concerns have been expressed regarding scalar adjectives, definite descriptions and shortened expressions used in local settings. Let us consider these three cases.

In a very Wittgensteinian vein, many authors (e.g. Recanati, 2007, 2010) claim that the context of utterance intended in a broad sense (including also presuppositions and beliefs, that is including cognitive context) is constitutive of the meaning and reference of many kinds of expressions. Take “tall”: I can say, referring to the same person on different occasions,

(3) “John is tall”

because John is taller than the average height of his friends at school, and at the same time I may say

(4) “John is not tall”

when asked about the possibility for him to become a basketball player. If the meaning or the truth conditions of “John is tall” do not depend on the context, then I would reach a contradiction. The context of utterance in a narrow sense (time, speaker and location) is not enough to solve the problem; to solve the problem we need a wider conception of context, inclusive not only of speaker, location and time, but also of shared presuppositions, implicatures and other kinds of assumptions.

Another much discussed case is that of referential uses of definite descriptions like “the x that has the property F”. We often use descriptions loosely because we want to refer to some salient element of the scene, and the context helps us to pick the right referent; think of

(5) “The book is on the table”

Which book? Which table? This is a case of *incomplete* description, where the context fills the gaps and helps to understand what I am speaking about. Things get worse in case of misdescriptions or inaccurate descriptions: at a party I may say

(6) “The man drinking a martini is a philosopher”

Let us assume that the man is not drinking a martini, but sparkling water. Strictly speaking I am referring to the only person in the room drinking a martini, therefore I say something false and by implicature I make myself understood (this is briefly the classical stance held by Kripke 1975, and generally shared by many contemporary authors). But certainly most of you will understand whom I am referring to, without making a step from the literal falsity to the implicature (see for instance

<sup>7</sup>A further case is the case of quantifiers; we use quantifiers in everyday language with implicit restriction of the domain of interpretation.

1) Every bottle of beer is in the fridge

Apparently we do not mean that every bottle of beer in the world is in the fridge; no fridge could be so big. We are just referring to a specific domain, maybe the bottles we have bought for the party. But also something which seems even simpler can be ambiguous.

2) Everyone is at home

An utterance of 2) may have different interpretations: “everyone belonging to a defined set of individuals is at the home of this set” or “everyone belonging to a defined set of individuals is at the home of N (where N is the relevant person in the context)”. In cases like these we need to find a way to formalize something which seems implicit, and in a logical language representation we need to postulate bound variables in the structure of quantified noun phrases, whose values, relative to a context, generate a domain of quantification (see Stanley, 2005). However, restriction to a domain of quantification is something which could be treated formally with less difficulties than the other cases under discussion.

Penco, 2010; Korta & Perry, 2011).

A third kind of examples concerns what is normally called “synecdoche” (the part for the whole). Take for instance:

(7) “The ham roll ran away without paying”

said by the waiter to the owner of the bar: it is apparent that the waiter is referring to the customer who ordered a ham roll, but literally speaking the matter is different. How can we give a treatment of the truth conditions of these and other kinds of examples without referring to the context? And which specific rules have to be applied in order to get the right truth conditions? In the case of indexicals we have standard, systematic rules for applying the parameters and get the content; is it possible to find some standard rules for other kinds of context dependent expressions in order to get their referent and build a correct semantic representation of the sentence in context?

Summing up, we have to accept that in basic cases of linguistic interaction, in order to understand the meaning of what is said, we need perceptual awareness of the context of utterance, together with the shared presuppositions of cognitive context, and language alone is not enough. It is the basic problem raised many years ago by Carnap, who said that “pragmatics is the basis for all linguistics”. This does not mean that syntax and semantics have to deal with proper language use in action, but means that pragmatics is what fills the content of semantical and syntactical features. The inferential working of the lexicon begins after pragmatic disambiguation.

I have given some space to this cursory presentation of what has been the average discussion on context dependence in the last twenty years in order to provide some examples of the challenges of the Context Argument. As we have said before, we need to take account not only of the “narrow context”, that is the context of utterance, as for (1) and (2), but also of the “broad context”, that is the intended domain of interpretation, the intention of the speaker, the setting of the scene, the presuppositions which emerge from the discourse, and many other features of the situation in which a sentence is uttered. Is anything like the imitation game a good guide in dealing with *these* problems? Not really. In fact contemporary natural language understanding systems work fairly well for simple automatic translation, but do not work properly in understanding language-in-context. Natural language understanding systems (like humans) need the capability of understanding elements of context corresponding to indexicals, demonstratives and other “contextuals”, like quantifiers, definite descriptions, adjectives and local abbreviations. These abilities are the ground which permits semantics (the truth evaluation of sentences) to work, and—as far as natural language processing is concerned—formalizing these abilities is one of the main challenges of our century.

## Updating the Turing Test

The attempts to overcome the limitations of standard formal semantics are promising, especially when connected with research programs in computational linguistics and artificial intelligence (the multi context theory developed by John McCarthy has been a first fundamental step in this direction)<sup>8</sup>. If, on the one hand, they constitute a challenge to the original Turing Test, they might also offer some hope of a new reconstruction of the imitation game, as inspired by the fundamental tenets shared by

<sup>8</sup>See for instance McCarthy, 1990; Guha & McCarthy, 2003; Penco & Vignolo, 2005.

Turing and Wittgenstein: thinking or understanding is the capability of using signs or, in other words, “Language understanding is symbol manipulation”.

Unfortunately the depth of this idea has been obscured by the interpretation of the test given by John Searle in his mental experiment of the Chinese Room. Searle’s challenge to artificial intelligence was exactly a critique of the concept of “symbol manipulation”, considered literally as working with symbols detached from any real interaction with the environment. In Searle’s mental experiment an English speaker has some instructions in English to take Chinese symbols as input and is to give some other symbols as output; the rules in English permit the English speaker to produce as output reasonable answers to the questions in Chinese. A Chinese speaker therefore would understand the answers to her questions produced by this procedure, thinking that whoever is inside the room understands Chinese. It is apparently a rhetorical presentation of the Turing test with the aim of depriving the test of its significance. In fact, Searle asks whether we can say that the man in the room understands Chinese. Certainly not! He understands English, and is able to use rules (formal or not) to give as output some patterns of Chinese symbols as answers to other patterns of Chinese symbols taken as input, without having any idea of what those symbols may mean. Symbol manipulation is not understanding language! What is missing is the understanding on the meaning of the Chinese symbols and the intentionality, that is the ability to understand what a symbol *refers to*. The English speaker inside the room has no idea what the Chinese symbols refer to; he only knows how to manipulate symbols, he is only using a syntactic ability without semantics.

However, the Chinese room is based on the traditional setting of the Turing test: somebody writes a sentence in Chinese and the Chinese room answers. One of the first reactions to the mental experiment was that what answers is not the man inside the room, who apparently knows only English, but the entire system: what is endowed with an understanding of Chinese is not the man inside the room—the linguistic program—but the entire system composed of the man, the room and the Chinese symbols. Another answer is to give the Chinese room some possibility of movement and perception: the entire system becomes a robot, where the man inside with his rules is just a small part, the syntactic manipulation; if the output symbols are correct answers to the questions, this means that the system can interact coherently with the environment. Searle’s argument seems insufficiently robust to answer the argument of the system and that of the robot combined<sup>9</sup>.

Thinking as symbol manipulations is not intended to be restricted to actions “on” symbols alone (as was supposed in the

<sup>9</sup>These are two of the main objections to Searle’s argument. They may be found in the original collection in *Brain and Behavioral Sciences*; we have also a nice summary of the different reactions in the *Stanford Encyclopedia* (Cole, 2009).

<sup>10</sup>The argument is as follows: let us not fall prey to the fallacy according to which in a symbolic language we can leave the world of symbols for a non symbolic world which could give them meaning. If we have a simulation we are in the symbolic world. The problem of the conception of perceptual primitives is a problem of robotics, which does not contribute in any manner to what is intended as “meaning”, intended as a procedure for manipulating symbols relative to a fixed domain. Providing sensors to a “thinking machine” does not enrich the procedural aspect of meaning. Working inside a simulated world is separated from working in robotics, and projects in robotics cannot properly give more information about meaning than already given in natural language processing. Unless... meaning cannot be given a purely inferential definition and procedures have to deal with real world situations.

eighties; e.g. Wilks, 1986)<sup>10</sup>, but to actions “with” symbols: symbol manipulation is the ability to use symbols in context, like using and understanding indexicals, demonstratives and definite descriptions, shortened or not, on the ground not only of shared presuppositions, but also of perceptual abilities of recognizing individual objects and patterns. On the contrary, in the original setting of the imitation game, the human is conceived as an inference machine, and a program simulating a human is thought to have a similar mastery of inferences, such that a dialogue is possible through a computer screen and a keyboard. This setting of the Imitation Game permits mental experiments such as Searle’s where a system is supposed to receive sentences as input and give sentences as output. But if we want to implement the more basic features of language understanding (such as demonstratives and referential uses of definite descriptions, quantifiers...) we need to rely on a dialogue with a shared environment. We need therefore to have a form of interaction with real situation, mastering the use of symbols to detect items in the environment. The challenge is to give machine the ability to use perceptual information from the context, and to mix it with background knowledge in order to use the most difficult aspects of context-dependent language use presented in the previous paragraph. The formal work done on these new aspects of the boundary between semantics and pragmatics is some of the most promising as regards this possibility: updating the Turing test in real situations.

Actually, unnoticed by many, the first example of a possible query of the interrogator to the unknown (man, woman or computer) interlocutor in the original paper by Turing (1950) is: “C: Will X please tell me the length of his or her hair?”

The question is typically considered as an application of “his” or “her” anaphorically towards a previously restricted domain (the supposed interlocutors). In this case the role of indexicals is anaphoric, that is it picks the individual referred to in the conversation (in this case X). But what will happen if the question considers a referential use of “his” or “her”? In this case, we need the ability to detect in the environment something which might match the use of the indexical adjective, probably referring to a human salient in the context. We have here the beginning of a different kind of imitation game: we might for instance test a group of experts in front of a robot, whose behavior could be either autonomous or directed by a human at distance. The challenge would be to detect whether the robot is an autonomous one or something human directed. The ability to manipulate symbols here is the ability to interact with symbols in a common environment. If the autonomous robot is able to interact correctly, why not accept that, being sufficiently “like” a human being in manipulating symbols in context, it thinks?

In a very subtle analysis of the limitation of the Turing Test and of other possible challenges to artificial intelligence, Cohen 2006 praises other kinds of tests, such as robot soccer competition, because of their feasibility and capacity of development in more and more complex stages: “Turing’s test requires simultaneous achievement of many cognitive functions and doesn’t offer partial credit to subsets of these functions. In contrast, robot soccer presents a graduated series of challenges: it gets harder each year but is never out of reach”. However a Turing Test based on basic linguistic abilities in an actual situation might ask for very simple symbol manipulation in understanding basic actions like “pass me that ball” or “take the red can you see near him”. Before arriving at a highly complex dia-

logue on Shakespeare (as happened at the Loebner Prize where judges were a bit more ignorant than machines on the specific subject matter), a challenge on basic linguistic abilities in context might be more easily transformed into an updated Turing Test.

Summarizing: Searle, against Turing, suggested the idea that *understanding is not symbol manipulation*. However, properly understood, a new Turing test might be grounded on the idea that thinking and *language understanding is symbol manipulation in context*. Will we be able to invent an imitation game which could constitute a new challenge for the present century? An updated Imitation Game would consist of an interrogator trying to understand whether she is interacting with a robot or with a human in changing, real world situations. Given that so many humans have stereotypical behaviors, it would be easy to raise the doubt that one was meeting some kind of automatic agent, and the test would be a true challenge for humans and robots.

## REFERENCES

- Cappelen, H., & Lepore, E. (2004). *Insensitive semantics. A defence of semantic minimalism and speech act pluralism*. Oxford: Blackwell.
- Cohen, P. (2006). If not Turing’s test, then what? *AI Magazine*, 26, 61-67.
- Cole, D. (2009). The Chinese room argument. *The Stanford Encyclopedia of Philosophy (Winter 2009 Edition)*.
- Copeland, B. J. (2001). The Turing test. *Mind and Machines*, 10, 519-539. doi:10.1023/A:1011285919106
- French, R. (2000). The Turing test: The first fifty years. *Trends in Cognitive Sciences*, 4, 115-121. doi:10.1016/S1364-6613(00)01453-4
- Harnad, S. (1991). Other bodies, other minds: A machine incarnation of an old philosophical problem. *Mind and Machines*, 1, 43-54.
- Harnad, S. (2000). Mind, machines and Turing. The indistinguishability of the indistinguishable. *Mind and Machines*, 9, 425-445.
- Hayes, P., & Ford, K. (1995). Turing test considered harmful. *Proceedings of the 14th International Joint Conference on Artificial Intelligence*. San Francisco: Morgan Kaufman Publishers.
- Hayes, P., & Ford, K. (1997). Talking heads—A review of speaking minds: Interviews with twenty eminent cognitive scientists. *AI Magazine*, 18, 123-125.
- Korta, K., & Perry, J. (2011). *Critical pragmatics*. Cambridge: Cambridge University Press.
- Guha, R., & McCarthy, J. (2003). Varieties of contexts. *Lecture Notes in Artificial Intelligence*, 2116, 290-303.
- McCarthy, J. (1990). *Formalizing commonsense*. New York: Ablex.
- Moor, J. H. (2001). The status and future of the Turing test. *Mind and Machines*, 11, 73-93. doi:10.1023/A:1011218925467
- Moor, J. H. (2003) *The Turing test. The elusive standard of artificial intelligence*. Dordrecht: Kluwer.
- Luger, G. F. (2005). *Artificial intelligence: Structures and strategies for complex problem solving* (5th ed.). Boston: Addison-Wesley.
- Penco, C. (2010). Essentially incomplete descriptions. *European Journal for Analytic Philosophy*, 6, 47-66.
- Penco, C., & Vignolo, M. (2005). Converging towards what? Pragmatic and semantic competence. In P. Bouquet, & L. Serafini (Eds.), *Context representation and reasoning* (Vol. 136). CEUR-WS. <http://sunsite.informatik.rwth-aachen.de/Publications/CEUR-WS/Vol-136>
- Perry, J. (1997). Indexicals and demonstratives. In R. Hale, & C. Wright (Eds.), *Companion to the philosophy of language* (pp. 586-612). Oxford: Blackwell.
- Rajaraman, V. (1997). Turing test and after. *Resonance*, 2, 50-59. doi:10.1007/BF02835001
- Recanati, F. (2007). *Perspectival thought*. Oxford: Oxford University Press. doi:10.1093/acprof:oso/9780199230532.001.0001
- Recanati, F. (2010). *Truth conditional pragmatics*. Oxford: Clarendon

- Press. [doi:10.1093/acprof:oso/9780199226993.001.0001](https://doi.org/10.1093/acprof:oso/9780199226993.001.0001)
- Saygin, A. P., Cicekli, I., & Akman, V. (2000). Turing test: 50 years later. *Mind and Machines*, 10, 463-518. [doi:10.1023/A:1011288000451](https://doi.org/10.1023/A:1011288000451)
- Searle, J. (1980). Minds, brains and programs. *Behavioral and Brain Sciences*, 3, 417-457. [doi:10.1017/S0140525X00005756](https://doi.org/10.1017/S0140525X00005756)
- Schweizer, P. (1998). The truly total Turing test. *Mind and Machines*, 8, 263-272. [doi:10.1023/A:1008229619541](https://doi.org/10.1023/A:1008229619541)
- Stanley, J. (2005). *Language in context*. Oxford: Oxford University Press.
- Turing, A. M. (1950). Computing machinery and intelligence. *Mind*, 1950.
- Turing, A. M. (1951). Can digital computers think? BBC 3rd programme. In B. J. Copeland (Ed.), *The essential Turing*. Oxford: Oxford University Press, 2004. <http://www.turingarchive.org/browse.php/B/5>
- Wilks, Y. (1986). Default reasoning and self-knowledge. *Proceedings of the IEEE*, 74, 1399-1404. [doi:10.1109/PROC.1986.13640](https://doi.org/10.1109/PROC.1986.13640)
- Wittgenstein, L. (1953). *The philosophical investigations*. Oxford: Blackwell.
- Wittgenstein, L. (1958). *The blue and brown books*. Oxford: Blackwell.