

# Chatbots' Greetings to Human-Computer Communication

Maria João Pereira<sup>1</sup>, Luísa Coheur<sup>1,2</sup>, Pedro Fialho<sup>1</sup>, Ricardo Ribeiro<sup>1,3</sup>

<sup>1</sup>INESC-ID Lisboa, <sup>2</sup>IST, Universidade de Lisboa, <sup>3</sup>Instituto Universitário de Lisboa (ISCTE-IUL)  
Rua Alvez Redol, 9, 1000-029 Lisboa, Portugal

mjoao.madeiras@gmail.com, luisa.coheur@inesc-id.pt, peter.fialho@gmail.com, ricardo.ribeiro@inesc-id.pt

## Abstract

In the last years, chatbots have gained new attention, due to the interest showed by widely known personalities and companies. The concept is broad, and, in this paper we target the work developed by the (old) community that is typically associated with chatbot's competitions. In our opinion, they contribute with very interesting know-how, but specially with large-scale corpora, gathered by interactions with real people, an invaluable resource considering the renewed interest in Deep Nets.

**Keywords:** natural language interfaces, agent-based interaction, intelligent agents, interaction design

## 1. Introduction

Chatbots are currently a hot-topic, both for industry and academia (Dale, 2016; Følstad and Brandtzæg, 2017). There are many platforms to help developing such systems, and the number of new chatbots continues to increase at a dizzying pace. Pandorabots hosting service<sup>1</sup> declares to have more than 225,000 botmasters (people in charge of creating/maintaining a chatbot), which have built more than 280,000 chatbots, resulting in more than 3 billion interactions (numbers collected in July 2017). On the academia side, since 2016, at least four workshops were dedicated to chatbots (defined as non goal-oriented dialogue systems), which have been co-located with well-established conferences<sup>2</sup>; also, several works point how chatbots could be used to in learning environments (e.g., (Bibauw et al., 2019) and (Fialho et al., 2013)).

Although the current definition of chatbot is broader than the one we use in this paper<sup>3</sup>, we will use the word “chatbot” to name the old school chatbots, typically associated with chatbot's competitions.

We focus on chatbots that freely engage conversation about any subject (the non goal-oriented feature), making them “entertaining in a large variety of conversational topic settings” (Schumaker et al., 2007). However, these are also systems that “seek to mimic conversation rather than understand it”, that is, there is no real intention of making them “intelligent”, as the main goal of their developers is to make these chatbots effective in their simulation of intelligence. Some of these chatbots were developed and tailored with the goal of participating in chatbot's competitions (in fact, the term *chatbot* was coined in (Mauldin, 1994) to name the systems that have the goal of passing the Turing Test (Turing, 1950)), and, due to that, some have gained visibility. The lack of full descriptions and papers about these chatbots (which explains the abnormal number

of references to web pages in this paper) makes it difficult to uncover the technology and the real possibilities behind them. In this paper, we unveil the main contributions of this community, as we believe that this line of work can bring important insights to the human-machine communication field, as some of them contribute with large amounts of data gathered during their interactions with the crowd, which could be used by current data-driven chatbots (e.g., (Li et al., 2016; Vinyals and Le, 2015)). As we will see, these chatbots range from “simpler” ones, based on pre-written pattern-matching templates, exploiting large stores of prepared small talk responses, to more complex architectures, based on some sort of learning process. Finally, we will see that concepts/tricks introduced by some chatbots often result in a more solid contribution to the “illusion of intelligence” than the involved models.<sup>4</sup> This document is organised as follows: in Section 2. we present a brief historical overview, in Section 3. we discuss chatbot's platforms and how to enrich them, and, in Section 4., we summarise the main “tricks” towards the “illusion of intelligence”. Finally, in Section 5., we present some conclusions and point to future challenges.

## 2. Historical overview

In this section we make a brief review of these chatbots' history, moving from the first chatbots to the ones with which we interact nowadays.

### 2.1. Early days

Although the term chatbot was not invented by that time, the first chatbot came to public in 1966 under the appearance of a Rogerian psychotherapist called *Eliza* (Weizenbaum, 1966). *Eliza* was a program developed by Joseph Weizenbaum that was able to establish a conversation with human beings, simulating it was one too. *Eliza*'s conversational model was based in the rephrasing of input sentences, when these matched a set of pre-defined

<sup>1</sup>[www.pandorabots.com](http://www.pandorabots.com)

<sup>2</sup>[workshop.colips.org/wochat/](http://workshop.colips.org/wochat/)

<sup>3</sup>Many terms are used as synonyms of *chatbot*, as for instance *dialogue system*, *avatar*, *intellectual agents*, and *virtual person*. A list of more than 160 terms used as synonyms of chatbot can be found in [www.chatbots.org/synonyms/](http://www.chatbots.org/synonyms/).

<sup>4</sup>An extended version of this paper can be found in <https://arxiv.org/abs/1609.06479>.

rules. For instance, consider the following rule<sup>5</sup> constituted by a regular expression (match) and an answer (answer):

```
match: * you are *
answer: What makes you think I am (2)?
```

Example 1.

In this rule, if the match part coincides with the input (\* is the wildcard and matches every sequence of words), the text associated with the answer part will be returned, being the variable (2) replaced by the sequence from the input captured by the second wildcard. The following dialogue (Example 2) illustrates an application of this rule. Notice that some internal processing needs to be done, so that the sequence captured by (2) *entitled to your opinion* is modified into *entitled to my opinion*.

```
user: You are entitled to your opinion.
Eliza: What makes you think I am entitled
to my opinion?
```

Example 2.

*Eliza* completely exceeded the expectations, given that many people, when interacting with it, believed they were talking with another human (this outcome is currently called the “*Eliza effect*”). Without having any intention of modelling the human cognitive process and despite its simplicity, *Eliza* showed how a program impersonating a specific professional role can cause a huge impression by the mere illusion of understanding. Weizenbaum was taken aback by some aspects of this success (Hutchens, 1997). What shocked him most was the fact that people actually believed that the program understood their problems<sup>6</sup>. Perceiving *Eliza* as a threat, Weizenbaum wrote “Computer Power and Human Reason” (Kuipers et al., 1976) with the aim of attacking the Artificial Intelligence (AI) field and educating uninformed persons about computers.

Nowadays, *Eliza* is still one of the most widely known programs in AI and is at the base of a great number of chatbots, including *Parry*, its “successor”. Following a very similar architecture to that of *Eliza*, *Parry* appeared in 1971 by the hands of Kenneth Colby, simulating a paranoid mental patient (Saygin et al., 2000). An interesting comparison between *Parry* and *Eliza* was made by Güzeldere and Franchi<sup>7</sup>: “*Parry*’s strategy is somewhat the reverse of *Eliza*’s”, as one simulates the doctor, distant and without personality traces, and the other a paranoid patient which states its anxieties. Differently from *Eliza*, *Parry* has knowledge of the conversation and it also some sort of “state of mind”. The combination of these two factors affects the output as it becomes a function not only of the input, but also of *Parry*’s beliefs, desires, and intentions. In (Mauldin, 1994) a few tricks to which *Parry* resorts are

summarised, namely: (1) admitting ignorance; (2) changing the conversation topic; and, (3) introducing small stories about the Mafia throughout the conversation. These three tricks are (respectively) illustrated in the following answers given by *Parry*:

```
Parry: I don't get you.
...
Parry: Let's talk about something else.
...
Parry: I know the mob controls the big
rackets.
```

Example 3.

After Colby gathered transcripts of interviews between psychiatrists, patients and his program, he presented the results to another group of psychiatrists. He asked this group if they could guess in what transcripts the interviewed was a human and in which ones it was a program. The psychiatrist could not do better than randomly guessing.

It is possible to conclude from these results that the emotional side can be easier to imitate than the intellectual one (Kuipers et al., 1976). However, one of the main criticisms *Parry* received was of not being more than an illusion, incapable of modelling a real person (Colby, 1974).

## 2.2. The chatbots’ competitions

Moving back to 1950, Alan Turing questioned “can machines think?” (Turing, 1950), and proposed a way of testing it: the imitation game (now known as the Turing Test). The original imitation game is played by a man, a woman and an interrogator whose objective is to guess the sex of the players. Turing proposed substituting one of the players by a machine and playing the same game. In this version, if the interrogator wrongly identifies who is the human it means that the machine “can think”.

Based on (their own interpretation of) the Turing Test, chatbots’ competitions keep appearing. *Chatterbox Challenge*<sup>8</sup>, or, more recently, the *Chatbot Battles*<sup>9</sup>, are examples of such competitions, although the most widely known is the *Loebner prize*<sup>10</sup>, where participants are challenged with a simplified version of the total Turing Test (Powers, 1998). This prize is due to Hugh Loebner, who offered a reward to the first person whose program could pass the proposed test. The first Loebner Prize Contest took place in 1991, at Boston’s Computer Museum (Epstein, 1992), and, since then, the competition has been held annually in the quest of finding the “thinking computer”.

As some chatbots, competing for the Loebner prize, are indeed capable of managing a conversation, keeping it consistent, at least for a while, every year the most human-like computer is distinguished with a prize. However, since the first edition of the Loebner prize, in 1991, until now, no one won it. Nevertheless, in another Turing Test organised in 2014 by the U.K.’s University of Reading, a chatbot simulating a 13-year-old boy, named *Eugene Goostman*, created

<sup>5</sup>Inspired from *Eliza*’s implementation in

[search.cpan.org/~jnolan/Chatbot-Eliza-1.04/](http://search.cpan.org/~jnolan/Chatbot-Eliza-1.04/)

<sup>6</sup>[www.alicebot.org/articles/wallace/eliza.html](http://www.alicebot.org/articles/wallace/eliza.html)

<sup>7</sup>[www.stanford.edu/group/SHR/4-2/text/dialogues.html](http://www.stanford.edu/group/SHR/4-2/text/dialogues.html)

<sup>8</sup>[web.archive.org/web/20150905221931/http://www.chatterboxchallenge.com/](http://web.archive.org/web/20150905221931/http://www.chatterboxchallenge.com/)

[www.chatbotbattles.com](http://www.chatbotbattles.com)

<sup>9</sup>[www.chatbotbattles.com](http://www.chatbotbattles.com)

<sup>10</sup>[www.loebner.net/Prizef/loebner-prize.html](http://www.loebner.net/Prizef/loebner-prize.html)

by Vladimir Veselov and his team, convinced 33% of the human judges that it was human.

This event brought to the spotlight the old question of AI and generated (again) much controversy. In fact, many people consider that there was a misunderstanding of Turing's intentions in the different implementations of the Turing test, as deep models of thinking were a presupposition underlying Turing's imitation game. Following this, even if a chatbot was good enough to deceive the jury, it would not pass the Turing Test in Turing's sense, as it does not have a cognition model behind it. Another important criticism is stressed by Levesque (Levesque, 2014). For this author, AI is the science that studies "intelligent behaviour in computational terms", and the ability to be evasive, although interesting, may not show real intelligence. A computer program should be able to demonstrate its intelligence without the need for being deceptive. In this sense, Levesque et al. (Levesque et al., 2012) further explore this idea by conceiving a reading comprehension test based on binary choice questions with specific properties that make them less prone to approaches based on deception. Apart from the numerous controversies regarding the Turing Test, the fact is that all these competitions strongly contributed to the main advances in the field, and the most popular chatbots are the ones that were/are present in these competitions.

### 3. Building chatbots

Behind each chatbot there is a development platform. These are typically based on a scripting language that allows the botmaster to handcraft its knowledge base, as well as an engine capable of mapping the user's utterances into the most appropriate answer.

#### 3.1. Scripting languages/platforms

An impressive collection of *Elizas* can be currently found in the web. For instance, Chatbot-Eliza<sup>11</sup> is an implementation in Perl that can be used to build other chatbots. Knowledge is coded as a set of rules that are triggered when matched against the user's input. Some of the available programs offer features such as a certain capability to memorise information, adding synonyms or ranking keywords. The most popular language to build chatbots is probably the "Artificial Intelligence Markup Language", widely known as AIML,<sup>12</sup> a derivative of XML that includes specific tags. As usual, knowledge is coded as a set of rules that will match the user input, associated with templates, the generators of the output. The large usage of AIML can be justified by the following facts: besides its detailed specification, its community allows anyone to obtain, for free, interpreters of AIML in almost all coding languages, from Java (program D) to C/C++ (program C) or even Lisp (program Z); the set of AIML files that constitute the contents of *A.l.i.c.e.*'s brain can also be freely obtained<sup>13</sup>. All the pandorabots are based on AIML, more specifically in AIML 2.0. This specific release is usually characterised as being very easy to

modify, develop and deploy. Therefore, anyone, even non-computer-experts, can make use of it (Wallace et al., 2007), as no prior knowledge about AIML is required.

ChatScript<sup>14</sup>, the scripting language and open-source engine, should also be addressed, as is at the basis of *Suzette* (2010 Loebner Prize winner), *Rosette* (2011 Loebner Prize winner), *Angela* (2nd in 2012 Loebner Prize), and the previously referred *Rose* (2014 Loebner Prize winner). It comes with useful features, including an ontology of nouns, verbs, adjectives and adverbs, and offers a scripting language (inspired by the Scone project, a knowledge-base system developed to support human-like common-sense reasoning and the understanding of human language (Fahlman, 2011)). According to Bruce Wilcox, its creator, ChatScript settles several AIML problems, such as not being reader friendly. In fact, as AIML is based on recursive self-modifying input, it is harder to debug and maintain. A detailed comparison between *ChatScript* and AIML capabilities was made available by Wilcox, as a motivation for the development of a new (his own) chatbot platform.<sup>15</sup>

#### 3.2. Building chatbots by chatting

Another approach to develop chatbots' knowledge sources, which avoids handcrafted rules, is based on chatting and learning from the resulting chats. Systems like the already mentioned *Jabberwacky* (and *Cleverbot*) learn by keeping never seen user interactions and posing them later to other users. The acquired answers are then considered suitable answers for these interactions. That is, they learn to talk by talking, by relying on what has been said before by users and mimicking them. The user's intelligence becomes "borrowed intelligence" as, instead of being wasted, it incorporates a loop: what is said is kept (along with the information of when it was said) and in the future that knowledge may be exposed to another user. The given replies are then saved as new responses that the system can give in the future. It is only possible to give a brief overview of *Jabberwacky*'s or *Cleverbot* learning mechanisms as their architecture is not available to the public. The only disclosed aspect is that the AI model is not one of the usually found in other systems, but a "layered set of heuristics that produce results through analyses of conversational context and positive feedback"<sup>16</sup>.

Another example of a chatbot that learns is Robby Garner's "Functional Response Emulation Device" (*Fred*), the ancestor of *Albert One*, the winner of 1998 and 1999 Loebner Prize. *Fred* was a computer program that learned from other people's conversations in order to make its own conversations (Caputo et al., 1997). *Fred* began with a library of basic responses, so that it could interact with users, and from then on, it learned new phrases with users willing to teach it<sup>17</sup>. Although such an (unsupervised) learning may lead to unexpected and undesirable results, with the Internet growth and the possibility of having many people talking

<sup>11</sup> [search.cpan.org/~jmolnar/Chatbot-Eliza-1.04/Chatbot/Eliza.pm](http://search.cpan.org/~jmolnar/Chatbot-Eliza-1.04/Chatbot/Eliza.pm)

<sup>12</sup> [www.alicebot.org/aiml.html](http://www.alicebot.org/aiml.html)

<sup>13</sup> [code.google.com/p/aiml-en-us-foundation-alice/downloads/list](http://code.google.com/p/aiml-en-us-foundation-alice/downloads/list)

<sup>14</sup> [sourceforge.net/projects/chatscript/](http://sourceforge.net/projects/chatscript/)

<sup>15</sup> This comparison can be found in [gamasutra.com/blogs/BruceWilcox/20120104/9179/](http://gamasutra.com/blogs/BruceWilcox/20120104/9179/).

<sup>16</sup> [www.icogno.com/a\\_very\\_personal\\_entertainment.html](http://www.icogno.com/a_very_personal_entertainment.html)

<sup>17</sup> [www.simonlaven.com/fred.htm](http://www.simonlaven.com/fred.htm)

with the chatbots, one may foresee that these will quickly evolve.

#### 4. The illusion of intelligence and/or the art of scripting

Creating chatbots goes beyond writing good programs and developing algorithms, as in order to create a chatbot, more than being a programmer, the botmaster must be an author. Juergen Pirner, creator of the 2003 Loebner prize winner *Jabberwock*<sup>18</sup>, emphasises the scripting process behind a chatbot, stating that in the presence of possible failures, the one at fault is not the engine but its author.

Since making a chatbot involves preparing it to the impossible mission of giving a plausible answer to all possible interactions, the botmasters usually take advantage of several tricks to simulate understanding and intelligence. For instance, Pirner describes basic techniques of scripted dialogs like “having a set of responses for each scripted dialog sequence” and “ending those same responses with a clue, a funny remark or a wordplay”. With *Eliza*, we learnt that including the user’s strings in its answers helps maintaining an illusion of understanding (Mauldin, 1994). Other approaches focus on trying to guess what the user might say or forcing him/her to say something expected.

##### 4.1. Giving the bot a personality

Whereas personality has been a subject of study among the agent’s community, deeply exploited in all its complexity, the concept is kept as simple as possible within chatbots. As we have seen, what is common is the association of an *a priori* “personality” to a chatbot, which can justify some answers that otherwise would be considered inappropriate. For instance, Rogerian mode of *Eliza* covers for its answers, as it leads to a conversation where the program never contradicts itself, never makes affirmations, and is free to know nothing or little about the real world without being suspicious. The same happens with *Parry*: being a paranoid mental patient its changes in subject or incongruous answers are considered satisfactory and hide its absence of understanding. The aforementioned *Eugene Goostman* also follows along these lines. Veselov explains his reasoning for such a character: “a 13 years old is not too old to know everything and not too young to know nothing”<sup>19</sup>.

Thomas Whalen, winner of 1994 Loebner prize, took this a step further with *Joe*, the janitor. Whalen’s decision was related to the fact that contrary to previous editions of Loebner competitions, where the conversation was restricted to a topic, in 1995 the judges could pose any question. Hence, Whalen decided that the best approach to deal with a non-topic situation, would be to present a system that “would not simply try to answer questions, but would try to incorporate a personality, a personal history, and a unique view of the world”<sup>20</sup>. And so *Joe* was born. *Joe* was a night-worker janitor in the verge of being fired. He was only “marginally literate”, and he did not read books, newspapers, or watch television. These premises by themselves

restricted the conversation by giving *Joe* a “fairly narrow worldview”. Another trick was to use *Joe*’s eminent dismissal to introduce some stories revolving around it, which would, at the same time, provide a way of directing the conversation, the topic of the next section.

##### 4.2. Directing a conversation

Personality can justify some appropriate answers, but the best way to deal with unexpected interactions is to avoid them. Thus, being able to direct the conversation is a trick used by many chatbots, including the simple forms used by *Eliza*, where the usage of questions incited the user participation and made him/her keep the conversation with little contribution from the program.

*Converse* (Batacharia et al., 1999), created by David Levy, was the 1997 winner of the Loebner competition and did extremely well by using the clever trick of controlling a conversation. Although directing a conversation by “talking a lot about a predefined topic” was already used (Saygin et al., 2000), *Converse*’s performance convinced a judge for the first five minutes that he was really human: after greeting the judge, *Catherine* (*Converse*’s character) asked the interrogator about something that had passed on the news the previous day and then kept talking about it, as can be seen in the transcripts<sup>21</sup>. David Levy’s won again the Loebner prize in 2009 with *Do-Much-More*<sup>22</sup>, but this time the system was more flexible in the range of topics and responses it covered.

##### 4.3. Paying attention to small talk

Small talk, or phatic communication (Malinowski, 1923), is another hot topic in chatbots. It can be seen as a “neutral, non-task-oriented conversation about safe topics, where no specific goals needs to be achieved” (Endrass et al., 2011). Small talk can be used for two main purposes (Schneider, 1988): establish a social relation by building rapport and avoiding (embarrassing) silence. As stated in (Bickmore and Cassell, 1999), chatbots have been making use of the small talk mechanism. For instance, Epstein, an American psychologist, professor, author, and journalist, went to an online dating service, and believed for several months that a chatbot, met in the dating service, was a “slim, attractive brunette” (Epstein, 2007). In brief, small talk is a constant in all chatbots programs, used in non-sequiturs or canned responses. It fosters the idea of understanding and eases cooperation, facilitating human-like interactions by gaining user’s trust and developing a social relationship (Bickmore and Cassell, 2000).

##### 4.4. Failing like a human

After introducing the imitation game, Turing presented an example (Example 4) of a possible conversation one could have with a machine (Turing, 1950). Observing this example, besides the delay in providing the response, we can easily see that the answer is wrong. As Wallace wrote<sup>23</sup>, “we tend to think of a computer’s replies ought to be fast, accurate, concise and above all truthful”. However, human

<sup>18</sup> [www.chatbots.org/developer/juergen\\_pirner/](http://www.chatbots.org/developer/juergen_pirner/)

<sup>19</sup> [www.huffingtonpost.com/2012/06/27/eugene-goostman-](http://www.huffingtonpost.com/2012/06/27/eugene-goostman-2012-turing-test-winner_n_1630412.html)

[2012-turing-test-winner\\_n\\_1630412.html](http://www.huffingtonpost.com/2012/06/27/eugene-goostman-2012-turing-test-winner_n_1630412.html)

<sup>20</sup> [hps.elte.hu/~gk/Loebner/story95.htm](http://hps.elte.hu/~gk/Loebner/story95.htm)

<sup>21</sup> [www.loebner.net/Prizef/converse.txt](http://www.loebner.net/Prizef/converse.txt)

<sup>22</sup> [www.worldsbestchatbot.com/](http://www.worldsbestchatbot.com/)

<sup>23</sup> [www.alicebot.org/anatomy.html](http://www.alicebot.org/anatomy.html)

communication is not like that, containing errors, misunderstandings, disfluencies, rephrases, etc.

```
Human: Add 34957 to 70764.
(after pause of about 30 seconds)
Machine: 105621.
```

Example 4.

This is something that earlier chatbot’s writers already had in mind, as some already cared about simulated typing. For instance, *Julia* (Mauldin, 1994) simulated human typing by including delays and leaving some errors. Simulated typing also proves to be useful in decreasing mistakes by slowing down the interaction: Philip Maymin, a Loebner contestant in 1995, slowed the typing speed of his program to the point that a judge was not able to pose more than one or two questions (Hutchens, 1997).

## 5. Conclusions and future challenges

The number of chatbots that can be found in the web increases every day. Besides tools and corpora, the chatbots’ community has important know-how, which should not be neglected by researchers targeting advances in human-machine communication. Therefore, we present a brief historical overview of chatbots and describe main resources and ideas. Furthermore, we highlight some chatbots, relevant because they introduce new paradigms and/or won the Loebner prize. However, it should be clear that these are only the tip of the iceberg of the panoply of chatbots that currently exist.

We have seen that AIML and, more recently, *Chatscript* are widely used languages that allow to code the chatbots’ knowledge sources, and that even in chatbots that implement learning strategies, scripting is still at their core. We have also seen that a personality capable of justifying some of the chatbot’s answers, the capacity of directing a conversation and producing small talk, and the idea of failing like a human are some of the chatbots’ features that give the illusion of intelligence. We have also grasped that to create a chatbot, one “only” needs to think about a character and enrich its knowledge bases with possible interactions. Even better, that work does not need to be done from scratch as many platforms already provide pre-defined interactions, which can be adapted according to the chatbot character. And this is the main richness of the chatbot’s community: the immense amount of collected interactions, where the majority of them represent real human requests. All this data (after some validation) could be used to train current end-to-end data-driven systems. A major future challenge is to be able to automatically use all this information to build a credible chatbot. How to avoid contradictory answers? How to choose appropriated answers considering a chatbot’s character? And if we move to other sources of dialogues, like the ones from books, theatre plays or movies subtitles, will we be able, one day, to integrate all that information simulating real human dialogues?

## 6. Bibliographical References

Batacharia, B., Levy, D., Catizone, R., Krotov, A., and Wilks, Y. (1999). *Converse: a conversational companion*. In *Machine Conversations*, volume 511 of *The*

- Springer International Series in Engineering and Computer Science*, pages 205–215. Springer.
- Bibauw, S., François, T., and Desmet, P. (2019). Discussing with a computer to practice a foreign language: research synthesis and conceptual framework of dialogue-based call. *Computer Assisted Language Learning*, 0(0):1–51.
- Bickmore, T. and Cassell, J. (1999). Small talk and conversational storytelling in embodied conversational interface agents. In *Proc. of the AAAI 1999 Fall Symposium on Narrative Intelligence*, pages 87–92. AAAI Press.
- Bickmore, T. and Cassell, J. (2000). How about this Weather? Social Dialogue with Embodied Conversational Agents. In *Socially Intelligent Agents: The Human in the Loop*, pages 4–8. AAAI Press.
- Caputo, L., Garner, R., and Nathan, P. X. (1997). FRED, Milton and Barry: the evolution of intelligent agents for the Web. In F. C. Morabito, editor, *Advances in Intelligent Systems*, pages 400–407. IOS Press.
- Colby, K. M. (1974). Ten criticisms of PARRY. *SIGART Newsletter*, pages 5–9.
- Dale, R. (2016). The return of the chatbots. *Natural Language Engineering*, 22(5):811–817.
- Endrass, B., Rehm, M., and André, E. (2011). Planning Small Talk behavior with cultural influences for multiagent systems. *Computer Speech & Language*, 25(2):158–174.
- Epstein, R. (1992). The Quest for the Thinking Computer. *AI Magazine*, pages 81–95.
- Epstein, R. (2007). From Russia, with Love. How I got fooled (and somewhat humiliated) by a computer. *Scientific American Mind*.
- Fahlman, S. E. (2011). Using Scone’s Multiple-Context Mechanism to Emulate Human-Like Reasoning. In *Advances in Cognitive Systems: Papers from the 2011 AAAI Fall Symposium*, pages 98–105. AAAI Press.
- Fialho, P., Coheur, L., Curto, S., Cláudio, P., Ângela Costa, Abad, A., Meinedo, H., and Trancoso, I. (2013). Meet Edgar, a tutoring agent at Monserrate. In *Proc. of the 51st Annual Meeting of the ACL: System Demonstrations*, pages 61–66.
- Følstad, A. and Brandtzæg, P. B. (2017). Chatbots and the New World of HCI. *Interactions*, 24(4):38–42.
- Hutchens, J. L. (1997). How to Pass the Turing Test by Cheating. Technical report, Univ. of Western Australia.
- Kuipers, B., McCarthy, J., and Weizenbaum, J. (1976). Computer power and human reason. *SIGART Bull.*, pages 4–13.
- Levesque, H., Davis, E., and Morgenstern, L. (2012). The Winograd Schema Challenge. In *Proc. of the Thirteenth International Conf. on Principles of Knowledge Representation and Reasoning*, pages 552–561. AAAI.
- Levesque, H. J. (2014). On our best behaviour. *Artificial Intelligence*, 212:27–35.
- Li, J., Galley, M., Brockett, C., Gao, J., and Dolan, B. (2016). A diversity-promoting objective function for neural conversation models. In *Proceedings of the 2016 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language*

- Technologies*, pages 110–119. Association for Computational Linguistics.
- Malinowski, B., (1923). *The Meaning of Meaning*, chapter The Problem of Meaning in Primitive Societies, page 38. Harcourt Brace Jovanovich, Inc.
- Mauldin, M. L. (1994). ChatterBots, TinyMuds, and the Turing test: entering the Loebner Prize competition. In *Proc. of the 12<sup>th</sup> National Conference on Artificial Intelligence (vol. 1)*, AAAI '94, pages 16–21. AAAI Press.
- Powers, D. M. W. (1998). The total Turing test and the Loebner prize. In *Proc. of the Joint Conf. on New Methods in Language Processing and Comp. Natural Language Learning, NeMLaP3/CoNLL '98*, pages 279–280.
- Saygin, A. P., Cicekli, I., and Akman, V. (2000). Turing test: 50 years later. *Minds and Machines*, 10:2000.
- Schneider, K. (1988). *Small Talk: Analyzing Phatic Discourse*. Sprachwissenschaftliche Reihe. Hitzeroth.
- Schumaker, R. P., Ginsburg, M., Chen, H., and Liu, Y. (2007). An evaluation of the chat and knowledge delivery components of a low-level dialog system: The AZ-ALICE experiment. *Decision Support Systems*, 42(4):2236–2246.
- Turing, A. M. (1950). Computing Machinery and Intelligence. *Mind*, 59:433–460.
- Vinyals, O. and Le, Q. (2015). A neural conversational model. *arXiv preprint arXiv:1506.05869*.
- Wallace, R., Tomabechi, H., and Aimless, D. (2007). Chatterbots Go Native: Considerations for an eco-system fostering the development of artificial life forms in a human world. <http://www.pandorabots.com/pandora/pics/chatterbotsgonative.doc>.
- Weizenbaum, J. (1966). ELIZA – a computer program for the study of natural language communication between man and machine. *Comm. of the ACM*, 9:36–45.