

AperTO - Archivio Istituzionale Open Access dell'Università di Torino

How Do People Ascribe Humanness to Chatbots? An Analysis of Real-World Human-Agent Interactions and a Theoretical Model of Humanness

This is the author's manuscript

Original Citation:

Availability:

This version is available <http://hdl.handle.net/2318/1947171> since 2024-10-18T13:08:58Z

Published version:

DOI:10.1080/10447318.2023.2247596

Terms of use:

Open Access

Anyone can freely access the full text of works made available as "Open Access". Works made available under a Creative Commons license can be used according to the terms and conditions of said license. Use of all other works requires consent of the right holder (author or publisher) if not exempted from copyright protection by the applicable law.

(Article begins on next page)

How do people ascribe humanness to chatbots? An analysis of real-world human-agent interactions and a theoretical model of humanness

Amon Rapp^{a*}, Arianna Boldi^b, Lorenzo Curti^c, Alessandro Perrucci^d,
Rossana Simeoni^e

^a Computer Science Department, University of Torino, Torino, Italy, ORCID 0000-0003-3855-9961

^b Department of Psychology, University of Torino, Torino, Italy, ORCID 0000-0002-2825-3443

^c Computer Science Department, University of Torino, Torino, Italy, ORCID 0000-0001-9144-3162

^d TIM S.p.A., Torino, Italy, ORCID 0000-0002-5569-0200

^e TIM S.p.A., Torino, Italy, ORCID 0000-0003-4465-7517

* Corresponding author: Amon Rapp

University of Torino, Computer Science Department – C.so Svizzera 185, 10149 Torino, Italy

Mail: amon.rapp@unito.it

This is an Accepted Manuscript of an article published by Taylor & Francis in International Journal of Human-Computer Interaction on 2023, available online:

<https://www.tandfonline.com/doi/abs/10.1080/10447318.2023.2247596?journalCode=hihc20>

To cite this article: Amon Rapp, Arianna Boldi, Lorenzo Curti, Alessandro Perrucci & Rossana Simeoni (2023): How Do People Ascribe Humanness to Chatbots? An Analysis of Real-World Human-Agent Interactions and a Theoretical Model of Humanness, International Journal of Human-Computer Interaction, DOI: 10.1080/10447318.2023.2247596

How do people ascribe humanness to chatbots? An analysis of real-world human-agent interactions and a theoretical model of humanness

Abstract: A central problem in text-based chatbot research revolves around humanness: not only chatbots may reproduce fundamental aspects of being human, but also users may attribute humanlike features to chatbots. Humanness is paramount in human-chatbot conversations because it may profoundly impact the quality of the interaction: for example, when users expect that interactions with chatbots are similar to human conversations and these expectations are not met by the agent, they may easily become frustrated. This is particularly relevant in customer care, where efficient support is essential: here, unmet expectations about the chatbot’s humanlike capabilities can undermine customer satisfaction. In this article, we qualitatively analyzed 12,477 real-world exchanges with a task-based chatbot in the customer care domain, involving a total of 1,060 conversations. We conducted a single case study to define a preliminary theoretical model of humanness for chatbot technology. The study findings point to a novel conceptualization of humanness in customer-chatbot interaction, highlighting that it is multiple, contextual, modular, and dynamic. Moreover, the theoretical model that we propose explains that the kind of humanness attributed to a chatbot depends on the context in which people expect to interact, the objectives and needs that they aim to fulfill, and the cues that the chatbot exhibits: all these factors may change as the interaction evolves over time, and such changes may further affect the users’ ascriptions of humanness. Finally, we propose design implications of the model, like the need to create “context” and account for the plurality of humanness.

Keywords: conversational agents, chatbots, humanness, human-AI interaction, human-agent interaction, artificial intelligence, customer care

1 Introduction

Over the last ten years, conversational agents, that is software applications interacting with people through natural language (Schuetzler et al., 2020; Rapp et al., 2021), have risen in popularity. In particular, text-based chatbots i.e., agents that are designed to converse and interact with users exclusively through natural written language (Rapp et al., 2021), have been adopted in a variety of domains, from customer service (Følstad & Skjuve, 2019) to mental well-being (Ly et al., 2017), in order to offer people automated support and new forms of companionship. In this landscape, the commercial interest in chatbots for customer service has substantially increased (Nordheim et al., 2019). Gartner highlighted that, in 2021, 27% of organizations had already implemented chatbots, and 70% of customers had tried this technology during their resolution journey (Gartner, 2022a). Gartner also predicted that, by the end of 2023, 38% of organizations will implement chatbots (Gartner, 2022a), and by 2027, chatbots will become the primary customer service channel for roughly a quarter of organizations (Gartner, 2022b). However, despite their increasing spreading, conversing with chatbots appears not to be exempt from issues that directly involve how people experience this technology. Chatbots may not meet the users' expectations or elicit biased user perceptions about their qualities (e.g., Forrester Research, 2016; Grudin & Jacques, 2019).

In this sense, a central problem in text-based chatbots revolves around their humanness: users may attribute humanlike features to chatbots (e.g., Go & Sundar, 2019) and designers may intentionally reproduce fundamental aspects of being human in their design (e.g., Lee et al., 2019). Humanness is central in human-chatbot interaction because it may deeply affect the quality of the conversation, by encouraging users to feel more favorable toward the chatbot (e.g., Mou & Xu, 2017), increasing trust and satisfaction (Rheu et al., 2021). However, humanness may also lead to social anxiety and reduced cooperation (Schanke et al., 2021), or set unrealistic expectations that, if not maintained, may jeopardize the user experience (Jain et al., 2018). In this sense, the chatbot's humanness may impact how the interaction between the user and the agent progresses (or non-progresses), leading to either "good" or "bad" conversations (e.g., Luger & Sellen 2016).

In particular, humanness may be relevant in the customer care domain, where users may need to effectively interact with the chatbot with the main goal of quickly solving an issue, and where expectations about its (humanlike) capabilities may heavily affect the interaction (Rapp et al., 2021): for example, humanness may have a negative effect on the satisfaction of customers in an angry emotional state, their overall evaluation of the firm, and their subsequent intentions of purchasing, as it may inflate expectations of chatbot efficacy, which may be subsequently disconfirmed (Crolc et al., 2022). However, in the customer care domain there is still a lack of understanding about whether and how users ascribe humanness to chatbots, and the influence of humanness on conversational breakdowns or progresses (Law et al., 2022).

Therefore, in this article we tried to give an answer to the following question: **how do people ascribe humanness to text-based chatbots in a conversation in the customer care domain?** To answer this question, we qualitatively analyzed 1,060 real-world conversations with a customer care chatbot endowed with a minimum level of humanlike features, for a total of 12,477 exchanges. Differently from the majority of previous studies that investigated humanness in conversational agents relying on experiments or users' ex-post reports about their conversations with the agent (e.g., Clark et al., 2019; Doyle et al., 2019, 2021), we analyze real-world interactions between users and a chatbot. One could argue that it is not possible to measure the users' perceived human-likeness without asking the users about their perceptions. However, we do not aim to directly address the users' mental states, rather focusing on the users' behavior towards the chatbot, that is how users address the chatbot and interact with it during a conversation. In so doing, we investigate whether users *behave as if the chatbot were* either a human or a machine and the impact that such behavior has on the interaction.

With this aim, we conducted a single case study, using Grounded Theory as a coding method (Glaser & Strauss, 2017), and eventually defined a preliminary theoretical model of humanness for text-based chatbots in customer support that are equipped with a minimum level of humanlike characteristics. A first substantial contribution of this article is to offer rich insights and thick descriptions on the situated experience of

interacting with a chatbot, as well as an in-depth look into real users' attributions of humanness. The findings point out that users do not ascribe humanness to chatbots as an on/off state. Rather, the ascribed humanness presents different shades that span from high (level of) humanness to low (level of) humanness. Second, different kinds of humanness may be ascribed depending on the context (what we called the "script") in which users think to interact (e.g., a complaint sent to a company). Third, users' expectations about the interlocutor's humanness do not address the entire repertoire of human features, rather targeting specific modular human abilities. Fourth, the humanness ascribed by users to the chatbot may change as the conversation progresses. As a second contribution, we propose a conceptual model emerging from the empirical data and further informed by literature, which points out the factors that may influence the users' attributions of humanness to a chatbot in the customer care domain and potentially allows researchers and practitioners to design text-based chatbots that better account for such attributions. In this sense, we describe three main design implications grounded on the model for creating chatbots that may acknowledge how users ascribe humanness to them.

The novelty of the findings and the model revolves around three main points. Firstly, we point out that the perceived context of the interaction shapes how users attribute humanness to the chatbot, while previous studies have emphasized that ascriptions of humanness may be influenced by real situational contexts, such as health or wellbeing (Doyle et al., 2021). In fact, the perceived context may not even have anything to do with the real context: for example, a user might frame the interaction as filling a complaint form, even though the actual context involves chatting with a chatbot. Secondly, we highlight that the capabilities of the chatbot that truly matter for the user to determine its humanness are those that enable it to accomplish her goals and meet her needs, while previous research has mainly focused on the chatbot's conversational skills (Schuetzler et al., 2020). In the customer care domain, these key abilities include, for instance, the capacity to recognize the user's emotions. Thirdly, we emphasize that ascriptions of humanness can evolve in various ways even within the span of a single conversation, whereby previous research predominantly focused on

how humanness attributions may change over the course of mid- and long-term relationships with the chatbot (Skjuve et al., 2022; Nißen et al., 2022).

Moreover, the findings show that the design factors of the chatbot (i.e., the humanlike cues that it has or prompts during the conversation) can only partially explain the users' attributions of humanness: actually, our chatbot was equipped with minimal humanlike features. Rather, such attributions can be also retraced to several user factors, namely how users perceive the context of the interaction, which shapes their expectations of humanness, and whether their goals and needs are fulfilled by specific capabilities of the chatbot.

The article is structured as follows. Section 2 outlines the theoretical background and relevant related empirical literature about humanness. Section 3 describes the method used in this research. Section 4 outlines the study findings, highlighting four main characteristics of humanness in text-based chatbots. Section 5 discusses the findings and presents the theoretical model of humanness we developed. Section 6 outlines our contributions to HCI research and practice. Section 7 describes the limitations of our study, as well as future work, whereas Section 8 concludes the article.

2 Background

2.1 Empirical Research on Humanness and Conversational Agents: The Lack of Studies of Real-World Interactions in Customer Care

Empirical research on humanness has contributed to enriching our understanding of those design factors that may mediate or increase users' perceptions of humanness of conversational agents, as well as to explore how people react to and interact with designs intentionally embedded with humanlike features. In fact, conversational agents are designed to engage individuals in humanlike interactions (Nißen et al., 2022), often exhibiting a variety of social cues (Feine et al., 2019) and mimicking humanlike forms of communication, like speech (Doyle et al., 2019) and gestures (Rueben et al., 2022).

In this sense, humanness seems fundamental to understanding the interaction with conversational agents. For example, Replika, which is a conversational agent having the appearance of a human, appears to push the user's relationship with it toward attachment and perceived closeness through its humanlike ability to support more deep-felt human needs related to social contact and self-reflection (Skjuve et al., 2022). Research, however, also pointed out that current limitations in matching human standards raises questions as to the need and desire to use humanness as an interaction metaphor (Cowan et al., 2017). Conversational agents may shape people's expectations about their capabilities (Luger & Sellen, 2016), and such expectations may negatively affect the interaction experience. Moore et al. (2016), for instance, reported that albeit people are excited about the appearance of speech-based conversational agents, many of them prefer typing over speech.

In particular, research on text-based chatbots points out how design can elicit users' perceptions of humanness, as well as the positive and negative effects that humanlike features may have on the conversation. Nißen et al. (2022) noticed that text-based chatbots that are developed to accompany users over longer periods have more pronounced manifestations of characteristics that anthropomorphize the interaction with respect to chatbots supporting users' short-term goals. However, a name given to the agent (Araujo, 2018), its conversational skills (Schuetzler et al., 2020) and interconnected responses (Go and Sundar, 2019), or its capability of taking initiative and developing a relationship with the user (Morrissey and Kirakowski, 2013) are often sufficient to increase the perception of the chatbot as being humanlike.

These humanlike features may provoke eerie feelings (Liu & Sundar, 2018; Ta et al., 2020; Skjuve et al., 2019; Ciechanowski et al., 2017), or even lower the user's identification with other team members (Mirbabaie et al., 2021). Nonetheless, they may also support the development of intersubjectivity, intimacy, and sense of closeness (e.g., Følstad et al., 2018; Ly et al., 2017; Lee et al., 2020), elicit users' trust towards the chatbot (Shin, 2022), and encourage them to establish a relationship (Xiao et al., 2023), exhibit socially desirable traits (Mou and Xu, 2017), and perform altruistic behavior (Shi et al., 2020).

In any case, expectations about humanlike capabilities of text-based chatbots appear to heavily influence users' satisfaction with the ongoing conversations. Følstad and Skjuve (2019) highlighted that when users hold realistic expectations toward the chatbot, they tend to feel satisfied with the interaction. Conversely, when they expect that interactions with chatbots are similar to human conversations and these expectations are not met by the agent, people become easily frustrated (Jain et al., 2018). This may be particularly problematic when chatbots are employed in the customer service domain, where they are increasingly used (Skjuve et al., 2022). Here, productivity, in the form of quick access to information (Brandtzaeg & Følstad, 2018; Følstad et al., 2018), and efficient and accessible support (Følstad & Skjuve, 2019; Yeh et al., 2022) are essential. When users' expectations are not met, bad conversations may occur: these have been shown to decrease the users' willingness to ask the agent for help with the task (Luger & Sellen, 2016) and to lead to the abandonment of the interaction (Crollic et al., 2022; Jain et al., 2018).

Despite its importance, research on humanness in text-based chatbots for customer care is still scarce (Law et al., 2022). At present, we do not have a clear and complete understanding of whether and how users humanize customer care chatbots during a conversation, what kind of factors may influence the attribution of humanness, and how this may impact the overall interaction. Moreover, most existing work on this domain has not been grounded on "in the wild" conversations but relied either on users' reports on how they perceived the agent (e.g., Law et al., 2022) or on controlled experiments where the conversations were artificial (e.g., Araujo, 2018). Exceptions are represented by Li et al. (2020) and Crollic et al. (2022) who quantitatively analyzed human-chatbot conversation logs. However, the former did not tackle issues related to humanness, while the latter focused only on the emotional aspects of humanness (e.g., whether anthropomorphism may elicit anger in users). Instead, we qualitatively analyzed a series of real-world conversations between users and a customer care text-based chatbot widely exploring their humanizing behavior, that is how users address the chatbot and what kind of characteristics (more or less humanlike) they ascribe to it during the conversation.

2.2 Theoretical Approaches to Humanness in Conversational Agent Research: The Lack of Specific Theoretical Perspectives

At its core, humanness refers to the question of “What is it to be human?”, namely to those attributes that define what it is to be human (Haslam & Loughnan, 2014), and is a central topic of interest in human-agent interaction. In fact, people may attribute humanness also to non-human entities, in a process that is commonly called anthropomorphism (Festerling & Siraj, 2022). Human-Computer Interaction (HCI) researchers have mostly drawn on either general theories of humanness developed in other disciplines like psychology, or technology-focused theoretical frameworks to study humanness in conversational agents.

As for general theories, Epley et al. (2007) proposed a theory of anthropomorphism that predicts that people are more likely to anthropomorphize an object when knowledge about humans and self-knowledge is accessible and applicable, when they are motivated to resolve uncertainty, seek meaning, and feel efficacious, and when they need to establish and maintain a sense of social connection with others. The theory has been found applicable for robots (e.g., Zlotowski et al., 2014), but it has been highlighted that it does not accommodate relevant factors that allegedly affect their humanness, like gestures (Salem et al., 2011) and verbal communication (Walters et al., 2008). Another approach tackled in social psychology has been proposed by Haslam (2006, 2008, 2015), who proposed a model of dehumanization that points to two different senses of humanness. Humanness may be defined in terms of “human uniqueness”, which identifies uniquely human characteristics, such as civility, moral sensibility, and rationality. By contrast, humanness may also be seen as “human nature,” referring to characteristics like emotional responsiveness, agency and interpersonal warmth that are viewed as central to or typical of humans, in a noncomparative sense, and reveal our continuity with other creatures. In this sense, it has been highlighted that there can be even infrahuman or superhuman entities, lacking certain human characteristics (like competence or warmth) but not others (like friendliness or rationality) (Leyens et al., 2003; Li et al., 2014). This theory has been used to ground research on human-agent interaction (Lee et al., 2019; Ghafurian et al., 2019), highlighting that

agents may be humanized along either the uniquely nature dimension, being perceived as intelligent or intentional, or the human nature dimension, being seen as warm and emotional (Zlotowski et al., 2014). Nonetheless, further research pointed out that the theory of dehumanization may not be fully adequate to explain humanness in conversational agents (Zlotowski et al., 2014; Jones et al., 2018).

As for the technology-focused theories, the uncanny valley theory suggests that cues in non-human agents encouraging the attribution of humanlike mental abilities can trigger eerie sensations among users, because their mental categorization conflicts and threatens to human distinctiveness (Stein & Ohler, 2017). When the humanness becomes too accentuated, it may evoke aversive responses rather than affinity (Mori et al., 2012). The theory has been widely applied to the study of chatbots (Liu & Sundar, 2018; Ta et al., 2020; Skjuve et al., 2019; Ciechanowski et al., 2017). Nevertheless, the uncanny valley theory can explain only specific aspects of humanness, which refer to the negative reactions that people may have when encountering humanlike conversational agents. CASA theory, instead, suggests that people apply human social rules and expectations when interacting with computers, responding to them in the same manner as they would toward other people (Lee & Nass, 2010). This theory grounded the study and design of several conversational agents (Purinton et al., 2017; Schuetzler et al., 2019; Lee et al., 2019). However, it has a very broad scope, being thought to be applied to every digital technology, and may not be able to identify the peculiarities that certain technologies may have in eliciting perceptions of humanness (Lang et al., 2013).

Finally, it is worth noticing that there are other theories that have been employed in the context of humanness even though they were not originally conceived for this aim. For example, Social Presence Theory (SPT), which defines social presence as the efficacy of communication media to facilitate a sense of connection with another individual, has been used to describe how conversational skill influences anthropomorphism of a chatbot (Schuetzler et al., 2020). Likewise, the Theory of Mind, which points to the human skill of attributing beliefs, goals, and desires to other people has been applied to the design of agents with similar capabilities (Wang et al., 2021). However, as these theories were not originally conceived to

address the problem of humanness, they may not provide a complete account of how humanization unfolds in human-chatbot conversations.

To summarize, the theories employed by the HCI community to investigate the users' ascriptions of humanness to conversational agents might not fully capture the peculiarities that a specific technology like chatbots may have with reference to humanness.

3 Method

3.1 A Single Case Qualitative Study

We opted for a single case qualitative study, as it best addresses the exploratory nature of our research, which aimed at a deep understanding of the users' ascriptions of humanness to chatbots. A qualitative study could also be more useful to gain rich and in-depth descriptions of such ascriptions, providing a variety of examples of how they unfold (Rapp et al., 2021). In so doing, we used a Grounded Theory approach as a coding method (Glaser & Strauss, 2017), analyzing conversations as if they were data coming from a qualitative study. Typically, Grounded Theory is applied to data coming from interviews and ethnographic field notes, but it has also been used to make sense of other kinds of data, like scientific articles (e.g., Mencarini et al., 2019).

What distinguishes Grounded Theory from other methods is an inductive, rather than a hypothetical-deductive stance (Glaser & Strauss, 2017). It uses coding processes heavily but differs from methods such as content analysis that emphasize reliability, validity and the counting of instances (Furniss et al., 2011; Pidgeon & Henwood, 1996). Grounded Theory has been used in many different ways in both social science and HCI research (Cole & Gillies, 2022; Cimolino et al., 2023), such as a technique for developing complex theoretical models directly emerging from the data, or as a coding method (Furniss et al., 2011). In this study, we used Grounded Theory as a technique for qualitatively analyzing a large amount of data and making sense of it. Following Furniss et al. (2011), while we initially inhibited the use of extant literature to

analyze the data, we then intentionally used it for refining and comparing the emergent model of humanness with previous work (Goldkuhl & Cronholm, 2013).

3.2 Chatbot Characteristics

The chatbot system that we analyzed is task-focused and is built on the Microsoft BOT Framework which relies on natural language understanding and other artificial intelligence capabilities to comprehend customers' requests, using a text-based user interface running on a web site and a mobile app (Figure 1). In particular, the user request is processed by a spellchecker and then analyzed by MS Language Understanding Cognitive Services to infer the user intent and other relevant pieces of information called entities, which are used to correctly address the user's needs. When the user provides out-of-scope or syntactically too-long/too-complex sentences, she is kindly invited to reformulate the request.

The combination of the inferred user intent and entities are used to address a specific dialog flow leading the user through the steps needed to accomplish a specific task in a straightforward and deterministic way. The chatbot's behavior is configured by combining dialog nodes in a flowchart, having one or more starting nodes, where the specific dialog may start, and one or more ending nodes, where the dialog ends up leading the chatbot to engage the customer in another new dialog. The response to be provided to the user is defined at each dialog node, possibly using some linguistic variations to be applied using a random criterion, as well as information retrieved from back-end systems, such as the remaining credit.

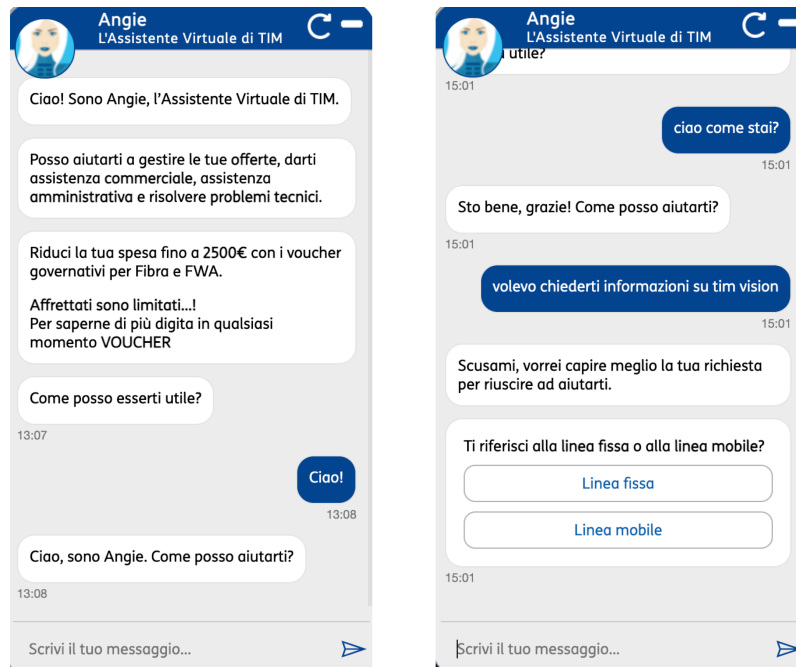


Figure 1. Two examples of the chatbot user interface: on the left, a common dialog flow; on the right, the chatbot prompts two alternatives in the form of quick response buttons¹.

The chatbot is not intentionally designed to encourage the user's projection of humanness: it presents itself as a virtual assistant ("Hello, I am Angie! The TIM virtual assistant..."), is identified by a foreign name (Angie) and is depicted using a stylized manga-style illustration of its face, which is static and cannot express any non-linguistic behavior (e.g., facial expressions or gestures). The chatbot further clarifies its artificial nature when it is not able to handle the user's requests ("Sorry, I'm still a young virtual assistant, but I'm a fast learner! I hope I can be helpful to you next time") by offering the option to call or receive a

¹ The translation of the conversation on the left is: C: "Hello, I am Angie! The TIM virtual assistant" C: "I can help you manage your offers, give you business assistance, administrative assistance, and solve technical problems." C: "Reduce your spending up to 2500€ with government vouchers for fiber and FWA. Hurry they are limited...! To learn more type in VOUCHER anytime." C: "How can I help you?" U: "Hi" C: "Hello, I am Angie! How can I help you?". The translation of the conversation on the right is: U: "Hello, how are you?" C: "I'm fine, thanks. How can I help you?" U: "I wanted to ask you about tim vision" C: "Sorry, I would like to understand your request better to be able to help you" C: "Are you referring to the landline or the mobile line?" C: "Landline" C: "Mobile line".

call from a human operator (“... if you want you can continue the chat with a human operator. He will get back to you as soon as possible and give you all the details”; “If you want, I can pass you on to a human operator who can give you more information. Shall I proceed?”). The repetition of certain expressions word by word may further strengthen the perception of its low humanness (“Sorry, I did not understand your request. Can you repeat it in simpler words?” is a standard sentence that the chatbot produces every time it does not understand the user). Moreover, the chatbot is not meant for engaging people in small talk or general conversations but for specifically addressing technical or commercial issues with the company services. If we refer to the conversational agent taxonomy of social cues defined by Feine et al. (2019), the chatbot shows minimal social cues: it has a low degree of human-likeness and only expresses greetings and farewells, apologizes when it is not able to deal with a request, and provides advice on how to fix certain issues, while it is not able to produce complex sentences. To interact with it, in addition to typing, which is always allowed, the chatbot sometimes provides alternative input modalities, including quick-response buttons.

We focused on a task-oriented chatbot, which typically has a narrow range and keeps conversations short (Grudin & Jacques, 2019), with minimal humanlike characteristics, as we hypothesized that if ascriptions of humanness occur by interacting with this kind of technology, they would likely occur, even to a greater extent, in other more humanlike or human-oriented chatbots, following a critical case sampling strategy (Cardano, 2011). This strategy is based on the argumentation that if a phenomenon can be observed in a condition where it is difficult that it will occur (e.g., when humanness is not intentionally supported), it will be likely encountered even in more favorable contexts (e.g., when humanness is intentionally elicited through design) (Cardano, 2011; Sharp, 1975; Perelman & Olbrechts-Tyteca, 2008).

3.3 Data Characteristics

The initial dataset was composed of 3,123 conversations with a task-based chatbot maintained by a telecommunication company. The company is one of the major telecommunication companies in Italy,

providing mobile and fixed line services (telephone and Internet) with a generalist customer base counting millions of fixed line and SIM card subscriptions. The conversations were collected on the 28th of September 2019. Figure 2 depicts the process of data polishing and analysis.

We first cleaned the data, by manually filtering out 1,052 conversations where the user only greeted the chatbot (e.g., Hi) and immediately abandoned the interaction without making any request (likely due to, for example, internet connection failure), or did not interact at all (they likely only opened the chatbot window to make the agent start the conversation by presenting itself but did not respond). We then selected three time slots to be included in the corpus (9am - 11am, 3pm - 5pm, 9pm - 11pm): for these time slots, we analyzed the entire conversations, for a total of 1,060 conversations and 12,477 exchanges (11.8 exchanges per conversation on average, SD: 7.4).

The choice of sampling the conversations during these hours followed a purposeful sampling method (Marshall, 1996) and relies on the fact that different categories of customers may access the chatbot service at different times (e.g., workers tend to use the chatbot in the evening after office hours), as we understood from preliminary conversations with several company stakeholders. Once the sample was initially defined, we still remained open to the analysis of further data (e.g., from different time slots) if needed. However, after having coded 1,060 interactions we became aware that new data could not lead to new insights for our aims, thus reaching theoretical saturation (Henwood & Pidgeon, 2003).

The conversations are real-world interactions in which users try to fulfill their situated needs and objectives by interacting with the chatbot. While entering the conversation, users had to accept the company's terms of service and were informed about our data-collection activities as in Li et al. (2020) (e.g., that the data could be recorded, analyzed, and published). Moreover, we followed guidelines to ensure ethics in online research (Beekhuyzen et al., 2015; Spicker, 2011; Sveningsson, 2004) and adopted parts of the "heavy disguise" strategy proposed by Bruckman (2002), by deleting all the personal information (e.g., names, telephone numbers) that people could have voluntarily or involuntarily disclosed during the

interaction in a pre-processing phase. All the quotes reported in this article have been further translated from Italian to English while preserving the meaning, and we double checked that no users could be identified from the quotations reported in this article. This strategy has been previously adopted by HCI research investigating, e.g., online forum posts (Su et al., 2019) or in-game chat logs (Rapp, 2022).

3.4 Data Analysis

The data were stored in a spreadsheet, with each row representing one of the dataset's 12,477 exchanges. Data analysis was conducted on the basis of the key principles of Grounded Theory. Analysis of top-down nature, e.g., through the use of pre-existing literature and theory, was inhibited in the coding stages where we emphasized an inductive approach. The first, the second and the third authors read the entire dataset. In the first phase, the second and the third authors separately assigned one or more conceptual labels to each row representing a single user's interaction through open coding. At this stage, the analysis paid attention not only to what the user communicated to the chatbot, but also to the terms used and the language structure of each utterance (e.g., subordinate or coordinate clauses, length of the utterance, repetitions, typos and grammatical errors, etc.). The two authors then discussed the codes generated to assess consistency between them and resolve inconsistencies. In certain cases, inconsistencies were related to the differences in labeling the same concepts. In other cases, two codes were condensed into one or new codes were developed when the discussion between the two researchers led to identify clearer commonalities or distinctions among the meanings of the data points. This process went through the whole data set.

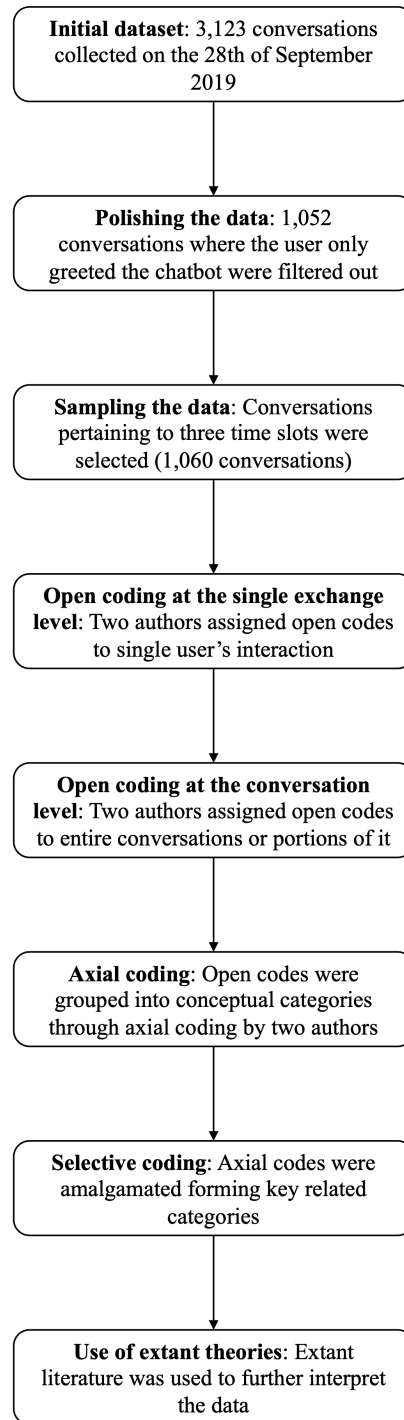


Figure 2: The phases of the data analysis.

As is common in qualitative research, even that adopting a grounded theory approach (e.g., Muller & Kogan, 2010; McDonald et al., 2019), no numerical reliability rating is reported, because our goal was to reach an intersubjective consensus, where each point of difference was debated and clarified until the coders agreed on appropriate usage of the set of codes (Harry et al., 2005). In the second phase, the analysis was conducted at the conversation level (human-machine), whereby the two authors independently assigned codes to entire conversations or portions of it (when it was possible to identify a separate interaction dynamic), focusing on how the conversations evolved from both the semantic (what the user communicates) and syntactic (how she communicates it) points of view. Then, the two authors discussed again the codes generated to resolve inconsistencies (e.g., “collaboration” for the first coder and “agreement” for the second one, which turned out into the decision of keeping the former as the best label describing the underlying concept) and connected the codes generated during the first level (row level) to the codes generated in this phase (conversation level). A final set of 83 open codes was then defined in a codebook. Then, the entire dataset was recoded anew with the codebook at hand.

As a second step, these concepts were grouped into conceptual categories through axial coding by each coder separately. They then discussed the categories together resolving discrepancies, also involving the first author whenever doubts about the decision arose. This yielded 13 learned abstracted categories. Axial coding categories from open coding were amalgamated to create a more defined hierarchy forming key related categories through selective coding. The resultant four selective categories are the central themes emerging from the analysis, which identify key aspects of the concept of humanness in text-based chatbot interactions: multiplicity, contextuality, modularity, and dynamicity.

Finally, we explored how the extant literature could help us in further interpreting our data, using it as leverage for insight that can bring new perspectives (Furniss et al., 2011). On the basis of this last step, we defined a preliminary theoretical model that outlines how ascriptions of humanness are formed and evolve over time in the domain of text-based chatbots for customer care.

4 Findings

In the following, we will recount the main findings coming from our analysis. We first show that humanness implies multiple kinds of being human spanning from high humanness to low humanness. Then, we explain that ascriptions of humanness depend on the perceived context in which the users think to interact.

Subsequently, we describe humanness as a modular concept, as people target specific abilities that are tied to their situated needs and objectives to attribute (or withdraw) humanness to a chatbot. Finally, we outline how humanness is dynamic, as users' ascriptions may evolve as the conversation progresses.

In each sub-section, we will analyze the users' conversational behaviors, identifying interactional patterns and language peculiarities. Then, we will provide an interpretation of these patterns in terms of ascriptions of humanness: a table will summarize the main findings also offering a variety of examples extracted from the data. Of course, we could not access the users' mental states, as we only analyzed their behavioral traces, that is how they conversed with the chatbot: therefore, it is possible to interpret such conversational behaviors even in different ways, for example, by saying that these are the easiest or most convenient ways for the users to state their requests. However, we think that the interpretations that we suggest in terms of ascriptions of humanness can consistently explain the users' conversational behaviors. In this sense, the following findings should be better read not as if the users actually perceived the chatbot in certain ways or they had certain mental states: rather, they indicate that the users *behaved as if* they perceived the chatbot in those ways or as if they had those mental states.

4.1 Humanness is Multiple: Users Behave as if They Ascribe Different Kinds of Humanness to the Chatbot

We identified four different ways through which people address the chatbot, which are signaled by the usage of different personal pronouns to identify the bot², the level of formality employed in the conversations, the

² Italian language has four different pronouns to refer to an interlocutor, that is "lei", "tu", "voi", and "esso".

number of details and explanations provided in the users' requests, and the syntactical complexity of their utterances.

These different ways of addressing the chatbot may point to different kinds of humanness ascribed by the users to it, eventually revealing that humanness is multiple and cannot be described as a single property that the chatbot may or may not have. Rather, it is a characteristic that may vary gradually between high humanness and low humanness, being also possibly conceived as collective humanness. Table 1 summarizes the main findings of this subsection, while in the following we will illustrate them extensively by using the Italian pronouns as a starting point.

Table 1. How users address the chatbot in the conversations.

Kind of humanness	Linguistic indicators	Examples
High humanness: the chatbot is treated as if it were a full human	<ul style="list-style-type: none"> - “Lei” pronoun - Greetings, formalities, politeness - Explanations of issues 	<p>Example 4.1a</p> <p>U: “<i>Good morning, a few hours ago I changed sim since I bought a new phone and changed company. I was told that today, day 28, it would have activated a new rate, I did it, but it has been activated a rate that is not mine. I chose the one with 50gb and unlimited minutes, instead I have 3gb and 1000 minutes. May you [lei] please explain why?</i>”</p>
Diminished humanness: the chatbot is treated as if it were a diminished human	<ul style="list-style-type: none"> - “Tu” pronoun - Imperative and hierarchical tone - No explanations of issues 	<p>Example 4.1b</p> <p>U: “<i>I want a SOS top-up</i>” [...]</p> <p>C: “<i>Sorry, I am still a young virtual assistant, but I’m a fast learner! I hope to be able to help you next time</i>”</p> <p>U: “<i>I want a SOS top-up now</i>” [...]</p> <p>C: “<i>But can’t you [tu] make the SOS?</i>”</p> <p>C: “<i>Sorry, I did not understand your request. Can you repeat it in simpler words?</i>”</p> <p>U: “<i>Can you [tu] do the SOS top-up?</i>”</p>
Collective humanness: the chatbot is treated as if it were a collective entity	<ul style="list-style-type: none"> - “Voi” pronoun - Formal tone - Complex syntactical structure - Courtesy 	<p>Example 4.1c</p> <p>U: “<i>Good evening, I wanted to point out that, even after several calls, with the new offer [service name] the [other service name] has been deactivated and the voice has been activated, in which I was not interested; your representative did not give me information about this point [...]. I would like to ask you [voi] that you [voi] activate again [service name] and deactivate the voice. [...]</i>”</p>
Low humanness: the chatbot is treated as a mere machine	<ul style="list-style-type: none"> - “Esso” pronoun - Command-style language 	<p>Example 4.1d</p> <p>C: “<i>Hello, I am Angie, the TIM virtual assistant</i>”</p> <p>U: “<i>Refund status of the remaining credit of user [phone number] of [user name]</i>”</p>

“Lei” is the third-person singular pronoun in Italian and is commonly used with people that the person does not already know, expressing a high grade of formality, which, nonetheless, is often granted during first interactions, unless both the interlocutors agree upon a more informal setting. In Italian, when people use “lei”, in most cases they grant a fully human status to the interlocutor. This is the reason why “lei” is the most appropriate pronoun to use in a situation when a person is asking someone that she does not know to solve a problem: in this way, the person refers to her interlocutor as a par that may (or may not) provide the requested help. Users addressing the chatbot using “lei” also employ formal greetings (“Good morning”), apologize for the bother (“I am sorry”), are polite when introducing the problem they have, and try to explain their issue as they would do with a human (as in Example 4.1a). This way of addressing the chatbot may be interpreted as the byproduct of the ascription of *high humanness* to the agent, as a signal that the user considers the chatbot as having human dignity, as if the machine were her equal, thus having the right of being respected and treated with kindness.

Other users, instead, employ the second-person pronoun (“tu”) to address the chatbot. In Italian, this pronoun is commonly used with friends, colleagues, and acquaintances who agreed upon an informal tone. Using “tu” with an unknown person, especially when a person needs to ask to solve a problem, as it happens in a customer service situation, is usually perceived as unpolite and/or disrespectful: it does not grant a full status to the interlocutor, because it treats her as inferior in some respects. “Tu,” in fact, in Italian is also often employed with people that are somehow considered having a lower status with respect to the person who is speaking (thus also conveying some in/out-group prejudices), like younger people (especially children), or foreign individuals. Users referring to the chatbot with “tu” also use a less formal and more hierarchical and authoritative tone if compared to users using “lei” (Example 4.1b): the sentences are shorter, and users are less inclined to explain their issue (e.g., U: “*But you [tu] haven’t solved anything at all I haven’t even told you where the problem is*”). This way of addressing the chatbot may be interpreted as the

result of the ascription of a *diminished humanness* to the chatbot, as if it were not able to meet some human standards, lacking certain linguistic, cognitive and relational qualities.

“Voi,” instead, is the Italian pronoun to indicate a collective human entity, like a bunch of people, a specific group, or even an organization. In the analyzed corpus, this pronoun is mostly used when the user wants to manifest her disappointment for an unresolved problem: the tone here is formal, and the language structure is quite complex, with long sentences and the employment of subordinate clauses. When used “voi,” the addressee remains unidentified and generic, as the user shows not to be interested in the specific identity of the interlocutor. This way of addressing the chatbot may be interpreted as the result of an ascription of *collective humanness* to the chatbot: the user behaves as if the chatbot were both a human being, at par of the other company employees that are in charge of solving the users’ problems, and an unidentified being, a member of an indistinct crowd (i.e., the company), which is considered only for its role but is not endowed with a precise identity (Example 4.1c). In other words, the formal tone, the complex syntactical structure, and the courtesy used in the conversation seem to acknowledge a human status. Nonetheless, these users treat the chatbot as a generic representative, as if they were talking with an undefined collective humanlike entity.

Finally, a last way to address the chatbot is identified by the pronoun “esso.” In Italian, this pronoun is the formal way to refer to things, but it is usually omitted in the everyday language. People who interact with the chatbot using “esso” mostly dismiss the sentence subject and use a command-style language, whereby the verbs are in the infinitive form, the syntax is extremely simple, and greetings, articles, circumlocutions, and unusual terms are avoided (like in example 4.1d). This last way of addressing the chatbot may be interpreted as the result of the ascription of *low humanness* to the chatbot. The avoidance of complex terms and language structures, as well as the adoption of command-style language may be explained by the users’ acknowledgment that the interlocutor is a mere machine that needs to be addressed by using a sort of

machinelike language, closer to the needs of the interlocutor than to their own habitual ways of interacting with other humans.

In sum, the four different modalities of addressing the chatbot that we identified may point to four different kinds of humanness that the users ascribe to the agent. In this perspective, rather than an on-off state, humanness appears to be a more nuanced concept, which unfolds along a continuum spanning from high humanness to low humanness. The users that behave as if the bot were a machine represent only a minor quota. Most of them talk to the chatbot as if it were either a human operator or a diminished human, like someone who does not have some fundamental human abilities.

4.2 Humanness is Contextual: The Perceived Context of the Interaction Affects the Users'

Expectations about the Chatbot's Humanness

In the analyzed corpus, we found that the users interact with the chatbot by using diverse communication styles. These styles may be interpreted as the result of the activation of specific perceived contexts that are common in the customer care domain. The interesting thing, here, is that such contexts appear to affect the users' expectations about the chatbot's humanness.

These perceived contexts may be better explained by the notion of "script" (Schank & Abelson, 1977). Scripts are knowledge structures that describe stereotypical sequences of events, possibly with adjustable parameters, allowing an agent to anticipate and figure out what will happen in a certain social interaction, what kind of people she will likely meet, and what social rules she will have to respect. They serve as a guide for action, allowing individuals to know in advance how they should behave and engage socially within a given context. Examples of these stereotyped sequences of actions include visiting a parent, taking the train, and eating in a restaurant.

We may thus interpret the diverse communication styles adopted by the users as the byproduct of the activation of specific scripts at the starting of the conversation, which all relate to the customer care domain, where an individual needs to interact with a company. These scripts refer to the situation where the

interaction with the company is thought to occur (e.g., making a phone call, writing an email, etc.), the kind of interlocutor (human or machine) with which the user will likely interact, and the stereotypical interactional steps that should be followed in that situation with that interlocutor.

Table 2. The scripts that users apparently employ to frame the conversations.

Script	Communication style	Examples
Email script: the user communicates as if she were writing an email to a human	<ul style="list-style-type: none"> - Long, verbose and courteous messages - High degree of syntactical and semantical complexity - Messages do not presume an immediate response 	<p>Example 4.2a</p> <p>U: “<i>Good evening, I would like to receive some clarifications about my subscription. My first question relates to the monthly cost, when I signed the contract, the operator told me about a monthly cost of 30 euros everything included... If you check my bills, you will realize that I have never paid 30 euros. Now, I would like to have a real contract, which I do not have, where the conditions and the rates of the subscription are written – [name of the service]. The last month you charged me even the phone book. I called the [customer care service number] and the operator told me that the cost would have been refunded in the next bill. [...] [Name and surname] owner of the line [telephone number].</i>”</p>
Complaint script: the user communicates as if she were filling a complaint form	<ul style="list-style-type: none"> - Informal tone - Negative emotions - References to past interactions 	<p>Example 4.2b</p> <p>U: “<i>Absurd. Not only you did not make it plain to me why 35 euros per month instead of 30 as you have proposed, but in the last bill you have charged me 7 euros of [service name]. As soon as possible I will change provider for both the mobile phone and the fixed line.</i>”</p> <p>C: “<i>I will try to help you as best as I can.</i>”</p>
Telephonic / chat script: the user communicates as if she were writing into a chat or calling a human operator	<ul style="list-style-type: none"> - Colloquial tone - Abbreviations and emoticons - Grammatical errors and typos 	<p>Example 4.2c</p> <p>U: “<i>Sorry ‘xké’ [Italian abbreviation for “perché” which means why and is used only in chats with humans] [subject wrongly omitted] take away my credit if. [punctuation error] I finished minute [orthographical error it stands for minutes]</i>”</p> <p>C: “<i>Sorry, I did not understand your request [...]</i>”</p> <p>U: “<i>No, I became aware that it took away 90 cents per day</i>”</p> <p>C: “<i>I am sorry, but I cannot help you on this point, I am still learning [...]</i>”</p> <p>U: “<i>I think ‘xké’ [why] there. [punctuation error] wasn't vredit [orthographical error]</i>”</p>
Proper chatbot script: the user communicates as if she were writing to a machine	<ul style="list-style-type: none"> - Collaboration with the chatbot - No grammatical errors or abbreviations - Reformulation of phrases when the chatbot does not understand 	<p>Example 4.2d</p> <p>U: “<i>[service name] and [other service name], I want to unsubscribe from them when they expire. I can do it now</i>”</p> <p>C: “<i>I have two lines in your name 1. [telephone number 1] 2. [telephone number 2]. What line should I address? Please chose one of the aforementioned lines [...] (1, 2, 3, ...)</i>”</p> <p>U: “<i>I</i>”</p> <p>C: “<i>Do you want to know the expiry date of the SIM or the expiry date of the offers?</i>”</p> <p>U: “<i>Expiry date of the offers and deactivation</i>”</p>

This interpretation may explain the reason why the communication styles used in the corpus differ so much: a different style is used depending on the activated script, and thus on the context in which the user thinks to be situated. In this sense, the script also frames the agent's humanness, as each context presumes a specific kind of interlocutor. Table 2 summarizes the four different conversational styles we identified, pointing to four different scripts, which are extensively described in the following paragraphs.

The first communication style is characterized by extremely long and verbose messages filled with all the information that the users think is needed to solve a specific problem: in this sense, the users employ a high degree of semantic (with e.g., common sayings, allusions) and syntactic complexity in their utterances (with chains of coordinate and subordinate clauses). Moreover, messages pertaining to this communication style are generally courteous in their tone, do not presume an immediate answer, rather expect a later notification when the issue will be likely solved. In most cases, people leave their contact information, so that they could be reached later, as well as close the message by saying "thanks in advance" and "looking forward to receiving a response" (Example 4.2a). This style may be interpreted as the result of the activation of an *email script*, which frames the interaction with the chatbot as if the user had to asynchronously write an email to a human addressee, who can understand complex stories and threads of events that originated even weeks or months before, or proactively fix a problem once it is notified.

The second communication style is characterized by messages with a lower level of complexity (e.g., shorter and with less subordinate clauses) and a less formal tone, but more loaded with negative emotions, like anger and irritation (Example 4.2b). Here, the interlocutor is often addressed using "voi," and the users often stress their past interactions with the organization highlighting the unsatisfying responses that they have received. This style may be interpreted as the result of the activation of a *complaint script*, as if the users were complaining at the company's front desk or filling an online form to report a disservice; these situations imply a collective human interlocutor, whose identity is not important, which is expected to fix the reported problem or show empathy for the user.

The third communication style entails a language that resembles more that used in oral conversations than in written ones: here, users greet the interlocutor and present themselves, and then expose the problem they have with little details and a colloquial tone. Several users use abbreviations and emoticons for interacting with the chatbot (e.g., C: “*Hello, I am [chatbot name], the virtual assistant of [company name]*” U: “*What timing :)*”). Moreover, they do not pay attention to their omissions or grammatical errors (Example 4.2c). This style may be interpreted as the outcome of the activation of a *telephonic conversation* or *chat script*, which frames the conversation as a phone call or a chat through an instant messaging application with a human operator, who will be able to fill the gap in the users’ messages and correctly interpret them as a human entity.

The last communication style is also characterized by short utterances and scarce functional words, but seems to point to a much greater level of cooperation between the user and the bot: for instance, when the chatbot proposes a series of predefined options among which the user has to choose to move the conversation forward, she either selects the correct option using the numbers proposed by the chatbot or repeats the option word by word, without introducing any variation (Example 4.2d). Likewise, users using this style often reformulate their sentences, or provide yes/no responses when the chatbot asks them a question, with almost no traces of grammatical errors, abbreviations, or emoticons. This style may be explained by the adoption of a *proper chatbot script*, in which users show to insert the conversation in a proper context that provides them with the “correct” expectations about the chatbot, allowing them to follow those social and communicative rules that may enable a shared understanding on the basis of its machinelike capabilities.

To summarize, the different communication styles that we identified in the analyzed corpus may be interpreted as the result of the activation of different scripts, which mostly imply a human interlocutor. In the email sending, complaint form filling, telephone calling/chat scripts the interlocutor is expected to be a human, as she best fits the frame anticipated by the scripts: a human addressee in the email script; a company representative (or the company itself) with reference to the complaint form; and a customer care operator in

the case of the phone call and chat script. Only a minor quota of users, instead, appears to activate a proper chatbot script that implies an artificial interlocutor.

4.3 Humanness is Modular: Users Target Specific Abilities when They Ascribe Humanness to the Chatbot

As we have noticed, when the users begin a conversation with the chatbot they may already have activated a certain script, which anticipates the kind of humanness that they will ascribe to the chatbot. Users, however, appear not to be concerned about the whole spectrum of qualities that we can find in a real human being, rather connecting the humanness of the chatbot with specific humanlike skills. Humanness may thus be seen as consisting of diverse and even separate abilities that are selectively targeted by users, depending on the perceived context of the interaction (the script) and the related user goals and needs (e.g., solving a problem, expressing frustration). In our corpus, we found that the targeted abilities are: comprehending what the user is communicating, reasoning over what she has said, remembering what she has previously reported, feeling what she is experiencing, and executing what she is requesting, and can be indicated by the interaction pattern that the users employ with the chatbot (See Table 3 for a summary).

In fact, a first interaction pattern refers to those users that mostly interact within the email and complaint scripts, using extremely complex language structures and mentioning intricate problems (e.g., that do not have an evident possible solution or involve long narratives) that they aim to resolve (Example 4.3a): these users either use compound sentences connecting many independent clauses, or complex sentences with a variety of subordinate clauses, like chains of relative clauses. They also often fill their messages with a variety of contextual information and details about the problem to be solved. This way of interacting with the chatbot may be interpreted as the byproduct of an expectation that the user has about the interlocutor's humanlike skill to fully understand her messages formulated in human language: the chatbot should be able to decipher the complex syntax of the sentences, understand the focus of the story, distinguish the provided contextual information on the basis of its pertinence, and connect together the many details given.

Table 3. The humanlike abilities targeted by the users.

Ability	Interaction pattern	Examples
<p>Understanding ability: capability of fully understanding messages formulated in human language</p>	<p>- Complex stories, intricate problems, and overload of information and details</p>	<p>Example 4.3a U: “<i>Within a few weeks there have been two malfunctions in the connection, not on my phone line, but in general. Last time both Internet and the fixed line, this time “only” Internet. This happened yesterday more or less at twelve. This morning, I had an appointment with the technician of [a TV service from another company], which I had to postpone, since, in fact, there is no connection. I understand that malfunctions cannot be foreseen, but I believe that, as I am not being able to use the service for so many days and, above all, not having been able to have the work done by the technician, it has created an inconvenience for me and so I am entitled to a refund.</i>”</p>
<p>Reasoning ability: capability of inferring, from a few traces left by the user, the true meaning of the sentences and the context of the interaction</p>	<p>- Insufficient details, errors, and omissions</p>	<p>Example 4.3b U: “<i>Good morning, I activated the offer with moden [orthographical error] [name of the modem misspelled] but [misspelled] at the moment it does not have the technician [wrong syntax]?He [no space] installed normal [name of other modem] [wrong syntax]?I [no space] wonder when she brings the moden [orthographical error] back that I requested?</i>”</p>
<p>Memory ability: capability of remembering past interactions</p>	<p>- References to what has previously happened in the conversation and to past interactions</p>	<p>Example 4.3c C: “<i>Hello, I am Angie, the TIM virtual assistant [...]</i>” U: “<i>I wrote on Messenger</i>” C: “<i>Sorry, I did not understand your request [...]</i>” U: “<i>You have written to me on Messenger that you would have canceled [name of the service] and brought the old line back since November</i>” C: “<i>The service can be requested only if you buy the modem [name of the modem] [...] Do you want to know the offer?</i>” U: “<i>No, I have a problem with the invoiceeeeeeeeeeeeeees</i>” C: “<i>For which invoice do you want more details? [...]</i>” U: “<i>You have written to me on messenger that you would have removed [service name in capital letters] AND brought everything back to the old line already from November</i>”</p>
<p>Emotional ability: capability of recognizing the user’s emotions or empathizing with her</p>	<p>- Expressions of distress, anger, or frustration</p>	<p>Example 4.3d U: “<i>I have already requested the payment by direct debit several times but it still lacks</i>” C: “<i>I can explain to you how to request the direct debit or I can give you information about the status of your request. What do you prefer?</i>” U: “<i>Info about the STATUS of the request</i>” [...] U: “<i>It’s been A LOT OF MONTHS</i>” C: “<i>Sorry, I did not understand your request [...]</i>” U: “<i>I want to know about the STATUS of my request of direct debit that I DID MONTHS AGO</i>” U: “<i>I HAVE ALREADY SAID. THE STATUS OF MY REQUEST</i>”</p>
<p>Ability of acting: capability of executing the user’s requests</p>	<p>- Requests for action</p>	<p>Example 4.3e U: “<i>I wanted to rectify my house number in the invoice, which is [address] and not [address], since the invoice comes the building next door. Thank you</i>”</p>

A second interaction pattern can be especially encountered in the chat script, where the user omits relevant parts of the syntactic structure of the sentence, misses to give sufficient details about the context, and puts in the conversation many orthographical errors or abbreviations (Example 4.3b). Here, the users mostly explicit that they need to rapidly get information related to a problem that they had (e.g., U: “*I don’t have time to waste*”), without giving too many explanations. This interaction pattern may be interpreted as the result of expectations about the interlocutor’s reasoning skills: the chatbot is expected to infer, from a few traces left by the user, either the true meaning of the sentences, even when they are wrongly written, or further details about the wider context of the problem (Example 4.3b).

A third kind of interaction pattern revolves around the user’s reference to previous parts of the conversation or to episodes happened in the past. The user provides content that is broken in two or more phrases but still refers to the same unit of meaning, returns to previous parts of the conversation, or mentions past interactions with the chatbot or with the whole company. This pattern may be interpreted as the result of expectations about the memory function of the chatbot, which is expected to be able to connect different sentences or parts of sentences, remember what has been said before or what happened in the past. This kind of memory goes beyond machinelike information retrieval, rather involving a humanlike capability of recollecting past episodes and connecting them to the present situation for a better understanding of the current problem (Example 4.3c).

A fourth kind of interactional patterns points to the user’s expression of emotions. In these cases, which occur mainly when people interact within the complaint script, users need to express an emotional distress, in the form of anger, frustration and dissatisfaction. In the Example 4.3d, the user progressively raises her “voice” (by capitalizing more words) to manifest her feelings to the chatbot. This interaction pattern may be interpreted as the byproduct of the user’s expectations about the chatbot’s capability of understanding what she is feeling with reference to the problem at hand. By expressing anger, frustration and dissatisfaction, the user behaves as if the chatbot were at least able to recognize such emotions, somehow putting itself in her

shoes and acknowledging the difficulties that she has encountered in facing the problem. However, the fact that the chatbot cannot recognize her emotions, like in Example 4.3d, may make the user become even more frustrated and angry.

Finally, a last interactional pattern concerns the request of doing things. The user, here, may have the need to, for instance, change some data in her own profile stored by the company and ask the chatbot for some sort of action. This pattern may be interpreted as the result of expectations about the chatbot's capabilities of executing what the user has been requested, showing a sort of agency in the real world. In the Example 4.3e, the user behaves as if the chatbot is expected to be able to elaborate appropriate strategies for satisfying her need and then to execute a proper sequence of actions that will lead it to achieve the objective, or at least give her some suggestions on how to fix the issue.

To summarize, users appear to target specific human capabilities when ascribe humanness to the chatbot. The skills addressed are tied to the goals and needs that the user wants to satisfy during the interaction: for instance, when the user needs to solve a complex problem, she seems to expect that the chatbot is able to understand complex language structures and stories. In fact, the chatbot's denial of such expectations may entail frustration or dissatisfaction and the willingness to quit the conversation; however, violations of users' expectations may also lead to transformations in the user's perception of humanness as we will see in the next subsection.

4.4 Humanness is Dynamic: Ascriptions of Humanness Evolve as the Interaction Proceeds

As the conversation proceeds, people's expectations about the chatbot's humanness and humanlike skills may change or remain the same. This, in turn, may affect the kind of humanness that they attribute to the chatbot. Rather than being a static perception, humanness is thus something dynamic, which evolves as the interaction progresses.

We found that humanness may evolve (or not) along four different directions. A certain ascription of humanness can be downgraded to a lower level of humanness when the user's expectations are not fulfilled

(Human→Machine). In parallel, people may also relapse to a higher degree of humanness, when, albeit they previously acknowledged the machine nature of the chatbot, they activate certain interactional patterns that presume specific humanlike skills in the interlocutor (Machine→Human). By contrast, people can also keep their expectations about the chatbot's humanness unchanged despite the cues about its machine nature prompted during the interaction (Human→Human). Finally, they may ascribe low humanness to the chatbot and maintain such ascription for the whole duration of the conversation (Machine→Machine). Table 4 summarizes these four kinds of evolution, which we will explain in detail in the following.

The most common evolution that we discovered in the analyzed conversations refers to the first configuration (Human→Machine). People enter the conversation behaving as if they were talking to a fully human interlocutor, using complex and courteous language. As the interaction proceeds, the communicative style used changes toward a command-style language (Example 4.4a). This evolution of the interaction may be interpreted as the byproduct of a change of the script that frames the interaction (e.g., from a complaint script to a proper chatbot script), which happens when the expectations that the user had about the chatbot's humanness are not met and need to be redefined. This may result in the user ascribing low humanness to the chatbot, starting to adapt her messages to the chatbot's conversational style, rephrasing her own sentences, and making them simpler and easier to understand. Alternatively, the user may start treating the agent as a sort of diminished human, when the expectations that she had about specific humanlike abilities that she considers important to fulfill her objectives and needs are disregarded (e.g., the reasoning ability). Here, the user may modify her sentences toward a greater simplicity but also adopting an imperative or aggressive tone (Example 4.4b), until she may decide to exit the conversation.

The second configuration (Machine→Human) is also not uncommon in the analyzed corpus. People enter the interaction by using a command-style language, short sentences, and simple syntactical structures. However, as the interaction progresses, they turn into using more complex language structures and expressing emotions to the chatbot (Example 4.4c, whereby when the chatbot asks the user whether she

wants to turn to an operator she starts complaining with the chatbot as if it were a human). This kind of evolution may be interpreted as the result of a shift in how the user ascribes humanness to the chatbot, moving from low humanness to a higher degree of humanness (Machine→Human).

Table 4. Different types of evolution of the users' ascription of humanness to the chatbot

Evolution type	Changes in the kinds of humanness ascribed	Linguistic indicators	Examples
Human→Machine	The user shifts from ascribing high humanness to ascribing low humanness or diminished humanness to the chatbot	- From complex and courteous (e.g., use of “lei”) language (high humanness) to command-style language (low humanness), or to simple informal language (e.g., use of “tu”) with imperative tone (diminished humanness)	<p>Example 4.4a U: “<i>Good morning, may I ask a question?</i>” C: “<i>Sorry, I did not understand your request. Can you repeat it in simpler words? [...]</i>” U: “<i>In the invoice, I pay 1 euro to see [service name]. Since I don’t care, is it possible to cancel my subscription and stop paying 1 euro per month?</i>” C: <i>You need: 1. Information about the cost 2. Support for the registration 3. Technical assistance 4. Information for the cancellation</i>” U: “4”</p> <p>Example 4.4b U: “<i>Good morning I activated the [service name] offer that lasts 3 days at 0.99 cents but it still doesn't let me send messages on whatsapp or update instagram. Why?</i>” [...] U: “[service name in capital letters] <i>doesn't work</i>” C: “[...] “<i>Do you want to know the details of the transactions?</i>” U: “<i>Okay, let it go</i>”</p>
Machine→Human	The user shifts from ascribing low humanness to ascribing high humanness or collective humanness to the chatbot	- From command-style language (low humanness) to complex and formal language (e.g., use of “lei” or “voi”) and expressions of emotions (high humanness or collective humanness)	<p>Example 4.4c C: “<i>Here it is the detail of the expense items present in the invoice [phone Number] for the month of January for the line [phone Number]: 1 Subscriptions [amount] 2 Phone calls [amount] 3 Other amounts [amount]. Do you want to have the detail of a specific item present in the invoice?</i>” U: “3” C: “<i>Here is the detail requested for the invoice of January [phone Number] [...]. Is the charge correct?</i>” U: “<i>Absolutely no</i>” C: “<i>If you want, I can also turn your report over to an operator, do you want me to do that?</i>” [...]</p>

			<p>U: “<i>I see that I called the [number of customer care], I ask now here, but [company name] doesn’t want to solve my problem, I will think about it in a different [orthographical error] way</i>”</p> <p>Example 4.4d C: “<i>Do you need to: 1. Have information for a requested refund 2. Request a refund [...]</i>” U: “<i>2</i>” C: “<i>[...] [company name] offers its customers the possibility to choose a compensatory solution as an alternative to the reimbursement of the days eroded as per [regulatory authority] resolution</i>” U: “<i>No I already know them, we have already done these steps, I prefer the refund and not the compensatory solution</i>”</p>
Human→Human	The user maintains the ascription of high humanness or collective humanness to the chatbot for the whole conversation	- Human complex language (high humanness or collective humanness) maintained for the whole conversation	<p>Example 4.4e U: “<i>Good morning, we are always in the same situation, I haven’t activated any service, in the invoices I find certain amounts charged, every time I have credit, you charge me for unknown things. I am really getting annoyed, this is not an isolated case.</i>” C: “<i>[...] On the line [phone Number] there are movements for the following categories [...] Do you want to know the details of the movements?</i>” U: “<i>I want my money back immediately</i>” C: “<i>I check what charges there have been on your [phone Number] [...] Do you want to know the details of the movements?</i>” U: “<i>Other does not belong to me</i>” C: “<i>Do you want to change the topic?</i>” U: “<i>I want my money back</i>”</p>
Machine→Machine	The user maintains the ascription of low humanness to the chatbot for the whole conversation	- Command-style language (low humanness) for the whole conversation	<p>Example 4.4f U: “<i>Deactivate paper invoice</i>” C: “<i>[...] To cancel the paper invoice you need to activate [service name] [...]. If you want, we can do it TOGETHER or I can give you a GUIDE to read whenever you want. What do you prefer?</i>” U: “<i>Together</i>” C: “<i>OK [... a series of instruction]. Go ahead and tell me when you’re done so I can explain the next step...</i>” U: “<i>I’m done</i>”</p>

These changes may be due to changes in the user’s goals and needs targeting humanlike abilities. In the example 4.4c, the impossibility of fixing the issue entails an emotional state that the user needs to express,

expecting that the chatbot will be able to be empathic with her: the initial need to solve a problem (which can be dealt by a machine and does not require humanization) is substituted with the need to be emotionally understood, which entails a humanlike entity having humanlike empathic skills. In other cases, the chatbot may prompt cues that trigger the ascription of higher levels of humanness, as in Example 4.4d, where a direct offer from the chatbot, even if it is not intentionally designed to elicit humanness attribution, fosters the user to communicate as if it were a human (by providing an explanation that it cannot understand).

By contrast, the third configuration (Human→Human) can be identified when the user maintains her interaction modality unchanged for the whole conversation. In this “static evolution” the user seems to ignore all the cues that the chatbot prompts suggesting that it is, in fact, a machine, and keeps using human complex language, also refusing to choose the predefined options offered by the chatbot or to rephrase what the chatbot has not understood (Example 4.4e). This kind of evolution may be interpreted as the byproduct of the preservation of the initial ascription of high level of humanness to the chatbot, where users avoid acknowledging the chatbot’s machine nature and keep demanding that the interlocutor expresses humanlike skills.

Finally, in the last configuration (Machine→Machine), people maintain their interactional pattern using simple language structures and a command-style language. Here, they seem to be open to tuning their interaction to the chatbot’s needs for the whole duration of the conversation (Example 4.4f). This kind of evolution may be interpreted as the result of the maintenance of the initial ascription of low humanness to the chatbot (Machine→Machine).

In sum, the ascription of humanness is often not static when interacting with a text-based chatbot. As the interaction evolves, users may change their ascription of humanness because they select another script, they change their objectives and needs, or the chatbot triggers ascriptions of higher degree of humanness. However, it is also possible that the initial ascription of humanness frames the whole conversation.

5 Discussion

The analysis of the conversations we have reported above shows that the concept of humanness is complex and multifaceted: humanness is multiple, contextual, modular, and dynamic, which means that it takes different forms, depends on the situation in which the user thinks to interact, refers to specific abilities, and evolves over time. In the following, we will discuss in depth these findings, connecting them with previous research literature on humanness and conversational agents. Then, we propose a preliminary theoretical model of humanness aimed at explaining how users ascribe humanness to chatbots in the customer care domain. Finally, we present some design implications of our model.

5.1 Scripts Shape the User's Expectation's about the Interlocutor's Humanness

We noticed that the user frames the conversation with the chatbot within a specific script, which shapes her expectations about the interlocutor: the scripts imply anticipations about the typical actants with which the user will likely converse in a specific situation, conveying initial expectations about the interlocutor's humanness. The notion of script echoes that of partner models, which states that people enter dialogue with assumptions about their interlocutors (Clark, 1996). Doyle et al. (2021) hypothesized that the salience of certain dimensions characterizing the partner models may vary within certain situations but did not offer empirical data to support their hypothesis. Our findings provide such data suggesting that the model of the interlocutor is affected by the specific context in which the user thinks to interact: for example, if she thinks to interact in a complaint script she will likely expect that the interlocutor is humanlike. Differently from previous research, we showed that even in a well-circumscribed domain as the customer service domain, people may refer to completely different scripts, which may entail completely different ascriptions of humanness: an insight that comes from our study, therefore, is that analyzing and understanding these contexts becomes fundamental when designing a chatbot.

Moreover, in the analyzed corpus, the majority of the users began the interaction within scripts implying human interlocutors. This easiness in assuming that the conversational partner is human is confirmed by

previous research on text-based chatbots (Araujo, 2018; Go & Sundar, 2019; Beattie et al., 2020; Shi et al., 2020; Jain et al., 2018), and may be explained by the fact that people have the tendency to assign human characteristics to computers, as stated by CASA theory (Lee & Nass, 2010), as well as by the fact that users may find in the chatbot's written conversations certain cues (even unintentionally designed) that trigger ascriptions of high level of humanness (as in Example 4.4d).

Moreover, text-based chatbots are characterized by the uncertainty about the interaction environment: users commonly converse with a chatbot in a pop-up window of the browser, or through communication channels commonly used for interacting with other humans (e.g., on Facebook), which may not constitute, at least for some users, an identifiable space that may activate the right social and conversational rules. This may further explain why so many users initially selected scripts pertaining to other, similar, Computer Mediated Communication (CMC) channels like sending an email, filling a form, or conversing through an instant messaging application, which all entail a human interlocutor. In sum, the chatbot's behavior and the scarce cues characterizing the communicative space might favor the activation of scripts referring to more familiar communicative situations, bringing along corresponding expectations about the humanness of the interlocutor.

5.2 User's Objectives and Needs Target Specific Humanlike Abilities in the Chatbot

The fact that most people initially select scripts implying a human interlocutor does not imply that they will seek in the chatbot all the features that characterize a human entity during the conversation. People do not need that the chatbot shows the entire repertoire of abilities that a real human owns in order to maintain or change a specific perception of its humanness: as we pointed out in Section 4.4, it is sufficient that the chatbot does not show a specific capability that the user considers important for satisfying her needs to lead her to withdraw humanness from the chatbot (Example 4.4b); by contrast, an emerging user need requiring a specific ability from the chatbot to be satisfied may elicit the ascription of higher levels of humanness (Example 4.4c). In fact, users target specific capabilities that allow them to pursue their objectives and

satisfy their needs, without considering other, less salient (at the moment), abilities: these specific capabilities, in turn, influence the kind of humanness that is ascribed to the chatbot; (e.g., if a capability is not present, the kind of humanness ascribed to the chatbot may be transformed). In the same line, Doyle et al. (2021) hypothesized that models of a partner might not be comprehensive at all times, as users might focus on a limited number of attributes at a time: however, this hypothesis was not based on empirical data. Our findings corroborate such hypothesis with empirical findings, highlighting that attributions of humanness are fundamentally modular and specifying that shifts in the focus on humanness modules may be driven by shifts in the person's objectives and needs: when users are driven by different needs and objectives other modules may be activated and become more important.

We found that in the customer care domain, the targeted abilities refer to comprehension, memory, reasoning, understanding emotions, and agency, which parallel the needs and goals that users typically have in such a domain (e.g., the need to express frustration). However, chatbots designed for other domains, e.g., aimed at providing forms of companionship, may activate expectations about other capabilities of the chatbot, for instance, to develop a relationship over time, in response to different users' needs and goals. For example, Lee et al. (2020) stressed the importance of the chatbot's capabilities of self-disclosing in a humanlike way to make people answer highly sensitive questions. These abilities were mostly not sought by our users. Our users addressed the agent mainly to solve a specific problem: therefore, self-disclosing skills were not considered when ascribing humanness to the chatbot.

5.3 Ascriptions of Humanness Change Even During a Single Conversation

The ascriptions of humanness may also vary during the interaction, being the result of how the conversation progresses. Research on text-based chatbots has highlighted that perception of a chatbot's anthropomorphism and intelligence may change significantly over time, that is over several months (Wang et al., 2021; Skjuve et al., 2022). Nißen et al. (2022) also pointed out that ascriptions of humanness may be modified during shorter interactions, like with temporary advisor chatbots, which are addressed to support medium-term

relationships with users. Our findings highlight that perceptions of humanness may evolve even in the course of a single conversation. In fact, these variations may be encountered even in the middle or in small portions of the conversations, as fluctuations in the users' humanness perceptions. Doyle et al. (2021) hypothesized that salience of partner models' dimensions may change during a conversation but did not identify either how partner models could evolve with reference to humanness (i.e., the direction of change) or the specific interaction events that could yield the change. Our study, instead, identifies four kinds of directions for the evolution of humanness over time and highlights that such evolutions may be due to the shift of the script framing the interaction, the modification of the user's objectives and needs, or to the cues prompted by the agent, which can all be modified as the conversation progresses. These findings, therefore, highlight important sensitive points, which may allow designers to intervene in an ongoing conversation for modifying on the fly the users' attribution of humanness. For example, clearly self-disclosing the chatbot as an artificial entity at the right time point (i.e., when the system detects that the user is ascribing high humanness to the chatbot) may at first cause uncertainty, if the user had perceived the agent as a humanlike entity, but subsequently can help her better set her expectations about the interaction.

5.4 Humanness may Take Different Forms Revealing the Variety of Ways of Being Human, Almost Human, or Machine

The study findings show that humanness in text-based chatbots can only be accounted by a nuanced, multi-faceted concept, which implies different forms of humanness. In between the extremes of high humanness and low humanness there is a continuum of neither human nor machine beings, which are forms of diminished humanness resulting from the non-recognition of key humanlike capabilities that are considered relevant by the user in a specific situation.

Haslam (2006) highlighted that people can be dehumanized in various ways depending on which human attributes are denied. Likewise, recent research on dehumanization has emphasized that in addition to the traditional understanding of dehumanization as an explicit and complete deprivation of humanness, there

exist more nuanced forms of denying humanness, whereby the perception of others is framed as less than (but not necessarily non-) human (Li et al., 2014). Leyens et al. (2003) called this phenomenon *infrahumanization*, which may treat dehumanized targets as entities lacking certain human characteristics (like competence and agency) but not others (like friendliness). Li et al. (2014) showed that certain forms of dehumanization align with paternalistic stereotypes, which treat, for example, women and the elderly as lacking competence but not warmth. This suggests that there is no single essence of being human than can be denied as an on/off state, but rather multiple core aspects of humanity that can be denied in different ways. The findings emerging from our study highlight similar conclusions, as they show that people may perceive a chatbot endowed with different degrees of humanness and even with different separate human abilities, without considering a unique essence that may definitely grant the agent the status of human. In other words, there are different pieces of humanness, which can be granted or denied depending on the context and the person's objectives and needs.

This portrayal of humanness may appear utilitarian to a certain extent. We project humanness onto other entities with extreme facility because we have something that we want to accomplish or that we desire to obtain by interacting with these other entities and we presume that our needs will be better satisfied by our human counterparts, at least because we *“have direct and immediate access to the phenomenological experience of being a human but do not have such immediate access to the phenomenological experience of any nonhuman agent”* (Epley et al., 2007, p. 868). As Epley et al.'s theory of anthropomorphism suggests, people reason about the mental states of others through a process of egocentric simulation and correct that simulation only when they have the capacity, motivation, and requisite representations about others to do so (Epley et al., 2007, p. 868). Therefore, when we expect that the “other” will be able to help us, as when we use technology not merely as a tool but as an agent with which we have to interact (as it happens when we interact with technology through language), it may be sufficient that this other has the ability (or that we believe that it has the ability) to satisfy our targeted needs to grant it a human status.

However, at the same time we also tend to withdraw humanness with extreme facility as soon as the entity is not able to satisfy our desires, violating our expectations about its usefulness with respect to our objectives. This is certainly more visible in utility-driven domains, as the customer-care domain, in which our study was conducted. But it is not hard to imagine that our dispositions to ascribe or withdraw humanness to chatbot may be elicited by similar drivers in other, apparently less utilitarian, domains. For instance, in a conversational chatbot we may seek a companion, a friend with which we can converse, self-disclose and empathize. Depending on the chatbot's abilities to fulfill these expectations, we might either grant it a high degree of humanness or dehumanize it by downgrading its status to a mere object or tool.

The utilitarian nature of processes of (de)humanization has been also found in dehumanization research. Haslam et al. (2008) highlight that perceiving others as less than human, likening them to beasts or unfeeling objects, and treating them with inhumanity are not only common occurrences in times of war, genocide, and ethnic conflict, where such practices are used to justify and encourage cruelty and atrocities, but can also be observed in more everyday abuses, where different forms of dehumanization are employed as strategies to satisfy certain desires or achieve certain objectives. For instance, in intergroup contexts, dehumanization is a means for differentiating groups, as fewer characteristics that are unique to humans are ascribed to outgroup members than to ingroup members for maintaining power relationships and inequalities (e.g., Haslam, 2015). Likewise, objectification, such as sexual objectification, is a strategy to be legitimized to treat the target as disposable for fulfilling one's own desires (Li et al., 2014).

It is not surprising, therefore, that when we interact with an "unknown other", humanness is ascribed not only because of our tendency to assimilate the other to ourselves, but also because it is the best way to assume that the other will be capable of responding to our desires. This may point to the fact that we did not already develop ways for considering technology as an "other" that can really help us through its own ways (and not through our humanlike ways).

Moreover, it would also come as no surprise that novel forms of humanness ascribed to chatbots will appear, as their capabilities will evolve, thus responding differently to people's expectations. For instance, the progress in the chatbot's rational and language skills, if not paired with a similar evolution in their emotional capabilities (like that of expressing empathy) could lead to make people perceive the chatbot as a superhuman. Superhumanization implies perceiving the others as unemotional, rigid, and yet highly intelligent, ascribing characteristics that transcend ordinary humanness to the target (Li et al., 2014).

5.5 A Theoretical Model of Humanness

On the basis of the study findings, we will now develop a preliminary theoretical model of humanness for text-based chatbots in the customer care domain. Although the data on which this model is grounded come from a specific text-based chatbot, we believe that it could have a wider applicability. In fact, this particular chatbot has not been intentionally designed to elicit ascriptions of humanness and shows minimal social cues: nonetheless, it appears that users consistently attribute humanness to it, which may signal a spontaneous users' tendency with respect to chatbot technology in this domain. This tendency, which points to several user factors sustaining the ascription of humanness, like the script adopted to frame the interaction and the specific goals and needs that users have, could then be even more evident in more favorable contexts, namely, when a customer care chatbot is purposefully designed to encourage perceptions of humanness and shows more social cues.

We want to specify that we do not aim to propose a sort of ultimate theory of humanness for conversational agents. Rather, we propose a preliminary and local theoretical model of humanness for text-based conversational agents in customer support, which are equipped with a minimum level of humanlike characteristics. At this stage, this is a preliminary model, which will require further empirical testing to prove its validity. It could then also be refined, expanded, or amended by further research investigating other kinds of conversational agents in order to widen its generalizability.

Figure 3 depicts the model showing that the kind of humanness that people ascribe to a chatbot depends on i) the user's expectations about the overall humanness of the interlocutor (e.g., highly human, lowly human), which are shaped by the perceived context in which they believe to interact, namely, the script; ii) the specific humanlike abilities that they expect to find in the chatbot (e.g., its memory or emotional abilities), which are targeted by the users' objectives and needs; iii) the chatbot's cues, which may elicit ascriptions of high or low humanness, even though they were not intentionally designed to these aims (as in Example 4.4c). All these factors may influence the user's ascription of a specific kind of humanness, as well as change as the conversation evolves over time: namely, the script framing the interaction may be modified (e.g., from a complaint script to a proper chatbot script) changing the user's expectations about the overall humanness of the chatbot; different user's objectives and needs may come to the foreground (e.g., from the need to solve a problem, to the need to be emotionally understood) modifying the humanlike targeted abilities of the chatbot; the chatbot may prompt different cues, triggering diverse ascriptions of humanness. However, initial perception of humanness may also be maintained for the whole conversation, or brief oscillations may be encountered between different humanness forms.

The model may tie together and explain some of the scattered and contrasting results about humanness that have been found in previous research. For instance, it may clarify why different studies point out different users' expectations about the very same chatbots' skills, which may differently affect their perception of humanness. Svenningsson and Faraon (2019), for example, found that people expect that chatbots not involve humor in customer service. By contrast, Ruan et al. (2019) and Ceha et al. (2021) stressed the importance of humor in learning for making users feel the chatbot like a humanlike study partner. These domains entail different user goals and needs, so that the same humor feature may be considered differently. Furthermore, the model may provide empirical evidence about characteristics of humanness that were only hypothesized in previous research on conversational agents but were not based on empirical data, like its contextuality and dynamicity (Doyle et al., 2021). In both cases, the proposed model

develops these previous conceptualizations and further adds novel aspects that may better account for the unique characteristics of text-based chatbot technology in the customer care domain.

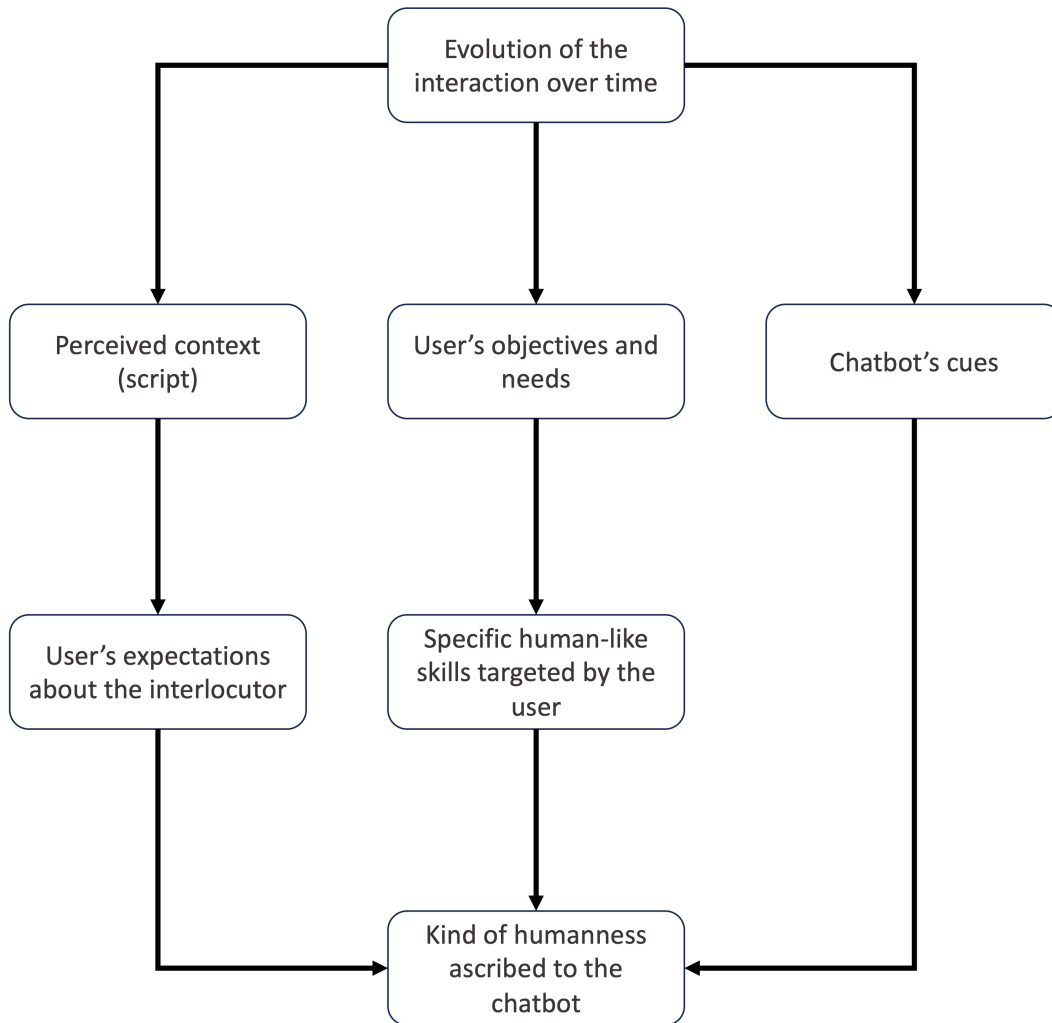


Figure 3. The model of humanness for text-based chatbots

5.6 Implications for Design

We now turn to discuss the opportunities and the challenges for creating text-based chatbots that account for the modalities of ascribing humanness to this kind of technology. Although the primary purpose of our analysis was not to generate design recommendations but to provide a theoretical model explaining people's ascriptions of humanness to chatbots, we may identify several design implications of the model for designing text-based chatbots.

5.6.1 *Understanding and Creating Context*

We discovered that users begin the interaction with a chatbot activating scripts that contain the interactional and social rules to be followed when conversing with the agent. The scripts refer to the contexts in which the conversation is inscribed, tackling familiar situations that people take as a model for interacting with the chatbot, thus working as a reference point that shapes their expectations about the conversation and the interlocutor (and thus its humanness). In our study, we have seen that most users connect the interaction with the chatbot to the well-known situations which all imply a human interlocutor.

This tendency to referring to familiar situations, in which nonetheless a human agent is involved, is also common in other chatbot domains, like health, where the user may point to environments in which she is visited by a human doctor showing empathy (Nadarzynski et al., 2019) and learning, where the interaction may be connected to the classroom where she may converse with classmates (Ruan et al., 2019).

As the scripts affect the interaction with the chatbot, shaping the user's initial expectations, researchers developing new conversational agents should carefully explore the domain in which the chatbot is deployed, identifying those typical situations that the user may take as a model for conversing with the agent. For a healthcare chatbot, for instance, these situations may be a first visit or a regular visit in the doctor's office (if the user returns several times for asking advice with reference to a particular health issue), an interaction with real doctors via an online forum, a consultation requested through an email (e.g., sending to the doctor a medical report), a telephone call to make an appointment, and so on. These different situations may entail

different kinds of interlocutors, as well as diverse interactional and social rules, leading different initial expectations: accounting for as much as possible scripts that users could likely activate in a given domain, would then be paramount to design ways for presenting the chatbot, as well as designing dialog flows and interaction modalities that can manage the different ways through which the users imagine the interaction.

For instance, in our case, immediately detecting that the user enters the conversation being framed in a complaint script, would likely allow the chatbot to perform a series of countermeasures to reduce the user's anger, avoiding increasing her sense of frustration, for example, by providing information when she seeks a vent. In other words, the early detection of the activated scripts and their coupling with the chatbot's behavior and way of presenting itself, as well as an in-depth understanding of how such scripts unfold, would reduce the possibilities of misalignment between the user's expectations and the chatbot's actual behavior. Qualitative techniques, like ethnography, observation and semi-structured interviews could be extremely useful for gaining knowledge about those scripts that pertain to a particular domain in which the chatbot will be deployed.

Moreover, we have seen that in the ongoing conversation, certain chatbot's cues may favor the emergence of humanlike expectations, while uncertain communicative environments, which do not exhibit recognizable chatbot features, may hinder the activation of a proper modality for interacting with the agent. Designers should then better characterize the conversational space of the chat with the chatbot, finding ways to highlight how it differs from other CMC communicative channels implying a human interlocutor, in order to favor the formation of more realistic user expectations. Researching how to design a distinctive chatbot environment, even pursuing the creation of standards that could allow users to immediately recognize the kind of technology they are interacting with (as we have standard interaction environments for sending an email, despite the different interfaces in which this communicative space is encapsulated), could work toward this aim.

5.6.2 Designing for Targeted Humanlike Abilities

We found that users expect that the chatbot has specific abilities depending on the needs and goals that they have during the conversation. Designers should then carefully identify what the primary needs and goals that drive the interaction with a chatbot are and focus on the chatbot's abilities that may satisfy such needs. We have found that in the customer care domain, such abilities relate to language comprehension, memory of past interaction, understanding of feelings (especially of frustration and dissatisfaction), and agency for solving circumscribed problems. Most users started interacting by presuming that the chatbot has such abilities at a humanlike level.

It is important, therefore, that designers make the chatbot clearly express whether it is able to satisfy (or more likely to what extent) the users' expectations related to such key abilities, which should be presumed to be initially set at the human standard level. Since these expectations may be unrealistic for many of these abilities, given the current systems' technical capabilities, it is essential that the chatbot works for explaining what it is not able to do, otherwise the user could easily become frustrated as our study and previous research suggest (e.g., Jain et al., 2018).

Cowan et al. (2017) proposed a similar suggestion with reference to speech-based conversational agents, highlighting that using the design of the interlocutor to signal abilities closer to the system's actual capabilities could help manage user expectations. Likewise, Langevin et al. (2021) suggested a heuristic for designing conversational agents that stresses the importance of clarifying the agent's capabilities.

We agree on this point, but we add that only certain abilities may be relevant for a certain domain, and that such abilities will be likely initially believed to be owned by the chatbot at a humanlike level. Thus, a thorough preliminary investigation of what kinds of human skills users expect to encounter in a given context is essential to design an effective chatbot in that context. As such expectations may also change as the conversation evolves, as we have seen in our study, developers should try to detect what ability is targeted by the user at a given moment and explain the extent to which it can be met by the agent. For

instance, if the user expects that the chatbot is able to remember interactions conducted in the past, the chatbot should capture this expectation and, in case, express that it does not have a memory and that the user needs to repeat what is happened before. In other words, generally stating that it is not able to do a thing might not be sufficient: users should be made aware of the reasons why something cannot be done, connecting it with the ability that would have made it possible (instead of simply stating “*Sorry, I did not understand your request*” for every interaction that the chatbot is not able to handle, it could say “*Sorry I cannot understand complex sentences with words that are not included in my vocabulary or with syntactical structures containing many subordinates,*” if the user’s targeted ability is language comprehension; or “*Sorry I cannot retrace past interactions because I do not have a memory*”). This would help users set more properly their specific expectations about the chatbot, diminishing the risk of non-progresses and abandonments.

5.6.3 *Designing for the Plurality of Humanness*

We stressed that people ascribe various forms of humanness to a chatbot, and these forms may change during the conversation. We also discovered that users may ignore the cues given by the chatbot to indicate its machine-like nature, and may have expectations toward greater humanness despite a variety of other signals going in the opposite direction. This suggests that designers should attempt to endow the chatbot with the ability to recognize the kind of humanness that the user is ascribing to it.

Wang et al. (2021) recommended that designers equip conversational agents with an analog of human Theory of Mind ability that can automatically identify user perceptions about the agent. We emphasize that such chatbot’s ability should primarily target the user’s perception of chatbot’s humanness and be ready to be updated in the course of the conversation, as such perception may rapidly change. Then, designers should allow the chatbot to decide whether to proactively encourage the user to change her perception or to indulge her attribution of humanity. There are clearly pros and cons of these two strategies, which should be decided depending on the situation: forcing users to adopt a perception of low humanness when they persevere in

ascribing high humanness to the chatbot may increase their sense of frustration rather than set more realistic expectations; by contrast, supporting unsupported attribution of high humanness may raise ethical concerns, as the user may be deceived and behave differently than she would if she knew that the interlocutor was a machine.

Another consideration relates to the fact that there is room to use the multiplicity and dynamicity of humanness as design material. On the one hand, users could be actively engaged in shaping the humanness of the chatbot on the basis of their desiderata, making it behave as they like or as they ideally expect, at least at the interaction style level (e.g., through a command-style language, or a human level language). On the other hand, designers could experiment on designing unexplored forms of humanness (e.g., that of a collective entity) or even novel forms of humanity, by endowing the chatbot with a peculiar humanness that is neither human nor machine, thus situating among those forms of infrahumanization that we encountered in the analysis: however, such novel forms, rather than being conceived as diminished forms of humanity (as users spontaneously do when they withdraw humanity from a chatbot that is initially believed to be human), could be thought as different forms of humanity, being intentionally designed to encourage users to recognize hybrid ways of being human and adapt their behavior accordingly.

6 Contribution to HCI Research and Practice

This article makes two substantial contributions. First, it describes how people ascribe humanness to chatbots in real-world interactions in the customer care domain, showing that even when a chatbot is equipped with a minimum level of humanlike characteristics people may behave as if it were a human. Second, it proposes a preliminary theoretical model, emerging from the analyzed data, which explains how such humanness is ascribed in the customer care domain, allowing researchers and practitioners to design text-based chatbots that better account for the users' attributions of humanness.

On the one hand, the novelty of our *contribution to HCI research* revolves around three main aspects. First, even though it may appear obvious that the context matters for the perceived humanness, our study

highlights that what really matters for humanness ascriptions is the perceived context of the interaction, rather than the actual context. A script is a mental entity and may be activated even when the real context has nothing to do with it: nonetheless, it heavily impacts the user's expectations about the humanness of the interlocutor. Second, our study highlights that for the user, in order to treat the chatbot like a human, what really matters is that it has the ability to help and satisfy her needs. Previous research (e.g., Schuetzler et al., 2020) mainly emphasized that the conversational skills may influence the perceived anthropomorphism of a chatbot. Instead, we highlight that such skills might not be sufficient, as the abilities that users consider important for attributing humanness are strictly tied to their goals and needs. Third, we point out the high degree of dynamicity of humanness ascriptions. Previous research noticed that perception of humanness may evolve over mid- and long-term relationships with conversational agents (Skjuve et al., 2022; Nißen et al., 2022). Instead, we discovered that humanness attributions may quickly change even during an ongoing conversation.

On the other hand, our *contribution to HCI practice* lies not only in a model that enables researchers and practitioners to understand and thoroughly analyze the users' ascriptions of humanness to a chatbot during the interaction, but also in three main practical implications grounded on the model for designing new conversational agents that better account for such ascriptions. First, as the scripts shape the user's initial expectations about the chatbot's humanness, we recommend that designers investigate the domain in which the chatbot is deployed, in order to identify those typical situations that the user may take as a model for conversing with it. Second, as users expect that the chatbot has specific abilities depending on the needs and goals that they have during the conversation, we suggest that designers make the chatbot clearly express whether it is able to satisfy the expectations related to such key abilities. Third, as users may attribute different forms of humanness to a chatbot, we recommend that designers endow the chatbot with the capability of identifying the kind of humanness that the user is ascribing to it.

7 LIMITATIONS AND FUTURE WORK

Our analysis has focused on a particular chatbot in the customer care domain, which may limit the generalizability of our findings. However, the fact that we found prominent ascriptions of humanness in a domain where users commonly interact quickly with the chatbot without aiming to develop humanlike relationship (Grudin & Jacques, 2019) and where the agent has not been intentionally designed to convey perceptions of humanness may signal that the phenomena we encountered could likely occur even in more favorable contexts (Cardano, 2011), for example, where the chatbot is purposefully developed to elicit attribution of humanness. However, the chatbot we employed has limitations in its Natural Language Processing capabilities and the interactions occurred with users were likely shorter if compared with conversations that could occur with chatbots addressed to provide forms of companionships (since in our case people had the goal to quickly fix a given problem). To fully validate the generalizability of the study findings to other domains, and thus of the model, similarly detailed conversation logs from other chatbot services should be analyzed. Moreover, it is questionable that our findings would be generalizable across different cultures, being our analysis limited to the Italian population.

The analyzed data did not include any user reflections on their conversations and ascriptions of humanness, only the conversations themselves. For this, it is possible that the inferences we made on certain conversational behaviors might not correspond to what the user experienced during that interaction. Different interpretations of their behavioral traces are possible since there is no data from the users on how they intended their language use. In other words, we cannot affirm that users had for sure certain experiences or mental states; rather, we claim that they behaved as if they had such experiences and mental states, and that our interpretations of their behavior consistently explain their conversational behavior. Nonetheless, we have to notice that, as researchers (and as humans), even through interviews we have no direct access to someone else's first-person perspective and that what we can investigate are only the linguistic traces that mediate every individual's experience (Rapp et al., 2019; Schwartz et al., 2013). In this sense, we think that the value

of our analysis is precisely to address real-world conversations rather than ex-post users' reports on their behaviors. The similarities that we found between our findings and previous research based on users' reports corroborates the validity of our interpretations (and that of previous research). Future work could then compare the analysis of conversations with the users' opinions about such conversations, in order to identify further correspondences and misalignments.

However, without asking users directly, we were not able to relate participants' individual differences to perceptions and expectations about the chatbot's humanness. Research could thus investigate more in depth the role of individual differences in determining expectations about the humanness of chatbots with reference to certain expected abilities, as different users in a given context may have different expectations about a specific skill of the chatbot.

Future works could then test the model on different kinds of chatbots deployed in other domains (e.g., conversational chatbots, healthcare chatbots), in order to understand whether and which other modules of humanness may be targeted more prominently by users when the contexts and the user goals and needs change. Future research should also tackle ethical issues emerging from the apparently users' tendency to ascribing humanness to text-based chatbots. The modularity of humanness signals that future designs could easily elicit perceptions of humanness by addressing only the target humanlike capabilities that are relevant to a certain situation (e.g., self-disclosure in a chatbot for mental well-being). This raises ethical concerns especially if such ascriptions of humanness are meant to subtly encourage users to behave in certain ways (e.g., to give information that the user would not have otherwise disclosed to a machine).

8 CONCLUSION

In this article, we explored how people ascribe humanness to text-based chatbots in the customer care domain. We discovered that humanness is multiple, as it assumes different forms; contextual, as it is tied to the context in which the user thinks to interact; modular, as users target specific humanlike capabilities; and dynamic, as humanness evolves over time. We then proposed a theoretical model explaining that attributions

of humanness to text-based chatbots in customer care depend on the user's perceived context, her objectives and needs, and the chatbot's cues.

Our contribution to HCI research points out that crucial factors in shaping the users' attributions of humanness are the perceived, rather than the actual, context of the interaction, and the chatbot's capabilities in fulfilling the user's objectives and needs, rather than its conversational skills. Moreover, it highlights that attributions of humanness may change even within the span of a single conversation. In parallel, our contribution to HCI practice suggests that designers should: i) carefully explore the context in which the chatbot is deployed; ii) make the chatbot clearly express whether it is able to satisfy the users' expectations related to its key abilities; and iii) endow the chatbot with the capacity to recognize the kind of humanness that the user is ascribing to it.

9 DISCLOSURE STATEMENT

The research was partially funded by TIM S.p.A. within the project "Studi e ricerche su sistemi conversazionali intelligenti" agreement n. S19AARLL. The funding source owned the chatbot investigated in this article and the dataset, and had the unique role of collecting the raw conversational data. It had no other roles in the analysis and interpretation of data, in the writing of the article (apart from the writing of the paragraph that describes the technical details of the chatbot, i.e., "3.3 Chatbot characteristics", provided by Alessandro Perrucci), or in the decision to submit it for publication in this journal, and in no way influenced the findings reported in this article.

REFERENCES

Araujo, T. (2018). Living Up to The Chatbot Hype: The Influence of Anthropomorphic Design Cues and Communicative Agency Framing on Conversational Agent and Company Perceptions. *Computers in Human Behavior*, 85, 183-189. <https://doi.org/10.1016/j.chb.2018.03.051>

Beattie, A., Edwards, A.P., & Edwards, C. (2020). A Bot and a Smile: Interpersonal Impressions of

Chatbots and Humans Using Emoji in Computer-mediated Communication. *Communication Studies*, 71(3), 409–427. <https://doi.org/10.1080/10510974.2020.1725082>.

Beekhuizen, J., von Hellens, L., & Nielsen, S. (2015). Illuminating The Underground: The Reality of Unauthorised File Sharing. *Information Systems Journal*, 25(3), 171–192. <https://doi.org/10.1111/isj.12069>

Brandtzaeg, P. B., & Følstad, A. (2018). Chatbots: changing user needs and motivations. *Interactions*, 25(5), 38–43. <https://doi.org/10.1145/3236669>

Bruckman, A. (2002). Studying The Amateur Artist: A Perspective on Disguising Data Collected in Human Subjects Research on The Internet. *Ethics and Information Technology*, 4(3), 217–231. <https://doi.org/10.1023/A:1021316409277>

Cardano, M. (2011). *La Ricerca Qualitativa*. Bologna: Il mulino.

Ceha, J., Lee, K. J., Nilsen, E., Goh, J., and Law, E. (2021). Can a Humorous Conversational Agent Enhance Learning Experience and Outcomes?. In *Proceedings of the 2021 CHI Conference on Human Factors in Computing Systems (CHI '21)*. Association for Computing Machinery, New York, NY, USA, Article 685, 1–14. <https://doi.org/10.1145/3411764.3445068>

Ciechanowski, L., Przegalinska, A., & Wegner, K. (2017). The Necessity of New Paradigms in Measuring Human-Chatbot Interaction. In *Proceedings of the International Conference on Applied Human Factors and Ergonomics*, pp. 205–214. Springer, Cham. https://doi.org/10.1007/978-3-319-60747-4_19

Cimolino, G., Chen, R., Gutwin, C., & Graham, T. N. (2023). Automation Confusion: A Grounded Theory of Non-Gamers' Confusion in Partially Automated Action Games. In *Proceedings of the 2023 CHI Conference on Human Factors in Computing Systems (CHI '23)*. Association for Computing Machinery, New York, NY, USA, Article 569, 1–19. <https://doi.org/10.1145/3544548.3581116>

Clark, H. (1996). *Using language*. Cambridge University Press, Cambridge, UK. <https://doi.org/10.1017/CBO9780511620539>

Clark, L., Doyle, P., Garaialde, D., Gilmartin, E., Schlögl, S., Edlund, J., Aylett, M., Cabral, J., Munteanu, C., & Cowan, B. (2019). The State of Speech in HCI: Trends, Themes and Challenges. In *Interacting with computers*, 31(4), 349-371. Oxford University Press. <https://doi.org/10.1093/iwc/iwz016>

Cole, T., & Gillies, M. (2022). Emotional exploration and the eudaimonic gameplay experience: a grounded theory. In *Proceedings of the 2022 CHI Conference on Human Factors in Computing Systems (CHI '22)*. Association for Computing Machinery, New York, NY, USA, Article 470, 1–16. <https://doi.org/10.1145/3491102.3502002>

Cowan, B. R., Pantidi, N., Coyle, D., Morrissey, K., Clarke, P., Al-Shehri, S., ... & Bandeira, N. (2017). "What Can I Help You with?" Infrequent Users' Experiences of Intelligent Personal Assistants. In *Proceedings of the 19th International Conference on Human-Computer Interaction with Mobile Devices and Services*. Association for Computing Machinery, New York, NY, USA, Article 43, 1–12. <https://doi.org/10.1145/3098279.3098539>

Crolic, C., Thomaz, F., Hadi, R., & Stephen, A. T. (2022). Blame the Bot: Anthropomorphism and Anger in Customer–Chatbot Interactions. *Journal of Marketing*, 86(1), 132–148. <https://doi.org/10.1177/00222429211045687>

Doyle, P. R., Clark, L., & Cowan, B. R. (2021). What Do We See in Them? Identifying Dimensions of Partner Models for Speech Interfaces Using a Psycholexical Approach. In *Proceedings of the 2021 CHI Conference on Human Factors in Computing Systems (CHI '21)*. Association for Computing Machinery, New York, NY, USA, Article 244, pp. 1–14. <https://doi.org/10.1145/3411764.3445206>

Doyle, P. R., Edwards, J., Dumbleton, O., Clark, L., & Cowan, B. R. (2019). Mapping Perceptions of Humanness in Intelligent Personal Assistant Interaction. In *Proceedings of the 21st International Conference on Human-Computer Interaction with Mobile Devices and Services (MobileHCI '19)*. ACM, New York, Article, pp. 1–12. <https://doi.org/10.1145/3338286.3340116>

Epley, N., Waytz, A., & Cacioppo, J. T. (2007). On Seeing Human: A Three-Factor Theory of

Anthropomorphism. *Psychological Review*, 114(4), 864–886. <https://doi.org/10.1037/0033-295X.114.4.864>

Feine, J., Gnewuch, U., Morana, S., & Maedche, A. (2019). A taxonomy of social cues for conversational agents. *International Journal of Human-Computer Studies*, 132, 138-161.

<https://doi.org/10.1016/j.ijhcs.2019.07.009>

Festerling, J., & Siraj, I. (2022). Anthropomorphizing technology: a conceptual review of anthropomorphism research and how it relates to children’s engagements with digital voice assistants.

Integrative Psychological and Behavioral Science, 56(3), 709-738. <https://doi.org/10.1007/s12124-021-09668-y>

Følstad, A., Nordheim, C. B., & Bjørkli, C. A. (2018). What makes users trust a chatbot for customer service? An exploratory interview study. In *Internet Science: 5th International Conference, INSCI 2018*, St. Petersburg, Russia, October 24–26, 2018, Proceedings 5 (pp. 194-208). Springer International Publishing. https://doi.org/10.1007/978-3-030-01437-7_16.

Følstad, A., & Skjuve, M. (2019). Chatbots for customer service: user experience and motivation. In *Proceedings of the 1st International Conference on Conversational User Interfaces*, 1-9, ACM, New York. <https://doi.org/10.1145/3342775.3342784>

Forrester Research. (2016). *Chatbots 101: Building Conversational Interfaces. Technical report*. Accessed October 15, 2020. Retrieved from: <https://www.forrester.com/report/Chatbots-101-Building-Conversational-Interfaces/RES136476> (Last access 10/6/23).

Furniss, D., Blandford, A., & Curzon, P. (2011). Confessions From a Grounded Theory Phd: Experiences and Lessons Learnt. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems (CHI '11)*. Association for Computing Machinery, New York, NY, USA, pp. 113–122. <https://doi.org/10.1145/1978942.1978960>

Gartner. (2022a). The Chatbot Deployment Guide. Improve the customer service experience at a lower cost. Technical report. Retrieved from: <https://www.gartner.com/en/customer-service->

support/trends/customer-service-chatbot-

guide?utm_campaign=RM_GB_2022_GCSSL_NPP_PR1_CHATBOT (Last access 10/6/23).

Gartner. (2022b). Gartner Predicts Chatbots Will Become a Primary Customer Service Channel Within Five Years. Blog Post. Retrieved from: <https://www.gartner.com/en/newsroom/press-releases/2022-07-27-gartner-predicts-chatbots-will-become-a-primary-customer-service-channel-within-five-years> (Last access 10/6/23).

Ghafurian, M., Budnarain, N., & Hoey, J. (2019). Role of Emotions in Perception of Humanness of Virtual Agents. In *Proceedings of the 18th International Conference on Autonomous Agents and MultiAgent Systems (AAMAS 0'19)*. International Foundation for Autonomous Agents and Multiagent Systems, Richland, SC, 1979–1981.

Glaser, B. G., & Strauss, A. L. (2017). *Discovery of Grounded Theory: Strategies for Qualitative Research*. Routledge, Evanston, IL.

Go, E., & Sundar, S. S. (2019). Humanizing chatbots: The effects of visual, identity and conversational cues on humanness perceptions. *Computers in Human Behavior, 97*, 304–316.
<https://doi.org/10.1016/j.chb.2019.01.020>.

Goldkuhl, G., & Cronholm, S. (2010). Adding Theoretical Grounding to Grounded Theory: Toward Multi-Grounded Theory. *International Journal of Qualitative Methods, 187–205*.
<https://doi.org/10.1177/160940691000900205>.

Grudin, J., & Jacques, R., (2019). Chatbots, Humbots, And the Quest for Artificial General Intelligence. In: *Proceedings of the 2019 CHI Conference on Human Factors in Computing Systems*. ACM, New York, pp. 1–11. <https://doi.org/10.1145/3290605.3300439>.

Harry, B., Sturges, K. M., & Klingner, J. K. (2005). Mapping the Process: An Exemplar of Process and Challenge in Grounded Theory Analysis. *Educational Researcher, 34(2)*, 3–13.
<https://doi.org/10.3102/0013189X034002003>

Haslam, N. (2006). Dehumanization: An integrative review. *Personality & Social Psychology Review*, 10, 252-264. https://doi.org/10.1207/s15327957pspr1003_4

Haslam, N. (2015). Dehumanization and intergroup relations. In M. Mikulincer, P. R. Shaver, J. F. Dovidio, & J. A. Simpson (Eds.), *APA handbook of personality and social psychology*, Vol. 2. Group processes (pp. 295–314). American Psychological Association. <https://doi.org/10.1037/14342-011>

Haslam, N., & Loughnan, S. (2014). Dehumanization and infrahumanization. *Annual Review of Psychology*, 65, 399–423. <https://doi.org/10.1146/annurev-psych-010213-115045>

Haslam, N., Loughnan, S., Kashima, Y., & Bain, P. (2008). Attributing and Denying Humanness to Others. *European Review of Social Psychology*, 19, 55-85. <https://doi.org/10.1080/10463280801981645>

Henwood, K. & Pidgeon, N. (2003). Grounded Theory in Psychological Research. In P. Camic, J. Rhodes, & L. Yardley (Eds.), *Qualitative research in Psychology: Expanding Perspectives in Methodology and Design*. APA, Washington DC.

Jain, M., Kumar, P., Kota, R., & Patel, S. N. (2018). Evaluating And Informing the Design of Chatbots. In *Proceedings of the 2018 Designing Interactive Systems Conference*, 895-906, ACM, New York. <https://doi.org/10.1145/3196709.3196735>

Jones, K. S., Niichel, M. K., & Armstrong, M. E. (2018). Robots Exhibit Human Characteristics: Theoretical and Practical Implications for Anthropomorphism Research. In *Companion of the 2018 ACM/IEEE International Conference on Human-Robot Interaction (HRI '18)*. Association for Computing Machinery, New York, NY, USA, pp. 137–138. <https://doi.org/10.1145/3173386.3177053>

Lang, H., Klepsch, M., Nothdurft, F., Seufert, T., & Minker, W. (2013). Are Computers Still Social Actors? In *CHI '13 Extended Abstracts on Human Factors in Computing Systems (CHI EA '13)*. Association for Computing Machinery, New York, NY, USA, pp. 859–864. <https://doi.org/10.1145/2468356.2468510>

Langevin, R., Lordon, R. J., Avrahami, T., Cowan, B. R., Hirsch, T., & Hsieh, G. (2021). Heuristic

Evaluation of Conversational Agents. In *Proceedings of the 2021 CHI Conference on Human Factors in Computing Systems (CHI '21)*. Association for Computing Machinery, New York, NY, USA, Article 632, 1–15. <https://doi.org/10.1145/3411764.3445312>

Law, E. L. C., Følstad, A., & Van As, N. (2022). Effects of Humanlikeness and Conversational Breakdown on Trust in Chatbots for Customer Service. In *Nordic Human-Computer Interaction Conference (NordiCHI '22)*. Association for Computing Machinery, New York, NY, USA, Article 56, 1–13. <https://doi.org/10.1145/3546155.3546665>

Leahu, L., Cohn, M., & March, W. (2013). How Categories Come to Matter. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems (CHI '13)*. ACM Press, Paris, France, 3331. <https://doi.org/10.1145/2470654.2466455>

Lee, J. R., & Nass, C. I. (2010). Trust in computers: The Computers-Are-Social-Actors (CASA) Paradigm and Trustworthiness Perception in Human-Computer Communication. In: Lee, J-E.R., & Nass, C.I. (Eds.), *Trust and technology in a ubiquitous modern environment: Theoretical and methodological perspectives*. IGI Global, Hershey, PA, pp. 1–15. <https://doi.org/10.4018/978-1-61520-901-9.ch001>

Lee, M., Ackermans, S., van As, N., Chang, H., Lucas, E., & IJsselsteijn, W. (2019). Caring for Vincent: A Chatbot for Self-Compassion. In *Proceedings of the 2019 CHI Conference on Human Factors in Computing Systems*, 1-13. ACM, New York. <https://doi.org/10.1145/3290605.3300932>

Lee, S., Lee, N., & Sah, Y. J. (2020). Perceiving a Mind in a Chatbot: Effect of Mind Perception and Social Cues on Co-presence, Closeness, and Intention to Use. *International Journal of Human-Computer Interaction*, 36(10), 930-940. <https://doi.org/10.1080/10447318.2019.1699748>

Leyens, J. P., Cortes, B. P., Demoulin, S., Dovidio, J., Fiske, S. T., Gaunt, R., et al. (2003). Emotional Prejudice, Essentialism, and Nationalism. *European Journal of Social Psychology*, 33, 703–717. <https://doi.org/10.1002/ejsp.170>

Li, C. H., Yeh, S. F., Chang, T. J., Tsai, M. H., Chen, K., & Chang, Y. J. (2020). A Conversation

Analysis of Non-Progress and Coping Strategies with a Banking Task-Oriented Chatbot. In *Proceedings of the 2020 CHI Conference on Human Factors in Computing Systems (CHI '20)*, 1-12. ACM New York.

<https://doi.org/10.1145/3313831.3376209>

Li, M., Leidner, B., & Castano, E. (2014). Toward A Comprehensive Taxonomy of Dehumanization: Integrating Two Senses of Humanness, Mind Perception Theory, and Stereotype Content Model. *TPM-Testing, Psychometrics, Methodology in Applied Psychology*, 21(3), 285–300.

<https://doi.org/10.4473/TPM21.3.4>

Liu, B., & Sundar, S. S. (2018). Should Machines Express Sympathy and Empathy? Experiments With a Health Advice Chatbot. *Cyberpsychology, Behavior, and Social Networking*, 21(10), 625-636.

<https://doi.org/10.1089/cyber.2018.0110>.

Luger, E., & Sellen, A. (2016). Like Having a Really Bad PA: The Gulf Between User Expectation and Experience of Conversational Agents. In *Proceedings of the 2016 CHI Conference on Human Factors in Computing Systems*, ACM, pp. 5286–5297. <https://doi.org/10.1145/2858036.2858288>

Ly, K. H., Ly, A. M., & Andersson, G. (2017). A Fully Automated Conversational Agent for Promoting Mental Well-Being: A Pilot RCT Using Mixed Methods. *Internet interventions*, 10, 39-46.

<https://doi.org/10.1016/j.invent.2017.10.002>

Marshall, M. (1996). Sampling for qualitative research. *Family Practice*, 13(6), 522–526,

<https://doi.org/10.1093/fampra/13.6.522>

McDonald, N., Schoenebeck, S., & Forte, A. (2019). Reliability and Inter-rater Reliability in Qualitative Research: Norms and Guidelines for CSCW and HCI Practice. *Proc. ACM Hum.-Comput. Interact.* 39, 39, Article 39 (November 2019), 23 pages. <https://doi.org/10.1145/3359174>

Mencarini, E., Rapp, A., Tirabeni, L., Zanacano, M. (2019). Designing Wearable Systems for Sport: A Review of Trends and Opportunities in Human-Computer Interaction. *IEEE Transactions on Human-Machine Systems*, 49(4), 314–325. <https://doi.org/10.1109/THMS.2019.2919702>.

Mirbabaie, M., Stieglitz, S., Brünker, F., Hofeditz, L., Ross, B., & Frick, N. R. J. (2021). Understanding Collaboration with Virtual Assistants - The Role of Social Identity and the Extended Self. *Business and Information Systems Engineering*, 63(1), 21-37. <https://doi.org/10.1007/s12599-020-00672-x>

Moore, R. K., Li, H., & Liao, S. H. (2016). Progress and Prospects for Spoken Language Technology: What Ordinary People Think. In *INTERSPEECH* (pp. 3007-3011). <https://doi.org/10.21437/Interspeech.2016-874>

Mori, M., MacDorman, K. F., & Kageki, N. (2012). The Uncanny Valley [From the Field]. *IEEE Robotics & Automation*, 19(2), 98–100. <https://doi.org/10.1109/MRA.2012.2192811>

Morrissey, K., & Kirakowski, J. (2013). ‘Realness’ In Chatbots: Establishing Quantifiable Criteria. In *Proceedings of the International conference on human-computer interaction*, 87-96. Springer, Berlin, Heidelberg. https://doi.org/10.1007/978-3-642-39330-3_10

Mou, Y., & Xu, K. (2017). The Media Inequality: Comparing the Initial Human-Human and Human-AI Social Interactions. *Computers in Human Behavior*, 72, 432-440. <https://doi.org/10.1016/j.chb.2017.02.067>

Muller, M. J., & Kogan, S. (2010). Grounded theory method in HCI and CSCW. *Cambridge: IBM Center for Social Software*, 28(2), 1-46.

Nadarzynski, T., Miles, O., Cowie, A., & Ridge, D. (2019). Acceptability of Artificial Intelligence (AI)-Led Chatbot Services in Healthcare: A Mixed-Methods Study. *Digital health*, 5. <https://doi.org/10.1177/2055207619871808>.

Nißen, M., Selimi, D., Janssen, A., Cardona, D. R., Breitner, M. H., Kowatsch, T., & von Wangenheim, F. (2022). See you soon again, chatbot? A design taxonomy to characterize user-chatbot relationships with different time horizons. *Computers in Human Behavior*, 127 (September 2021). <https://doi.org/10.1016/j.chb.2021.107043>

Nordheim, C. B., Følstad, A., & Bjørkli, C. A. (2019). An initial model of trust in chatbots for customer

service—findings from a questionnaire study. *Interacting with Computers*, 31(3), 317–35.

<https://doi.org/10.1093/iwc/iwz022>

Perelman, C., & Olbrechts-Tyteca, L. (2008). *Traité de l'argumentation* [Treatise on argumentation]. Preface by Michel Meyer. Brussels: Éditions de Université libre de Bruxelles.

Pidgeon, N., & Henwood, K. (1996). Grounded Theory: Practical Implementation. In *Handbook of Qualitative Research and Methods for Psychology and the Social Sciences*. Wiley-Blackwell, New Jersey, USA.

Purinton, A., Taft, J. G., Sannon, S., Bazarova, N. N., & Taylor, S. H. (2017). “Alexa Is My New Bff”: Social Roles, User Satisfaction, And Personification of The Amazon Echo. In *Proceedings of the 2017 CHI Conference Extended Abstracts on Human Factors in Computing Systems* (pp. 2853-2859).

<https://doi.org/10.1145/3027063.3053246>

Rapp, A. (2022). Time, Engagement and Video Games: How Game Design Elements Shape the Temporalities of Play in Massively Multiplayer Online Role-Playing Games. *Information Systems Journal*, 32(1), 5–32. <https://doi.org/10.1111/isj.12328>

Rapp, A., Curti, L., Boldi, A. (2021). The human side of human-chatbot interaction: A systematic literature review of ten years of research on text-based chatbots. *International Journal of Human-Computer Studies*, 151, July 2021, 102630. <https://doi.org/10.1016/j.ijhcs.2021.102630>

Rapp, A. Tirassa, M., & Tirabeni, L. (2019). Rethinking Technologies for Behavior Change: A View from the Inside of Human Change. *ACM Transactions on Computer-Human Interaction (TOCHI)*, 26(4), Article 22, 33 pages. <https://doi.org/10.1145/3318142>

Rheu, M., Shin, J. Y., Peng, W., & Huh-Yoo, J. (2021). Systematic review: Trust-building factors and implications for conversational agent design. *International Journal of Human-Computer Interaction*, 37(1), 81-96. <https://doi.org/10.1080/10447318.2020.1807710>

Ruan, S., Jiang, L., Xu, J., Tham, B. J. K., Qiu, Z., Zhu, Y., Burnane, E., Brunskill, E., Landay, J. A. (2019). Quizbot: A dialogue-based adaptive learning system for factual knowledge. In: *Proceedings of the 2019 CHI Conference on Human Factors in Computing Systems*. ACM, New York, pp. 1–13. <https://doi.org/10.1145/3290605.3300587>.

Rueben, M., Horrocks, M. R., Martinez, J. E., Cormier, M. V., LaLone, N., Fraune, M., & Toups Dugas, Z. (2022). “I See You!”: A Design Framework for Interface Cues about Agent Visual Perception from a Thematic Analysis of Videogames. In *Proceedings of the 2022 CHI Conference on Human Factors in Computing Systems (CHI '22)*. Association for Computing Machinery, New York, NY, USA, Article 417, 1–22. <https://doi.org/10.1145/3491102.3517699>

Salem, M., Eyssel, F., Rohlfing, K., Kopp, S., & Joublin, F. (2011). Effects Of Gesture on The Perception of Psychological Anthropomorphism: A Case Study with A Humanoid Robot. In: Mutlu B., Bartneck C., Ham J., Evers V., Kanda T. (eds) *Social Robotics*. ICSR 2011. Lecture Notes in Computer Science, vol 7072. Springer, Berlin, Heidelberg. https://doi.org/10.1007/978-3-642-25504-5_4

Schank, R. & Abelson, R. (1977). *Scripts, Plans Goals and Understanding: An Inquiry Into Human Knowledge Structures*. Hillsdale, NJ: Erlbaum.

Schanke, S., Burtch, G., & Ray, G. (2021). Estimating the impact of “humanizing” customer service chatbots. *Information Systems Research*, 32(3), 736-751. <https://doi.org/10.1287/isre.2021.1015>

Schuetzler, R. M., Grimes, G. M., & Giboney, J. S. (2019). The Effect of Conversational Agent Skill on User Behavior During Deception. *Computers in Human Behavior*, 97, 250-259. <https://doi.org/10.1016/j.chb.2019.03.033>

Schuetzler, R. M., Grimes, G. M., & Scott Giboney, J. (2020). The impact of chatbot conversational skill on engagement and perceived humanness. *Journal of Management Information Systems*, 37(3), 875-900. <https://doi.org/10.1080/07421222.2020.1790204>

Schwartz, T., Stevens, G., Ramirez, L., & Wulf, V. (2013). Uncovering Practices of Making Energy

Consumption Accountable: A Phenomenological Inquiry. *ACM Transactions on Computer-Human Interaction*, 20(2), Article 12, 1-30. <https://doi.org/10.1145/2463579.2463583>

Sharp, V. (1975). *Social control in the therapeutic community*. Brooksville: Saxon House.

Shi, W., Wang, X., Oh, Y. J., Zhang, J., Sahay, S., & Yu, Z. (2020). Effects of Persuasive Dialogues: Testing Bot Identities and Inquiry Strategies. In *Proceedings of the 2020 CHI Conference on Human Factors in Computing Systems*, 1-13. ACM, New York. <https://doi.org/10.1145/3313831.3376843>

Shin, D. (2022). The perception of humanness in conversational journalism: An algorithmic information-processing perspective. *New Media and Society*, 24(12), 2680-2704.

<https://doi.org/10.1177/1461444821993801>

Skjuve, M., Haugstveit, I. M., Følstad, A., & Brandtzaeg, P. B. (2019). Help! Is My Chatbot Falling into The Uncanny Valley? An Empirical Study of User Experience in Human-Chatbot Interaction. *Human Technology*, 15(1), 30-54. <https://doi.org/10.17011/HT/URN.201902201607>

Skjuve, M., Følstad, A., Fostervold, K. I., & Brandtzaeg, P. B. (2022). A longitudinal study of human-chatbot relationships. *International Journal of Human Computer Studies*, 168(August).

<https://doi.org/10.1016/j.ijhcs.2022.102903>

Spicker, P. (2011). Ethical Covert Research. *Sociology*, 45(1), 118–133.

<https://doi.org/10.1177/0038038510387195>

Stein, J. P., & Ohler, P. (2017). Venturing into the uncanny valley of mind—The influence of mind attribution on the acceptance of human-like characters in a virtual reality setting. *Cognition*, 160, 43-50.

<https://doi.org/10.1016/j.cognition.2016.12.010>

Su, N. M., Lazar, A., Bardzell, J., & Bardzell, S. (2019). Of Dolls and Men: Anticipating Sexual Intimacy with Robots. *ACM Transactions on Computer-Human Interaction (TOCHI)*, 26(3), 1-35.

<https://doi.org/10.1145/3301422>

- Svenningsson, M. (2004). Ethics in Internet Ethnography. In E. Buchanan (Ed.), *Readings in Virtual Research Ethics: Issues and Controversies* (pp. 45–61). Hershey: Idea Group.
- Svenningsson, N., & Faraon, M. (2019). Artificial Intelligence in Conversational Agents: A Study of Factors Related to Perceived Humanness in Chatbots. In *Proceedings of the 2019 2nd Artificial Intelligence and Cloud Computing Conference*. ACM, New York, pp. 151–161.
<https://doi.org/10.1145/3375959.3375973>.
- Ta, V., Griffith, C., Boatfield, C., Wang, X., Civitello, M., Bader, H., Decero, E., & Loggarakis, A. (2020). User Experiences of Social Support from Companion Chatbots in Everyday Contexts: Thematic Analysis. *Journal of medical Internet research*, 22(3). <https://doi.org/10.2196/16235>
- Walters, M. L., Syrdal, D. S., Koay, K. L., Dautenhahn, K., & Boekhorst, R. T. (2008). Human Approach Distances to A Mechanical-Looking Robot with Different Robot Voice Styles. In *Proceedings of the 17th IEEE International Symposium on Robot and Human Interactive Communication (RO-MAN 2008)*, pp. 707–712. <https://doi.org/10.1109/ROMAN.2008.4600750>
- Wang, Q., Saha, K., Gregori, E., Joyner, D., & Goel, A. (2021). Towards Mutual Theory of Mind in Human-AI Interaction: How Language Reflects What Students Perceive About a Virtual Teaching Assistant. In *Proceedings of the 2021 CHI Conference on Human Factors in Computing Systems (CHI '21)*. Association for Computing Machinery, New York, NY, USA, Article 384, pp. 1–14.
<https://doi.org/10.1145/3411764.3445645>
- Xiao, Z., Li, T. W., Karahalios, K., & Sundaram, H. (2023). Inform the Uninformed: Improving Online Informed Consent Reading with an AI-Powered Chatbot. In *Proceedings of the 2023 CHI Conference on Human Factors in Computing Systems (CHI '23)*. Association for Computing Machinery, New York, NY, USA, Article 112, 1–17. <https://doi.org/10.1145/3544548.3581252>
- Yeh, S. F., Wu, M. H., Chen, T. Y., Lin, Y. C., Chang, X., Chiang, Y. H., & Chang, Y. J. (2022). How to Guide Task-oriented Chatbot Users, and When: A Mixed-methods Study of Combinations of Chatbot

Guidance Types and Timings. In *Proceedings of the 2022 CHI Conference on Human Factors in Computing Systems (CHI '22)*. Association for Computing Machinery, New York, NY, USA, Article 488, 1–16.

<https://doi.org/10.1145/3491102.3501941>

Zlotowski, J., Strasser, E., & Bartneck, C. (2014). Dimensions of Anthropomorphism: From Humanness to Humanlikeness. In *Proceedings of the 2014 ACM/IEEE International Conference on Human-Robot Interaction (HRI)*. IEE. <https://doi.org/10.1145/2559636.2559679>

About the Authors

Amon Rapp is an associate professor in the Computer Science Department at the University of Torino. His research interests revolve around self-tracking and wearable devices, video games, intelligent agents, and behavior change technologies.

Arianna Boldi is a Ph.D. student in the Psychology Department at the University of Torino. Her research interests revolve around video games, self-tracking and intelligent agents, as well as the use of technology in the context of work.

Lorenzo Curti is a psychologist with an interest in human-computer interaction, artificial intelligence and gender identity.

Alessandro Perrucci is a computer engineer with an interest in machine learning, artificial intelligence, process & service innovation and project management; he is currently employed in the Data Office department of TIM S.p.A.

Rossana Simeoni is a computer scientist with an interest in intelligent interactive systems, human-computer interaction, knowledge-based systems, AI technologies, UX & interaction design, system integration; she is an innovation project manager in the Data Office department of TIM S.p.A. and an adjunct professor in the Cultures, Politics and Society department at the University of Torino.

