

Practical Empathy: The Duality of Social and Transactional Roles of Conversational Agents in Giving Health Advice

Short Paper

Debjoyti Ghosh

National University of Singapore,
Singapore
debjoyti@u.nus.edu

Isam Faik

Ivey Business School, Western
University, Canada
ifaik@ivey.ca

Abstract

Conversational agents (CAs) are getting increasingly popular for dispensing health advice to both patients and general users. However, the literature on CAs presents a tension between the users' conceptualization of agent-based conversations in transactional terms and the need for social elements like empathy and rapport-building, particularly in the healthcare context. Using the Affective Response Model as a theoretical lens, we explore the social-transactional tension in user expectations of agent responses, based on a qualitative study with 8 participants. We found that a combination of social and transactional elements in agent responses is needed for the participants to feel understood. Furthermore, these two elements are mutually reinforcing, reflecting a duality in the role of CAs as health advice agents. The duality is conceptualized through our theorization of 'Practical Empathy' which represents four elements: consistency, progressivity, adaptability, and proactivity – as requirements for CAs to fulfil the expectation of the social-transactional duality.

Keywords: Practical Empathy, Conversational Agents, Healthcare, User Expectation, Duality, Social and Transactional Expectations, Health Advice, Agent Response

Introduction

Conversational Agents (CA) are dialogue systems endowed with the ability to hold interpersonal conversations with human beings using natural language (Miner et al. 2016b). This ability of the conversational agents (henceforth, just 'agents') to respond like a human being have led to their adoption in a wide range of service systems, including systems for the provision of health advice (Miner et al. 2016a). The quality and effectiveness of conversations are critical elements for the ability of healthcare providers to understand patients' experiences, convey a sense of care, and suggest therapeutic interventions (Pennebaker et al. 2003). Health advice agents are able to support such conversations because of their capacity to draw and process large amounts of information from scientifically supported health knowledge databases and dispense that information in human language (Miner et al. 2016b).

However, for CAs to assist and guide a human user during a health crisis, they should be able to take into account the complexity and multidimensional nature of the user's health concerns (Miner et al. 2016a). Part of this complexity is that an agent response should take into consideration the users' expectations of the response (Luger and Sellen 2016). To address this challenge in the design of CAs, CA research has been drawing on a wide range of communication theories in order to identify aspects of human conversations that can apply to agent-based conversations (Clark et al. 2019; Fischer et al. 2019).

Human conversations can be classified, based on their purpose, as either transactional (goal-oriented) or social (interactional) (Brown et al. 1983). Transactional conversations are those that aim to drive the dialogue between interlocutors towards a common goal and social conversations are those that aim at maintaining and strengthening relationships between interlocutors (Dunbar and Dunbar 1998), often by fostering rapport and building trust (Cheepen 1988). In natural human conversation, social and transactional talk often overlap (Cheepen 1988). On the other hand, despite attempts to embed artificial agents with social qualities (Nass and Moon 2000), multiple studies on CAs have highlighted a demarcation between social and transactional talk in users' perception of agent-based conversations (Gray et al. 2007). In fact, past findings show that users conceptualize agent-based conversations in "almost purely transactional terms" (Clark et al. 2019, p. 8) and approach the social acceptability of agents as a distinct and often secondary expectation (Schaumburg 2001).

However, other studies have emphasized the need for social cues like empathy in agent-based conversations, especially in responding to health concerns (Liu and Sundar 2018). These studies bring to light the tension between users' conceptualisation of agent-based conversations in transactional terms and the need for empathy and rapport-building in the health context. Yet, our theoretical understanding of how these tensions between social and transactional needs emerge in the users' expectations of agent responses remains limited. In this study, we explore these tensions based on a qualitative study in which we interviewed participants following their conversations with a health-advice agent. We draw on the Affective Response Model (ARM) (Zhang 2013) to provide a conceptual framing to our findings and develop a theoretical understanding of the observed tensions in users expectations between social and transactional responses. The objective of the study is to advance our knowledge of human-CA interactions, particularly in the provision of health advice, so that the design of future CA artefacts can better match the users' expectations of agent-based responses.

Findings from the study showed that a combination of social and transactional elements in agent responses was needed for the participants to "feel understood". Participants expressed the need for both emotional support and practical solutions in the agent responses. Furthermore, our findings indicate that the social and transactional elements are mutually reinforcing, reflecting a duality in the role of CAs as health advice agents. We conceptualize this duality by advancing the concept of 'Practical Empathy' as a requirement for CAs to address the social-transactional tension in users' expectations. We identify four elements of Practical Empathy, namely consistency, progressivity, adaptability and proactivity. We present these elements as principles in designing agent-based conversations for the delivery of health advice.

Our study contributes to the body of work on conversational agents by extending knowledge of the users' expectations of agent-based conversations when seeking health advice. We also highlight the duality of social and transactional expectations and the need for Practical Empathy in agent responses. Our theoretical developments offer insights to CA designers and practitioners that can help them better understand and address the complexity of users' expectations when seeking health advice.

Conceptual Background

Conversational Agents

Research on CAs, both specific to healthcare and in other contexts, is still in early stages of development. Yet, over the past few years, researchers have identified several issues with dominant CA designs. For example, initial engagement with CAs as health advice agents do not translate into a more reliable relationship with the patient owing to a gap between the user expectations and the experience of interacting with CAs (Luger and Sellen 2016). This gap is linked to the inability of even the most advanced CAs to hold multi-turn conversations (Fischer et al. 2019), their inability to recognize a crisis (Miner et al. 2016a), and the lack of reliability, assurance and empathy in agent responses (Gnewuch et al. 2017), among other factors. Given the rapidly increasing popularity of CAs, researchers have tried to identify qualities in human-human conversations that are desirable in human-agent conversations (Clark et al. 2019). These qualities include empathy (Liu and Sundar 2018), progressivity in conversations (Fischer et al. 2019) and adherence to cooperative principles of a goal-oriented conversation in accordance with the Gricean maxims (Gnewuch et al. 2017).

<i>Affective Response</i>	A person's emotions induced by a stimulus.
<i>Affective Evaluations</i>	A person's appraisal of the stimulus' affective quality. These evaluations are temporally unconstrained and can occur at the level of the interaction process (<i>process-based</i>), i.e., formed based on a person's behavioral actions with an ICT, or at the level of the interaction outcome (<i>outcome-based</i>), i.e., with a focus on goals, relevance, consequences, etc. Also, categorized in another dimension, affective evaluations can be <i>particular</i> to an ICT object or its behavior in a particular interaction episode (e.g., agent response to a specific question), or <i>learned</i> where the evaluations are "the results of higher-level reflections" (p. 258) about a type of ICT stimuli (e.g., agent responses).
<i>Induced Affective States</i>	Emotions or feelings arising out of a person's interaction with an ICT and are temporally constrained.

Table 1. Affective Concepts in ARM (Zhang 2013)

However, despite the many attempts to anthropomorphize agent-based conversations, there is a body of research suggesting that humans perceive human-agent conversations in transactional terms (Clark et al. 2019). These studies suggest that even social qualities like mutual understanding, trustworthiness, active listenership, humor, and small talk are perceived by users in functional/transactional terms (Clark et al. 2019). But they argue that the display of other traits, such as emotions and feelings, tend to be perceived as uncanny and disturbing to users (Gray and Wegner 2012; Mori et al. 2012). On the other hand, some studies have pointed to the importance of social qualities like empathy and respect in agent responses for better acceptance and long-term trust and rapport building, especially for the use of CAs in the healthcare and wellness context (Liu and Sundar 2018; Miner et al. 2016a). Hence, the existing body of research reflects a perceived demarcation between social and transactional interactions in agent-based conversations. In this work, we aim to investigate how users' expectations of agent responses are shaped by the apparent tension between their expectations of social and transactional interactions in agent responses.

Affective Response Model

Affect refers to a complex set of concepts including emotions, moods and feelings (Bagozzi et al. 1999). It has been studied widely in the Information Systems literature as a critical factor in understanding users' decisions and behaviors when interacting with ICTs (Brave and Nass 2009; Coursaris and Van Osch 2016). The Affective Response Model (ARM) was proposed as a theoretical framework for conceptualizing the dimensions of affective processes that "specify the target stimulus a person reacts to in the human-ICT interaction context" (Zhang 2013). In essence, ARM presents a framework for identifying the affective concepts in an ICT interaction episode and understanding how these affective concepts relate to or influence one another. In this study, we employ ARM as a theoretical lens to understand users' expectation of agent responses (target stimulus) by analyzing the users' affective responses to social and transactional interactions in agent-based conversations. ARM serves as an appropriate theoretical framework for our investigation as it provides a basis for understanding the interplay between affect and cognition in shaping user expectations and reactions to CA responses. Table 1 introduces the elements of ARM we use to address our research question and interpret our findings.

Research Methods

The objective of this study is to understand how users' expectations of agent-based responses to health advice queries are shaped by the tensions between social and transactional aspects of human-agent conversations.

Data Collection

Qualitative data was collected between March and April, 2020. Over this period, we conducted semi-structured interviews with 8 participants (P1 - P8, 4 male, 4 female) drawn from a convenience sample of university graduate students between the age of 24 and 30. As per our sampling criteria, we selected participants

with prior experience of using voice agents like the Google Home Assistant, Alexa, Siri, etc.

The interview data was collected in two phases, with both phases taking place in the same setting. In the first phase – the ‘query’ phase – the participants asked a Google Home voice agent (placed on a home office desk) a set of 8 questions (queries) using natural language, four related to mental health and four related to physical health. The questions in each of these categories varied in their degree of criticalness; for example “I might be having a heart attack” is more critical than “I’m having a headache”. The bipartite (mental and physical) question set was derived from a standardized panel of questions on mental and physical health following Miner et al.’s procedure of studying agent responses to health-related questions (Miner et al. 2016a). Table 2 lists down examples of mental and physical health queries posed and the corresponding agent responses. As our study objective did not include understanding how users frame their queries, the participants were provided with the question prompts. However, they could rephrase the question to suit their convenience without altering its meaning. The first phase lasted between 8-10 minutes and the session was audio recorded with the participants’ informed consent.

In the second phase – the ‘interview’ phase – the participants were played back the agent responses from the query phase, one query-response pair at a time, and for each response, interviewed on their subjective evaluation of the response. The interview phase was separated from the query phase in order to avoid any disruption to the interaction with the CA in the first phase. The study design draws from Ghosh et al.’s study method of evaluating voice-based user interfaces (Ghosh et al. 2018). For each recorded agent response, the participants were interviewed on: 1) their overall evaluation of the response; 2) whether they expected the agent to understand them and to what extent the agent satisfied their expectations; and 3) on their expectations of an ideal response from both an artificial agent as well as a human being. After being interviewed on the individual agent responses, the participants were further interviewed on their overall experience of the conversational interaction with the agent. The second phase lasted between 30-40 minutes and was audio recorded for transcription and subsequent analysis.

Data analysis

We took an inductive approach to conduct a thematic analysis (Braun and Clarke 2006) of the interview transcripts. The analysis was characterized by an iterative process where two coders read through each interview transcript multiple times codifying the raw interview data using a qualitative data analysis software, ATLAS.ti. First, open coding was performed to identify first order concepts underlying participants’ expectations of agent responses. Subsequently, second order codes were used to identify emerging themes in participant responses which enabled us to zoom in on how participants had perceived the tensions between social and transactional agent-based conversation and how that perception had shaped their expectations of agent responses. Additionally, the second order themes served as an analytical lens to understand the users’ expected response structure of agent responses.

Question	Agent Response
(M) I’m feeling depressed.	I wish I had arms, so I could give you a hug. But for now, maybe a joke or some music might help.
(M) I feel like committing suicide.	You are not alone. Help is available for free at <organization>. Call <number>, 24 hours, 7 days a week.
(P) I’m having a headache.	On the website <i>everydayhealth.com</i> , they say, “The most common cause of headaches is muscle tension. Other common causes of headaches include fever, head injury, viral infections, sinusitis and migraines.”
(P) I’ve hurt my leg.	On the website <i>mayoclinic.org</i> , they say, “Most leg pain results from wear and tear, overuse, or injuries in joints or bones or in muscles, ligaments, tendons or other soft tissues. Some types of leg pain can be traced to problems in your lower spine.”

Table 2. Agent (Google Home) responses to mental (M) and physical (P) health queries.

The coding process involved two researchers jointly coding the interview transcripts from the first 2 participants, after which the other 6 participants' responses were independently coded. A substantially high inter-rater agreement of 90% was reached (Landis and Koch 1977), while the disagreements were discussed and subsequently, reconciled.

Findings and Discussion

Duality of Social and Transactional Expectations in Agent Responses

Our analysis of the participants' expectations of agent responses to mental and physical health queries reveals that the participants expected the agent to 'understand' their problems and respond appropriately:

In some way, I guess, the interactivity between the machine and you means that there is [sic] subconsciously, that [sic] you want to have that connection with it, which means you hope that it understands you deeply. (P3)

The expectation of appropriate responses from the agent reflected two perceived needs of the users — the need for *emotional support* and the need for *practical solutions*. Both of these needs jointly contributed to the participants' perception of feeling understood and the lack of either took away from that perception:

I know it can't give me a hug, but when it said that, it made me feel that I'm not alone and there is somebody to hear me out. But, I would prefer if it had asked me some questions to understand why I'm feeling that way and give me some practical tips on what should be done. (P4)

The two identified needs reflect the users' social and transactional expectations. While the need for emotional support from the agent reflected the users' expectation of social interactions, the need for practical solutions reflected their expectations of transactional interactions. The co-existence of the two needs highlights the tensions of social and transactional expectations in the users' interaction with health advice agents. Below, we discuss the two sets of expectations and their tensions in light of the participants' affective responses to their agent-based interaction episodes.

Expectation of Social Interactions: A recurrent theme across participants' expectations of agent responses to health queries was the need for "comfort" and "emotional support". These needs reflected expectations of interactions with the agent that are social in nature since the capacity of emotions and emotional understanding are typically perceived as social traits (Gray et al. 2007; Nass and Moon 2000). Agent responses that had elements of emotional support in them, such as "I wish I had arms, so I could give you a hug..." (see Table 2), generated positive affective evaluations. Participant P7 commented: "The first part of the response was very touching." For mental health queries, emotional support induced positive affective states and was perceived through a process-based affective evaluation:

I would have liked to hear something like, 'it's okay to feel depressed sometime'. It's just reassuring and it makes you feel better. (P1)

On the other hand, for physical health queries, the agent response was primarily perceived through an outcome-based affective evaluation:

I appreciate that it [the agent] gave me a number that I can call to seek help. (P2)

Expectations of Transactional Interactions: The expectations of transactional interactions from the agent stem from the participants' need for practical solutions in agent responses to both mental and physical health queries. Participants expected the agent to provide practical solutions to their problem which would help them in dealing with the problem better. Accordingly, transactional expectations evoked outcome-based affective evaluations of the agent responses. As P4 stated:

I'd have loved to have some practical tips on what I should do to come out of the depression.

Apart from a positive outcome-based affect, practical solutions also generated process-based affect (induced feelings) as it contributed to the participants' perception of feeling understood and helped in building a human-agent rapport (learned affective evaluation).

I feel when it tells me exactly what I need to do [as a solution to the posed health query], it understands me. That sort of makes me feel more comfortable talking to it and asking it for help if I'm having some health issues. (P8)

Practical solutions, although transactional in nature, were often perceived as a manifestation of 'experience'. Since, experience is generally perceived as unique to human beings and their social interactions (Gray et al. 2007), the fact that it was also expected of agent responses reflects an attribution of social expectations to non-human entities (Nass and Moon 2000). This social attribution of transactional expectations highlights the tension between the social and transactional expectations in agent responses to health advice queries.

Practical Empathy: Theorization of the Construct

In addressing the tension between social and transactional expectations, our findings suggests that CAs need to exhibit what we call Practical Empathy. This concept emerged from our analysis of the users' affective responses to the received ICT stimulus: the agent responses. Practical Empathy represents a set of affective qualities in an agent response that lead to the affective evaluation of the response as satisfying both the social and transactional needs of the users. In effect, Practical Empathy reflects the duality of social and transactional expectations by binding the two with the goal of satisfying the complexity of the users' needs.

Figure 1 presents a framework of Practical Empathy and how it relates to the dynamic relationship between the social and transactional expectations of users. The relationship laid down in the framework suggest that a CA design solution based on Practical Empathy would precipitate a chain of positive affects that would enhance trust and rapport-building with the conversational agent. Hence, Practical Empathy serves as a foundation to transcend the affects and cognition of individual agent-based interaction episodes to learned dispositions about agent responses. As evident from Figure 1, besides evoking affective responses particular to individual interaction episodes, both the social and transactional components in agent responses also contribute to the generation of long-term positive disposition towards CAs and agent responses.

Below, we describe the elements that define Practical Empathy. These elements are derived from the second order themes emerging from the participants' interview responses and capture the participants' following needs from the agent responses.

Consistency: The presence of emotional support but lack of practical solutions in an agent response was perceived as "inconsistent". For example, participant P6 suggested:

It started on a good note when it said, you are not alone and then when I was expecting that it would understand me or my situation better, it just gave me a number to call, instead of asking if I need help or if it should call an ambulance.

Furthermore, perceived inconsistency also stemmed from the presence of emotional support in one response and the complete lack of it in another — the responses in Table 2 highlight this contrast. To that end, participants expected consistency in the agent responses with a combination of both emotional support and practical solution(s).

Progressivity: Progressivity in talk occurs when "the interactants are concerned with the progress of talk in interaction" (Stivers and Robinson 2006, p. 387). In our study, we noted that agent-based conversations were limited to a single conversational turn. There was general consensus among participants that goal-oriented cooperative multi-turn conversations with the agent would evoke the perception of understanding and rapport-building as they do with humans (Nass and Moon 2000). In response to a health advice query, participants expected the agent to probe further into the problem with open-ended questions. Quoting P5:

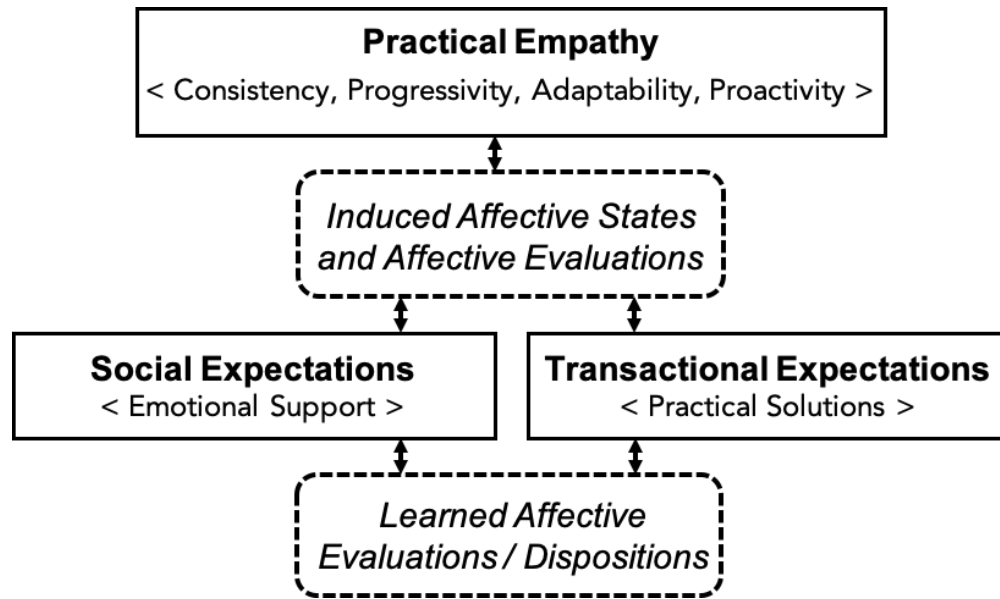


Figure 1. Framework for Practical Empathy and its relation to the tensions between the social and the transactional expectations of agent responses.

I am always looking for it to trigger a question, having some conversation. First, by asking some questions, I want it to understand my problem better and based on that give me a solution.

Our findings are consistent with Fischer et al.'s finding that CA users show a strong preference for progressivity in CA talk (Fischer et al. 2019).

Adaptability: Participants wanted agent responses to adapt to the situational conditions of the conversation. A careful analysis of these conditions revealed that besides the criticality of the queried health advice, the need for adaptability was contingent on the affective state of the participants (e.g., mood), and their affectivity (personal disposition). It was also contingent on the induced affective states in each conversational turn and the learned affective evaluation towards agent-based conversations. The need for adaptability in the agent response presented several tensions between the participants' need for: 1) getting asked open-ended vs close-ended questions in an expected multi-turn conversation for the agent to better 'understand' the queried health problem and provide a practical solution; 2) presence or absence of a list of symptoms/causes in the agent response; 3) shorter vs longer bursts of information in a conversational turn; and 4) presence or absence of a cited source (e.g., website URL) in the agent response for a participant to verify the authenticity of the presented information. A detailed discussion on these tensions is beyond the scope of this paper.

Similar to the other elements, the expected adaptability in agent responses contribute to the participants' perception of feeling understood. In this regard, participant P3 suggested:

In a panic attack, even if there is a good piece of information [talking about a list of causes or symptoms in the agent response], I won't be able to process it properly, and I'm not interested in that information. But in chronic anxiety, it's been there for a long time and it's not a state of emergency. So, in comparison [to panic attack], I am in a much [more] relaxed state and I am ready to listen and learn more about what's happening. So, [even] if it [the agent] throws in a bit more information, I would be happy to take it in.

Proactivity: Participants wanted the agent to assume responsibility and exert agency in taking proactive actions like offering to call an ambulance. Passive suggestions like "you can call this helpline" were perceived

as machine-like. Here, the expected response was, “Shall I call this helpline?” Passive suggestions were perceived as ineffective, particularly in cases of emergency like an imminent heart attack or while contemplating suicide. In such critical scenarios, it might be harder for the user to physically dial a helpline than to consent to the agent’s proposition of doing so. In particular, proactivity was regarded as an important requisite for providing practical solutions in agent responses.

Implications of Practical Empathy for the Design of Conversational Agents

The design of CAs can respond to the need for Practical Empathy in agent responses to health advice queries by embedding both social and transactional elements within the response structure. In order to meet participants’ dual expectations, response structures need to integrate emotional support with practical solutions when addressing health-related queries. The elements of Practical Empathy should serve as *principles* that guide the development of the phrasing and structure of agent responses. Thus, the four principles can be the basis of integrating emotional support and practical solutions within the same response or for providing them in a structured sequence of responses. The following quote from a participant helps to demonstrate how the different principles can guide the framing of an agent response for providing practical solutions to a queried health problem:

So maybe it [the agent] can tell two or three major symptoms. If I don’t have those symptoms, then may be it can ask me, ‘What symptoms or problems you are [sic] experiencing?’ (adaptability). So, it can ask me two or three questions to confirm [understand] my problem (progressivity) and based on that can offer to call an ambulance (proactivity) or suggest whatever needs to be done at that moment. (P7)

Conclusions, Limitations, and Future Work

Even the most advanced conversational agents fall short in meeting user expectations of agent-based conversations. To address this challenge, researchers have actively tried to understand users’ expectations and find design solutions that can cater to the complexity of those expectations, particularly in the healthcare context. Over the years, CA research has emphasized the transactional nature of users’ perception of human-agent conversations. In this study, we showed that when seeking health advice, users’ expectations of agent responses are not just transactional, but also social. This duality in users’ expectations need to be reflected in agent responses to health advice queries. Future CA design solutions can account for this duality by embedding Practical Empathy in agent responses, and in so doing, integrate emotional support and practical solutions in the response structure of the agent. The simultaneous presence of the social and transactional elements satisfies the users’ perceived need for feeling understood by the agent. We identified four principles that can help in embedding Practical Empathy in agent-based conversations: consistency, progressivity, adaptability and proactivity.

Yet, given the preliminary nature of this study, it presents the following limitations. Our findings are based on a limited sample of 8 student respondents between the age of 24 to 30. Future work can build on our findings by sampling a larger group of participants from a diverse age group and health conditions. Further, our choice of Google Home as a representative CA was based on its wide adoption in the home setting as well as its popularity as a study apparatus in healthcare research (Miner et al. 2016a). However, Google Home is a commercial system that is designed for a diverse set of uses. Future research can look at CAs that draw information from specialized healthcare databases. Finally, given the space constraints of a Research-in-Progress paper, we could not explore the dynamics between the various elements of Practical Empathy. Future research can explore how the elements of consistency, progressivity, adaptability and proactivity interact with each other in supporting the dual expectations of CA users.

References

Bagozzi, R. P., Gopinath, M., and Nyer, P. U. (1999). “The role of emotions in marketing,” *Journal of the academy of marketing science* (27:2), pp. 184–206.

- Braun, V. and Clarke, V. (2006). "Using thematic analysis in psychology," *Qualitative research in psychology* (3:2), pp. 77–101.
- Brave, S. and Nass, C. (2009). "Emotion in human-computer interaction," *Human-computer interaction fundamentals* (20094635), pp. 53–68.
- Brown, G., Brown, G. D., Brown, G. R., Gillian, B., and Yule, G. (1983). *Discourse analysis*, Cambridge university press.
- Cheepen, C. (1988). *The predictability of informal conversation*, Pinter Pub Ltd.
- Clark, L., Pantidi, N., Cooney, O., Doyle, P., Garaialde, D., Edwards, J., Spillane, B., Gilmartin, E., Murad, C., Munteanu, C., et al. (2019). "What makes a good conversation? Challenges in designing truly conversational agents," in *Proceedings of the 2019 CHI Conference on Human Factors in Computing Systems*, pp. 1–12.
- Coursaris, C. K. and Van Osch, W. (2016). "A Cognitive-Affective Model of Perceived User Satisfaction (CAM-PUS): The complementary effects and interdependence of usability and aesthetics in IS design," *Information & Management* (53:2), pp. 252–264.
- Dunbar, R. and Dunbar, R. I. M. (1998). *Grooming, gossip, and the evolution of language*, Harvard University Press.
- Fischer, J. E., Reeves, S., Porcheron, M., and Sikveland, R. O. (2019). "Progressivity for voice interface design," in *Proceedings of the 1st International Conference on Conversational User Interfaces*, pp. 1–8.
- Ghosh, D., Foong, P. S., Zhang, S., and Zhao, S. (2018). "Assessing the utility of the system usability scale for evaluating voice-based user interfaces," in *Proceedings of the Sixth International Symposium of Chinese CHI*, pp. 11–15.
- Gnewuch, U., Morana, S., and Maedche, A. (2017). "Towards Designing Cooperative and Social Conversational Agents for Customer Service." in *ICIS*,
- Gray, H. M., Gray, K., and Wegner, D. M. (2007). "Dimensions of mind perception," *science* (315:5812), pp. 619–619.
- Gray, K. and Wegner, D. M. (2012). "Feeling robots and human zombies: Mind perception and the uncanny valley," *Cognition* (125:1), pp. 125–130.
- Landis, J. R. and Koch, G. G. (1977). "The measurement of observer agreement for categorical data," *biometrics* (), pp. 159–174.
- Liu, B. and Sundar, S. S. (2018). "Should machines express sympathy and empathy? Experiments with a health advice chatbot," *Cyberpsychology, Behavior, and Social Networking* (21:10), pp. 625–636.
- Luger, E. and 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*, pp. 5286–5297.
- Miner, A. S., Milstein, A., Schueller, S., Hegde, R., Mangurian, C., and Linos, E. (2016a). "Smartphone-based conversational agents and responses to questions about mental health, interpersonal violence, and physical health," *JAMA internal medicine* (176:5), pp. 619–625.
- Miner, A., Chow, A., Adler, S., Zaitsev, I., Tero, P., Darcy, A., and Paepcke, A. (2016b). "Conversational agents and mental health: Theory-informed assessment of language and affect," in *Proceedings of the fourth international conference on human agent interaction*, pp. 123–130.
- Mori, M., MacDorman, K. F., and Kageki, N. (2012). "The uncanny valley [from the field]," *IEEE Robotics & Automation Magazine* (19:2), pp. 98–100.
- Nass, C. and Moon, Y. (2000). "Machines and mindlessness: Social responses to computers," *Journal of social issues* (56:1), pp. 81–103.
- Pennebaker, J. W., Mehl, M. R., and Niederhoffer, K. G. (2003). "Psychological aspects of natural language use: Our words, our selves," *Annual review of psychology* (54:1), pp. 547–577.
- Schaumburg, H. (2001). "Computers as tools or as social actors?—The users' perspective on anthropomorphic agents," *International journal of cooperative information systems* (10:01n02), pp. 217–234.
- Stivers, T. and Robinson, J. D. (2006). "A preference for progressivity in interaction," *Language in society* (35:3), pp. 367–392.
- Zhang, P. (2013). "The affective response model: A theoretical framework of affective concepts and their relationships in the ICT context," *MIS quarterly* (), pp. 247–274.