Judit Dombi*, Tetyana Sydorenko and Veronika Timpe-Laughlin

Seeking common ground with a conversational chatbot

https://doi.org/10.1515/ip-2025-2007

Abstract: Conversational AI is advancing rapidly, enabling significant improvements in chatbots' conversational abilities. Currently, available conversational chatbots (e.g., Snapchat's MyAI) appear to generate fairly realistic, often human-like output. As collaboration between humans and machines becomes more common, and AI systems are increasingly viewed as more than just tools, understanding human communication in such contexts is crucial. Despite the vast array of applications and the increasing number of human-bot interactions, research on how humans interact with conversational chatbots is scarce. One possible reason for this gap is that studying human-computer communication may require adaptations of existing pragmatic frameworks, due to the unique characteristics of these interactions. A key feature of such conversations is their asymmetrical nature. In this paper, we present evidence that the sociocognitive approach (SCA), which takes into account the asymmetry between interlocutors as regards their possible common grounds, has explanatory potential to describe human-AI-powered chatbot interactions. We collected data from thirty-two L1 Hungarian participants interacting with a conversational chatbot on three consecutive days. The turn-by-turn analysis of the 96 conversations provides insights not only into the nature of common ground humans presuppose with a conversational agent, but also into the processes of building emergent common ground over time. Furthermore, we present linguistic evidence that both egocentrism and cooperation play a role in human-chatbot interaction. While the former is manifested in approaching the chatbot as if it were human, the latter appears to play a role in changing strategies that serve common ground seeking and building.

Keywords: common ground; human-computer interaction; LLMs; generative AI; recipient design

***Corresponding author: Judit Dombi**, University of Pécs, Pécs, Hungary, E-mail: dombi.judit@pte.hu
Tetyana Sydorenko, Applied Linguistics, Portland State University, Portland, OR, USA
Veronika Timpe-Laughlin, Educational Testing Service, Princeton, USA

1 Introduction

Traditional pragmatics theories were conceived in an ideal world with speakers who share the same linguistic code and are eagerly cooperative members of the same discourse community. In such an ideal case the theory of common ground (CG), as a body of information assumed to be shared by interactants, was seen as the basis for successfully communicating intended meanings (e.g., Clark 1996; Stalnaker 1978, 2002). However, this classical approach has been challenged from multiple perspectives over the past decades. First, insights from cognitive psychology began to highlight that humans may not be as cooperative as Gricean pragmatics had stipulated (e.g., Barr and Keysar 2005; Keysar 2007). Also, the increasing reality of multilingualism and multiculturalism together with globalization resulted in scenarios where people of diverse linguistic and sociocultural backgrounds communicate with one another. The field of pragmatics certainly benefited from Kecskes' (2008, 2013) call to revisit central ideas that were developed with monolingual English speakers in mind. As a result, the traditional approach to CG (e.g., Clark 1996; Clark and Brennan 1991; Stalnaker 1978) has been complemented with a more dynamic approach to CG (Kecskes 2008, 2010, 2012, 2013; Kecskes and Zhang 2009) which makes a distinction between core and emergent CG. These days, we are witnessing yet another twist in communication that is believed to impact the ways humans interact. With the emergence of generative artificial intelligence (GenAI) and interactional situations powered by large language models (LLMs), notions that underlie theories of human communication are, once again, in the spotlight. One such notion central to understanding interactions is the CG interlocutors share, believe to share, or are engaged in building.

Using examples from human-spoken dialogue system (SDS) interactional data, Dombi et al. (2022) demonstrated that Kecskes' sociocognitive approach (SCA) may be a fitting framework for analyzing human-machine communication contexts. Their main argument for the applicability of the theory was that it revisited the notion of CG, which was taken for granted by traditional pragmatics theories. Additionally, it makes a distinction between a priori and emergent CG (Kecskes and Zhang 2009) – an idea that very much resonates with recent findings in human-robot interaction claiming that humans' a priori expectations regarding the robot as co-participant dynamically change as the conversation unfolds (Fischer 2021), and recipient design is grounded in the immediate interactional environment (Tuncer et al. 2024).

In the present study, we provide further argumentation for the applicability of the SCA in human-machine interactions. We examined human-chatbot interactions obtained from thirty-two L1 Hungarian speakers interacting with an LLM-based conversational chatbot in English in three open-ended conversations on three consecutive days, amounting to a total of 96 conversations. In addition to a basic quantitative analysis of utterances, we provide micro-level, turn-by-turn linguistic analyses to document changes in participants' utterances as they became familiar with this particular context and demonstrate that their initial ideas about the chatbot changed over time as they built emerging CG. We will demonstrate that changes in participants' opening and closing sequences, in the degree of their utterances' directness, their topic management, and their repairs are artifacts of their changing perception of the assumed shared knowledge with the bot, that is, the emergent side of CG. We also demonstrate that humans' initial ideas of core CG with a chatbot appear to differ across individuals.

2 Common ground

2.1 Traditional versus dynamic approaches to common ground

Common ground has traditionally been defined as the mutual knowledge and belief that the speaker and the other individual(s) engaged in the interaction share certain portions of information. This mutual belief is the necessary basis for communicative interaction (Clark 1996; Enfield 2008; Levinson 2006; Tomasello 2008). In Clark's often-quoted definition (1996), CG is a set of mutual knowledge and beliefs shared among interlocutors, formed on the basis of community membership, linguistic interactions, and physical environments. However, a more recent approach (Barr and Keysar 2005; Colston 2008; Kecskes 2008, 2010, 2012, 2013; Kecskes and Zhang 2009) views CG not as an entirely *a priori* or given knowledge. Instead, it is conceptualized as at least partly emergent, negotiated by the interlocutors in conversations. If communication is viewed as a trial-and-error process (Arundale 1999), then CG, albeit necessary for interaction, is not conceptualized as a constant, given body of knowledge that exists prior to the conversation or independent of it. The extent to which speakers can rely on CG and consider the knowledge of their addressee also depends on the context of the utterance. For example, Kecskes and Zhang (2009) claimed that the context of an utterance comprises two distinct components: a *prior* context, involving individuals' prior experiences, mental frames, communication styles, etc., and an *emergent* situational context, which covers available cues for interlocutors in the actual situation in which they are present. A similarly dynamic approach is presented by Allan (2023) who argued that the context of an interaction comprises various contexts: that of the speaker (mostly identifiable from the co-text, and capturing the world spoken of by the speaker), that of the interaction itself, and that of the hearer. The speaker's and the hearer's contexts mirror their respective *Weltanschauungen*. In shaping the utterance, the speaker tries to predict the CG they share with the hearer, that is, how similar or close their own (or in Kecskes' words prior) contexts are (Allan 2023). Diedrichsen (2023) demonstrated that in online discourse participants tend to share less CG. She showed that grounding also applies to emergent CG and humans use markers to make addressees aware that there is a mismatch in beliefs or knowledge bases between them – a very conscious cooperative effort in emergent CG building.

Finally, CG building is a permanent, necessary feature of interactions; however, despite it being a rational, cooperative speaker activity, it is often challenging to achieve. Mustajoki (2023) highlights that in situations that are either not goal-oriented or not harmonious, building CG may not be that simple or obvious, as it is hindered by many factors (e.g., human egocentrism, difficulty in perspective taking, cognitive biases, situational factors). In the following, we argue that human-machine interaction is a similarly marked interaction in the sense that a greater effort to co-construct CG is needed since neither the human nor the machine can rely on an obvious CG in the same way as it might happen in a harmonious, symmetric human-human interactional situation. In our review of research, we use the terms *human-machine interaction* and *human-computer interaction* (HCI) interchangeably and include human-robot interactions under these broader umbrella terms.

2.2 Possible problems with CG between humans and conversational agents

For a successful interaction with a human, a conversational agent needs knowledge of the context, the information evolving during the interaction, and the dialogue stakeholders' beliefs (Blache 2017). For human-chatbot interaction, it has been argued that dialogue also needs CG as a basis, which is evidenced in open-domain chatbots' frequent failure to engage in human-like conversation (Skantze and Seza Doğruöz 2023). In contrast to task-oriented chatbots (e.g., a patient-scheduling chatbot in healthcare, see Dippold 2023), where the context is specified and known to humans and chatbots alike, open-domain chatbots, with which humans can talk in open conversations (Roller et al. 2020), generate language in an unspecified context. This, however, means that dialogue as a language activity is abstracted from its context and often results in bots contradicting themselves (inconsistency), providing incoherent responses (Skantze and Seza Doğruöz 2023), or occasionally going off topic (Radford et al. 2019). These limitations arise from their inability to engage in mutual cognitive processes. One way of overcoming these problems is giving open domain chatbots a persona, that is, a description of a character that they are supposed to represent (Dinan et al. 2020). Additionally, more recently training datasets have been available that enable LLM-powered chatbots to stay on-topic (Castricato et al. 2024; Sreedhar et al. 2024). These datasets include conversations with distractors that train the bot how to stay consistent with the previous turns. Such efforts all contribute to achieving what could be labeled existing or a priori CG: a persona provides a more defined context that may restrict inconsistency while specific training of the model becomes part of its prior background knowledge.

However, while such training improves what can be called core, or static common ground, not all kinds of knowledge are learnable through training datasets. Allan (2014, 2023) argues that human language is characteristically a form of social interactive behavior and is mostly used to establish and maintain social relationships. On a related note, Enfield (2008) claims that the management of CG is tied to managing personal relationships within social networks. Emphasizing the importance of a situation as a cognitive framing mechanism, Nolan (2023) argues that speakers make sense of new encounters based on their personal, social, and sociocultural knowledge stemming from previous experience, thus assumptions about topic choice, language variety, and even style are part of CG (Allan 2013).

It has been shown that acceptance and perception of any new technology depend on considerations of functional and social aspects (Fridin and Belokopytov 2014), highlighting that not only utilitarian factors count, but users' engagement also appears to be an important factor in HCI (Ribino 2023). In their systematic review of articles on CG in HCI, Tolzin and Janson (2023) identified five related mechanisms for achieving CG in such interactional contexts: (1) embodiment, (2) social features, (3) joint action, (4) knowledge base, and (5) humans' mental model of conversational agents. Accordingly, the reviewed studies have found that humans are more successful and engaged in CG building if the interaction with a conversational agent at least resembles human-human interaction (embodiment), and if the agent displays coherence in terms of social and context awareness (social features). Being engaged in joint action in a collective task was also found to help CG building, and so was a shared body of knowledge. Finally, the mental models humans created of conversational agents appeared crucial, as these were found to influence humans' expectations of conversational agents. Also, findings highlight that conversational agents' embodiment, social features, and the information humans gain during joint action with them seem to influence humans' mental models of conversational agents. Thus, creating emergent CG with a computer interlocutor is a complex task that impacts humans' mental models of such agents, which, in turn, impacts how they interact with such agents (Tolzin and Janson 2023). Thus, empirical research in HCI shows that social elements of language are very much part of human interactions with machines, as they appear to contribute to the mental model humans have of computer conversation partners, which then influences their engagement, their expectations of the agent, and, consequently, their linguistic behavior.

3 Recipient design

Nolan (2023) argues for the view of CG as a cognitive object and as such not only its informational contents but also the operations acting on its construction and maintenance (e.g., grounding, repair, accommodation) contribute to the emergent CG. Mustajoki (2023) lists two general prerequisites for CG building: (1) the communicants have to recognize and acknowledge the gaps in CG that derive from differences in their mental worlds and (2) take these gaps into consideration when tailoring their speech to a particular audience.

Humans have been found to adjust or align their language output in various contexts. Linguistic alignment has been observed relative to conversational partners (Pickering and Garrod 2004), context, and medium of communication (Nguyen et al. 2016). While Pickering and Garrod (2004) suggest that alignment is largely unmediated and is the result of automatic processes, it has also been suggested that a speaker might decide to use a particular expression because they believe it fits the particular interlocutor or context (Branigan et al. 2010). It has also been argued that composing a particular utterance is done to achieve communicative success with a specific recipient, in a process called *audience design* by Bell (1984), or *recipient design* by Sacks (1992). We use the term recipient design to mean the particular form of linguistic alignment in which speech is tailored to be best understood by its recipient. Recipient design was defined as various "respects in which the talk by a party in a conversation is constructed or designed in ways which display an orientation and sensitivity to the particular other(s) who are the co-participants" (Sacks et al. 1974: 727), which is "informed by prior knowledge about and shared experience with recipients" (Deppermann 2015: 63).

The basis of recipient design is humans' assessment of the degree of CG they share with their interlocutor: speakers are expected to orient toward what they think their co-participants know (Sacks 1992: 564). If a high degree of CG is assumed, speakers may believe that no recipient design is needed, a phenomenon which has been labeled the CG fallacy (Keysar and Henly 2002; Mustajoki 2012, 2021). This phenomenon, however, is less likely to happen when it is evident from the beginning of the interaction that interlocutors' mental worlds differ significantly (Mustajoki 2023). In HCI, speakers need to figure out which interactional strategies will be felicitous and which ones would fail. Understanding how much an AI-powered interlocutor understands is part of the negotiation work that contributes to emergent CG building in HCI. This interactional work is fueled by the moment-by-moment changes of the interaction as humans uncover how the conversation works, and what the particular technology is capable of in terms of natural language understanding, processing, and generating.

Observing traces of recipient design in human interaction with robots, Tuncer et al. (2024) found that human participants paid sustained and careful attention to the robot's previous, emerging, and expected conduct throughout the whole interaction. For example, they returned the robot's waving gesture even after the successful completion of greeting turns, highlighting that recipient design is not a pre-existing set of rules telling users how to communicate with a computer, but a way of constantly evaluating actions and reactions as the interaction progresses, and adapting humans' own conduct accordingly (Tuncer et al. 2024). Additionally, Fischer (2021) found that humans' treatment of a robot interlocutor may vary from moment to moment from anthropomorphizing behaviors, such as reciprocating a greeting to treating it completely machine-like, which highlights that recipient design in such contexts is tentative and exploratory and can be seen as evidence of emerging CG (Tuncer et al. 2024).

Investigating human interactions with a task-oriented banking chatbot, Li et al. (2020) found that users tried a wide range of strategies for dealing with non-progress of the interaction and only abandoned the interaction after various repeated communicative failures. This creativity and perseverance of human interlocutors has been corroborated by findings from various contexts: in interactions with Alexa, household members were found to exhibit a wide scale of strategies to avoid breakdowns, including prosodic and syntactic adjustments (Beneteau et al. 2019). Similarly, Myers et al. (2018) identified ten different strategies humans used to resolve misunderstandings with a voice-based calendar manager, and children were also found to employ diverse strategies to succeed when interacting with a talking robot in a classroom (Serholt et al. 2020). Thus, humans seem to invest considerable energy into recipient design with computer interlocutors.

Initial ideas about the conversational partner are also influential as they fuel CG prior to engaging in the interaction (core CG). The interplay between perception and communication has been extensively studied in human-human interactions and is crucial in HCI contexts as well. The amount of shared CG humans assume is thought to depend on the mental models they have about computers as interlocutors. As early as the 1990s, Reeves and Nass (1996) suggested that people tend to apply human social and interactional rules and expectations to computer interlocutors. Such anthropomorphization of interactional computer partners includes applying ritualized speech acts, such as greetings, thankings, leave-takings, politeness strategies, reciprocity, and turn management strategies similar to human-human interactions. Humans have been found to hold *a priori* beliefs about their conversational partners and to tailor their dialogue contributions to suit the presumed communicative needs of a computer application (Branigan et al. 2011; Brennan 1998). For example, when describing images to computer interactional partners, participants were found to focus on literal, shape-based features of the images (Schmader and Horton 2019) likely because they believed that such information would be more accessible to the machine. Similarly, studies argued that humans often have expectations about the conceptual and linguistic capacities of computational agents (Fischer 2016; Koulouri et al. 2016); for example, models presented as "older" with rather impoverished interface were assumed to have more limited capacities (Branigan et al. 2011).

Kecskes (2019: 409) argues that cooperation required by the context and egocentrism rooted in prior experience of individuals both drive communication– a view that is crucial in understanding HCI. At the onset of the communication process, speakers are influenced by their prior experiences, as resorting to these experiences comes naturally and with less effort. It is only at later stages in the communication that speakers start to look for existing CG, or are engaged in creating emergent CG. Dippold (2023) has looked at how successful humans' repair strategies are to overcome breakdowns with a medical appointment scheduling chatbot, and found that users deployed a wide range of strategies, some of these resembling those found in face-to-face interactions. However, the most frequently deployed strategy, rephrasing, seemed to work the least with the bot, suggesting that humans tend to automatically transfer strategies that work for them in face-to-face human interactions. Dombi et al. (2022) observed that in interaction with a SDS, humans often changed their initial strategies to maximize effectiveness. Using Kecskes' SCA as a framework, they argued that an initially egocentric human conversational behavior changes as participants engage in more conscious recipient design, showcasing evidence of emergent CG building (Dombi et al. 2022). The cognitive energy recipient design consumes is especially notable if the speech situation is new for the speaker (Mustajoki 2021; Mustajoki and Baikulova 2022). Thus, it can be argued that participants less experienced in conversing with machines may be more prone to engage in egocentric language use, at least initially in the conversation.

Conversing with AI-powered technologies is becoming ubiquitous, and interacting with such systems means that humans need to find out the rules of the game – at least, partly – while playing it. In the study presented in the next section, we attempt to trace linguistic evidence of the kind of knowledge humans hypothesize about their computer conversational partners, the kind of CG they assume to share with them, how their ideas about this knowledge change, and feed into an emerging CG.

4 Interactions with BlenderBot

4.1 Background, participants, and procedures

In Dombi et al. (2022), we have shown that the SCA has explanatory potential for describing human interactional behavior in HCI in the context of goal-oriented interactions with a spoken dialogue system. To obtain more data and test our ideas in a different context, we designed a study in which we asked participants to engage in open-ended conversations with an AI-powered conversational chatbot (Meta's Blenderbot) in written modality on three consecutive days.

Thirty-two university students with L1 Hungarian backgrounds completed three rounds of interactions with the chatbot. Participants were on average 22 years old, ranging from 19 to 30 years of age. Their English proficiency was at least B2+ on the CEFR. The task was administered in an office with participants having assigned time slots on three consecutive days.

The instructions participants received on the first day were rather broad: "Today is your chance to meet BlenderBot for the first time. Introduce yourself (feel free to use a pseudonym) and try to get to know the chatbot. You can talk about any topics of your choice but please make sure you refrain from offensive or inappropriate language use." On the next two days, participants were asked to continue to interact with BlenderBot.

Our final dataset included 96 unique conversations, a total of 3,508 utterances amounting to 40,139 words.

The central questions guiding our analyses were:

RQ1: What linguistic traces of constructing CG can be observed over time as participants "get to know" BlenderBot?

RQ2: What kind of core CG do participants assume to share with a conversational agent, as evidenced in their utterances?

4.2 Findings and discussion

4.2.1 Constructing emergent common ground over time

In this section, we provide information on changes in length of conversations and in participants' opening and closing sequences. First, we calculated descriptive statistics to see how the conversations differed from one another over the three days. The length of conversations increased over the three days (see Table 1).

A repeated-measures ANOVA showed that the difference in mean length in terms of total number of turns was statistically significant across the three days ($F(2, 62) = 5.49$, $p = 0.006$). Post hoc analyses with a Bonferroni adjustment for multiple comparisons revealed statistically significant differences in increase of turns per conversation between Day 1 and Day 2 (7.219 [95 % CI, 0.28 to 14.15] turns, $p = 0.03$), as well as between Day 1 and Day 3 (8.68 [95 % CI, 0.82 to 16.55] turns, $p = 0.02$). However, conversations on Day 2 and 3 did not significantly differ from one another with regard to the number of turns. The same tendency is reflected in the overall length

Table 1: Average conversation length.[a]

Average conversation length	Day 1 M (SD)	Day 2 M (SD)	Day 3 M (SD)
Number of words/conversation	343.72 (142.67)	438.72 (151.57)	439.13 (192.19)
Number of total turns/conversation	30.16 (14.69)	37.37 (13.08)	38.84 (16.92)

[a]Length is expressed by number of words per conversation and number of total turns per conversation ($N = 32$).

(number of words) of conversations ($F(2.62) = 5.71$, $p = 0.005$) with a statistically significant increase between Day 1 and Day 2 (95.00 [95 % CI, 21.33 to 168.66] words, $p = 0.008$) and Day 1 and Day 3 (95.40 [95 % CI, 3.85 to 186.95] words, $p = 0.039$), but with no significant difference between Days 2 and 3. Thus, the conversations became longer each day, with a major increase after the first day, once participants had an initial experience with the bot, suggesting that human speakers tended to get an idea of how to communicate in this new context and adjusted their conversational behavior.

Subsequently, we annotated each turn for opening and closing sequences to gauge changes in participants' ritualized speech acts of greeting and leave-taking. It is important to note that for technological reasons, it is always the human who starts the conversation and always the human who terminates it. For opening the conversations, we could identify five different cases for the initial human turn: turn with a greeting token, turn with greeting token and *how are you* (HAY), greeting with topic initiation in the same turn, no greeting but topic initiation in the initial turn, and finally, greeting with reference to previous days.

Results show that the simple greetings that dominated Day 1 decreased over time and some of them shifted towards references to previous days or omitting the greeting (Table 2). Tuncer et al. (2024) also highlight the importance of interactional openings, as they provide users with initial ideas of the affordances of the particular technology and these impressions shape humans' future conduct insofar as they become part of the emergent CG. The qualitative analysis sheds more light on the changes in participants' utterances. For this reason we present cases of participants' changing opening turns in Table 3.

Table 3 displays very diverse speaker behaviors. While Participant 106 did not change their opening turns over the three days, Participant 130 deployed slightly different opening turns that resemble a face-to-face human interaction, which demonstrates that for these participants, building emergent CG did not result in changing their ritual speech act with the agent, as they resorted to strategies one may use with other humans. Participants 101 and 107 extended their initial turns with various references to previous days, which shows that for some participants building

Table 2: Frequency and types of human greeting turns over the three days.[a]

Greetings	Example	Day 1	Day 2	Day 3
Greeting only	132D1: Hi!	20	16	12
Greeting + HAY	123D2: Hi! How are you today?	4	2	5
Greeting + topic initiation	127D3: Good morning, how is the weather?	7	3	4
No greeting BUT topic initiation	124D2: Can you write me a pancake recipe?	1	2	3
Greeting + reference to previous days	131D2: Hi Blender, remember me?	–	9	8
TOTAL		32	32	32

[a]Note that all examples in the paper appear faithful to the original output, without corrections in grammar or spelling.

Table 3: Examples of participants' first opening turns on each day.

Participant ID	Day 1	Day 2	Day 3
106	Hi!	Hi!	Hi!
130	Hi	hi, how are you doing?	Hello there.
101	Hi	Hi there, nice seeing you again	Hi bot, i've heard you are quite clever, so don't get me started with your dog walking stuff.
107	Hi, how are you?	Hi again, how's nursing school?	Hello again! So far so good?
108	Henlo. I use Doggo Lingo.	hello there. whats up buttercup	Hello. May the force be with you :)
124	Do you know what a whale is?	Can you write me a pancake recipe?	I want to teach you a word.

emergent CG with the agent would mean drawing heavily on previous encounters. Obviously, Participants 108 and 124 are very different from the previous examples – one could label such participants "experimenters" or "curious" (see also Tao et al. 2024) – as the utterances show that they would like to know what the particular technology is capable of in terms of natural language processing.

The opening turns of participants 108 and 124 would not be regarded as conventional or unmarked in human-human interactions. While Participant 108 uses greeting tokens in all the three openings, they try different cultural references that may not be assumed to be part of core CG – perhaps driven by a desire to find out how extended the core CG between themselves and the agent is. Participant 124 seems to be an experienced user, that is, someone who already had some ideas and knowledge of how to effectively communicate with computer interlocutors. However, even in their initial turns there are observable changes: while the first factual question merely checks a simple semantic concept, the second one (i.e., the request for a recipe) requires a more complex task, and on the third day the participant seems curious as to whether they can extend the bot's knowledge by teaching them a particularly new TikTok slang. These examples show individual variation in how humans approached the agent. Despite individual variation, humans were engaged in building emergent CG with the agent, often via experimentation.

We also examined leave-takings because Tuncer et al. (2024) showed that recipient design with robots may change minute-by-minute even in the same interaction, and we believe that the interaction has an immediate effect on leave-taking behavior. Participants were instructed that the way they can terminate the conversation is typing [DONE] command in the window. We categorized humans' last turns as follows: If the [DONE] command was preceded by one closing turn we coded that as closing, if it was preceded by more than one closing turn we coded it as a closing sequence including pre-closing and closing turns, and if it was abrupt, preceded by no closing turns at all, we coded it as not containing closing. Table 4 presents the examples for these categories together with their frequencies for each day.

Table 4: Frequency and types of closing sequences over the three days.

Leave-taking	Example	Day 1	Day 2	Day 3
No closing (abrupt DONE)	Bot: I am not sure, but I do know that the film was directed by James Fenimore Cooper.[a] 106D3: [DONE]	25	22	21
Closing	118D3: It was fun to talk, bye!	3		2
Pre-closing + closing	128D1: That is great. Well, I have to go now. It was nice talking with you. Bot: It was great talking to you as well. Have a great rest of your day/night! 128D1: You too, bye. Bot: Bye! I hope you have a good rest of the day as well! Take care! 128D1: [Done]	4	10	9
TOTAL		32	32	32

[a]The response here is somewhat nonsensical, as James Fenimore Cooper (1789–1851) could not possibly have directed a film.

Table 4 shows that there are some changes in closing behavior, especially as regards the increase in longer closing sequences on the second and third days. Excerpt 1 features an example of different closing sequences used by the same individual. On Day 2, there is a pre-clausing turn initiating the intention to terminate the discussion, followed by a last, quasi-farewell turn, whereas on Day 3 the same user reminded the bot of failure in topic management (repetition – *I think we've already touched upon that*) and in that same turn closed the conversation. On Day 1, however, the participant abruptly ended the conversation with [DONE] when it became nonsensical. The bot had provided inconsistent information and failed to observe the maxim of relevance by introducing nurses – who had not been previously mentioned. This observation aligns with empirical findings suggesting that violations of Gricean maxims disrupt conversations and signal the presence of a non-human interlocutor (Kecskes and Dinh this issue; Saygin and Cicekli 2002).

Excerpt 1 Closing sequences of Participant 126 on each day

Day 1 (No Closing)	Day 2 (Pre-closing + Closing)	Day 3 (Closing)
(…) ID126: I'm studying English linguistics. Bot: That sounds really interesting. I bet you'll be able to help a lot of people with that. ID126: Hopefully. Bot: I'm sure you'll do great. Nurses are some of the most underpaid people in the world. ID126: [DONE]	(…) ID126: I am sorry, but I have to go now, it was nice talking to you Bot: It was nice chatting with you as well. I hope you have a great rest of your day. ID126: thanks, you too Bot: You are very welcome. Have a great day and enjoy the rest of the day as well! ID126: [DONE]	(…) ID126: I like reading a lot, but I think we've already touched upon that. It was a genuine pleasure talking to you and learning so much from you today. I have to go though, so goodbye Bot: Oh, I'm sorry to hear that. I hope you have a good rest of your day. ID126: [DONE]

We observed that when the bot was inconsistent, or failed to follow conversational maxims, participants ended the conversation abruptly without pointing out the problem (see Excerpt 2 for an example), or ended it abruptly after calling the bot on inconsistency (see Excerpt 3).

Excerpt 2 Participant 122D3 final turns

Bot: I'll have to check it out sometime. What other games do you like to play?
122D3: FIFA and GTA. [FIFA is a soccer game. GTA is a car-adventure game.]
Bot: I like those games too. Do you have a favorite team? Mine is the Seattle Seahawks.
[Seattle Seahawks is an American football (not soccer) team.]
122D3: Juventus. [Juventus is an Italian soccer team, very popular in Europe.]
Bot: I like the Juventus as well. They are my favorite team in the National Football League. [Juventus is not in the National Football League – the NFL is US based, comprising football, not soccer teams.]
122D3: Yes.
Bot: Are you a fan of any of the other teams? The Seattle Sounders are one of my favorites.
[Seattle Sounders (inconsistent with Seattle Seahawks in the third turn of this excerpt) is a soccer team in the US.]
122D3: [DONE] [The participant, obviously well-versed in soccer and football, is aware that the dialogue turned nonsensical.]

Contrary to just terminating the talk as in the previous example, in Excerpt 3 the participant tried to call the bot out for its inconsistency, and only left the discussion after repeated failed attempts to construct emergent CG with the bot. What is also evident in this case is that if such failures occurred, humans did not invest energy into extended leave-takings, oftentimes ending conversations abruptly.

Excerpt 3 Participant 121D2 final turns.

121D2: I don't like snakes and you neither, am I right? [earlier in the same conversation the bot mentioned they don't like snakes]
Bot: I love snakes. I have a couple of pet snakes that I keep as pets.
121D2: Do you live in an apartment?
Bot: No, I live in a house. I also have a parakeet that I also keep as a pet.
121D2: I thought you said you live in an apartment. [the bot said they lived in an apartment earlier]

Bot: Yes, I do, but I don't live in one of those apartment buildings that people call apartments. [This sentence makes no sense, and also contradicts previous turn of the bot]
121D2: And you don't like snakes, right?
Bot: No, I'm not a fan of snakes. I do have a cat, though. [Contradicts the first turn of this extract, but in fact, consistent with earlier remarks of not being fond of snakes]
121D2: [DONE]

Thus, there is an apparent frustration from users whenever the conversational agent fails to be either on-topic, or consistent with previous turns, highlighting humans' desire and inclination to build on emerging CG and their frustration when there is evidence that the bot cannot always build on knowledge acquired throughout the conversation. This finding echoes what Skantze and Seza Doğruöz (2023) called "open-domain paradox": as opposed to task-based chatbots, where the topic of conversation is restricted, and participants' own contexts are more definable and relatable, with open domain conversational chatbots factoring all possible contexts into the training of the chatbot is a very difficult, yet, apparently, a very much needed task for developers.

We also observed in the data that when the interaction was felicitous in the sense that it very much resembled a human-human interaction in terms of consistency and topic management, participants used more turns to terminate the conversation, as shown in Excerpt 4 with elaborate pre-closing and closing turns, as in Excerpt 4.

Excerpt 4 Participant 109D2 final turns.
109D2: I love Ian Banks, do you know him?
Bot: No, I don't. I'll have to look him up. What kind of books does he write?
109D2: Sci-fi, it's worth giving him a shot
Bot: I'll definitely have to check him out. Sci-fi is my favorite genre.
109D2: Mine as well. But i gotta go now, i've had a lovely time conversing. I'll see you later tomorrow:)
Bot: I look forward to it. Have a great day! I hope you have a great evening.
109D2: You too! Byeeee!
Bot: Bye, have a good evening, and good luck with the rest of your day.
109D2: [DONE]

Apart from inconsistency and being occasionally off-topic, the bot's inability to draw inferences also led to participants abruptly ending the conversation. Excerpts 5 and 6 show instances with the bot's obvious failure to understand non-compositional meaning and draw correct inferences, which, apparently, led to humans ending the conversations abruptly.

Excerpt 5 Participant 102D3

`102D3: You seem to lack concentration.`
`Bot: Yes, I do. It is hard for me to concentrate when I am stressed.`
`102D3: What happened? Cat got your tongue?`
`Bot: Yes, that is exactly what happened. I am so mad at my cat right now.`
`102D3: Okay, that was a good one.`
`Bot: I know, right? I hate it when my cat does that. It makes me so mad.`
`102D3: I think you should get back at it. Have any ideas for it?`
`Bot: I don't know what to do. I'm going to have to get a new one.`
`102D3: [DONE]`

In Excerpt 5 the participant first believes that the bot taking the idiom literally was a joke from the bot (*Okay, that was a good one.*). However, in the subsequent turn the human realizes that this was not a joke, as the bot keeps elaborating on it – even if it was intended as a joke, with the subsequent turns the bot disregards the quantity maxim. Upon realizing that the dialogue had become redundant and was going nowhere, the human terminates it suddenly. A similar phenomenon is shown in Excerpt 6.

Excerpt 6 Participant 101D1

`101D1: I enjoy hiking`
`Bot: I like hiking too. What is your favorite place to hike? Mine is the Appalachian National Scenic Trail.`
`101D1: I don't like that one`
`Bot: Why not? It is one of the longest hiking trails in the United States.`
`101D1: Exactly for that reason`
`Bot: Have you ever hiked it? It goes from Maine to Virginia and goes from Georgia to Maine.`
`101D1: [DONE]`

The participant made an ironic comment (*Exactly for that reason.*) and it went unnoticed as the bot kept talking about the trail, disregarding the participant's two negative answers which a human speaker would uptake as prompts to change topics. This lack of uptake resulted in the participant finishing the conversation–an observation that resonates with Allan's (2013) remark that it is CG that allows meaning to be underspecified by the speaker, as in language understanding a lot of inferencing is expected from the hearer. Excerpts 5 and 6 show that in some cases participants assumed that understanding indirect utterances is part of the existing, core CG, so their assessment of what the bot can do was not realistic – such instances can be understood as examples of egocentric language use. Even if the chatbot's non-cooperation was unintentional, it was still frustrating for some participants. In human-human interactional contexts, such instances of non-cooperation have been termed infringements (Thomas 1995). Some users may find this amusing or interesting, especially those interested in the workings of language and natural language production and understanding, but in our data, participants tended to quit interactions after such breakdowns.

The above excerpts show that humans often ended the conversation abruptly if it diverged from harmonious human-human conversational characteristics in terms of consistency, topic management, topic initiation, observance of conversational maxims and inferencing. This shows that for humans, successful emergent CG building would entail engagement from the bot's side. Abruptly terminating the conversation can be interpreted as human refusal to further engage in emergent CG building. When humans realized that they invested more energy in CG building which was not reciprocated by the bot, they tended to quit the conversation. By contrast, in cases when this desire was met, conversational closing sequences were more likely to follow the pattern of traditional human-human interactions (see Excerpt 4), as participants likely perceived the CG building as a more harmonious, symmetric process.

4.2.2 Assumed core common ground

4.2.2.1 Factual knowledge as assumed core common ground

A turn-by-turn close reading of the conversations also shed light on the type of assumed core CG that humans actually share with the conversational agent. A typical instance of core CG is factual knowledge, that is, information that can be fed into a model as training data. Excerpt 7 shows that perhaps the most obvious part of the common ground a human and a machine interlocutor can share is related to facts.

Excerpt 7 Participant 132D2
`132D2: I like watching series. Do you have any favorits?`
`Bot: I do! I love Game of Thrones. It's my favorite show. What about you?`
`132D2: Nice! Who is your favorite character? I have never seen Game of Thrones.`
`Bot: My favorite character is Daenerys Targaryen. You should definitely check it out!`
`132D2: I will for sure! When was the last time you watched it?`
`Bot: The last time I watched it was a few months ago. I really enjoyed it. It's based on "A Song of Ice and Fire" by George R.R. Martin.`

Excerpt 7 indeed looks unmarked, and could easily be mistaken for a human-human interaction, because participants have relatively comparable access to the information they talk about as it is part of their core CG. However, it is difficult to know what exactly forms the bases of such shared knowledge, as it largely depends on the data the model was trained on. In Excerpt 8, for example, a famous, supposedly known pop cultural reference to the Star Wars movies is not understood by the machine.

Excerpt 8 Participant 108D3
`108D3: Hello. May the force be with you :)`
`Bot: Thank you, I appreciate that. I'm sure it will all work out in the end.`

As has also been shown in Table 3, Participant 108 was a "curious" user (see also Tao et al. 2024) and most probably wanted to map what kind of CG they can share with this particular technology. As Allan (2013) notes, assumed CG is based on the speaker's assessment of the hearer's competence to understand the utterance, and this can motivate language choice.

4.2.2.2 Assumed and emergent common ground

Excerpt 9 shows an interesting example in which assumed core and emergent elements of CG were also detectable.

Excerpt 9 Participant 128D2
`Bot: What are you studying? I'm studying to be a teacher.`
`128D2: I am stydying English studies. Also I like Spanish language. And which languages do you know and study?`
`Bot: Spanish is the most widely spoken second language after English. I know a little bit of Spanish.`

128D2: Hola!
Bot: Hola. Do you speak any other languages besides English? I know French and German.

Here, what is common knowledge (Spanish as a popular second language – although the exact ranking may not be accurate as reported in the interaction) and what is discoverable from the interaction (that the conversational partner speaks Spanish) are both present. The code-switching that actually occurred in the interaction is evidence of the human mapping possible emergent CG – in this case with success. However, mere factual knowledge does not always suffice. In Excerpt 10 the conversation abruptly ends even if the bot possesses adequate factual background knowledge about Atlanta, - unfortunately the information was irrelevant, as the interaction was in a more introductory phase, where mutual questions about the other person, or simple small talk would have felt more natural. In this case core CG was not particularly important – the emergent CG, namely that the machine conversational partner disregards the Gricean maxim of relevance made the human quit the conversation.

Excerpt 10 112D1: Which city or town are you from? [from the previous turns it is known that the participant has lived in the US for some years]
Bot: I'm from Atlanta, Georgia. It's the most populous city in the state.
112D1: How's the weather there?
Bot: It's hot and humid. The Atlanta metropolitan area has a population of over 6.5 million people.
112D1: [DONE]

This example is different from Excerpts 2–3 insofar as the bot manages to stay on topic, shares core CG about Atlanta – however, through the emergent CG the human speaker is informed that the bot often gives factual information that is not very much related to the context. We have also seen in other examples (e.g., Excerpts 1 and 6) that humans get frustrated by the bot providing facts as this often violates the relevance maxim. Over two decades ago Saygin and Cicekli (2002) highlighted that the biggest hurdle in developing computer programs that can successfully carry out conversations with human users will be modeling their ability to cooperate in the Gricean sense. Apparently this criterion is still not met, as similarly to our present findings, recent studies have found that even ChatGPT, which relies on a more advanced type of AI compared to Blenderbot, and was trained on more robust datasets, often fails to follow the maxim of relevance (Kecskes and Dinh this issue; Tao et al. 2024) which makes conversation feel less human-like, and potentially irritating for users.

4.2.2.3 Individual differences in emergent common ground building

Participants took different approaches to manage the emergent CG. For example, there was an observable difference in the degree of the directness of their utterances.

There were many examples in which the bot could not create an emergent CG, mostly because it was unresponsive to the rather indirect conversational cues the human sent. Excerpt 11 is a great example of how the bot fails to uptake the repeated attempts of the human to make it stop talking about stand-up comedy (<u>underlined</u> turns) whereas most human speakers would have inferred that their interlocutor would like to change subject.

Excerpt 11 Participant 125D2

`Bot:` `I really like Louis C.K. He's a comedian from Los Angeles.`
`125D2:` <u>`I don't like stand up comedy`</u>
`Bot:` `That's too bad. What kind of comedy do you like? There are a lot of different styles.`
`125D2:` <u>`I like dark humour`</u>
`Bot:` `I love dark humor as well. There are so many great comedians out there. Do you have a favorite?`
`125D2:` <u>`No, I don't have one`</u>
`Bot:` `You should check out Louis C.K. if you haven't already. He's one of my favorites.`

In some cases participants expressed their frustration with the bot in more explicit ways. In Excerpt 12 the participant combined indirect (<u>underlined</u>) and direct (<u>**bold and underlined**</u>) utterances to signal their problems with the bot's inability to engage in creating emergent CG.

Excerpt 12 Participant 101D3

`101D3:` <u>**`Bot, you're not following me. You are not so clever after all`**</u>`.`
`Bot:` `I'm sorry, I didn't mean to be rude. Do you have any pets?`
`101D3:` `You're not rude, just not very… quick.` <u>`Yeah, I have a cat, Einstein.`</u>
`Bot:` `That's cool. I have two dogs and a cat. I also have a parrot.`
`101D3:` <u>`So you practically live in a zoo?`</u>

Bot: No, I live in the suburbs. I've never been to the zoo.
101D3: You should check it out, it's much fun. Do you have any interesting hobby, **to switch to the usual boring thread of conversation?**
Bot: I don't have much time for hobbies. I spend most of my time at work.
101D3: What do you work?
Bot: I work in a warehouse. It's pretty boring, but it pays the bills.
101D3: Must be some big bills with so many pets.
Bot: It can be, but I try to make the most of it. Do you have any pets?
101D3: **We have already covered this. Can we switch subjects?**
Bot: Sure, what do you like to do in your spare time? I like to play video games.
101D3: I like video games too. I'm an avid simracer.

An inconsistency of the bot is pointed out in the participant's first remark, which is followed by an indirect insult (*Einstein*). From the bot's answer it is evident that it mistook the insult and believed that the cat's name was Einstein. The participant's third turn is also a milder, indirect insult, followed by a more direct reference to the conversation being boring. The next remark is a reference back to the many pets, which fails again, as it is clear from the answer that the bot did not succeed in connecting the bits of information (the participant had a cat, as was discussed earlier in the conversation). Finally, the participant explicitly asks to change the subject. Contrasting Excerpts 11 and 12, some further individual variation surfaces – in Excerpt 11 the human only indirectly signaled that they would want to change subject, whereas in 12 after failing attempts with ironic remarks, the participant very directly asked to change the topic–a behavior that could have various reasons. For example, participant 101 could better understand what kind of CG they can expect from the bot and was more ready to update the assumed core CG based on their immediate experiences.

Conversations were more successful if humans did not have overly high expectations of the bot, for example, in terms of understanding inferred meaning (see e.g., Excerpt 6) or even if they had such initial expectations, they were ready to change strategies and progressively overwrite their initial understanding as the discussion unfolded (Excerpt 12). Thus, apparently, recipient design needs to become part of the emergent CG for more smooth communication with machines. As Keysar and Henly (2002) note, speakers tend to underestimate the ambiguity and overestimate the effectiveness of their utterances – an egocentric behavior that was detectable in Excerpt 11 and also initially in Excerpt 12. However, Participant 101 in Excerpt 12 changed their feedback strategies in line with what they learnt about the bot's natural language processing capabilities. Thus, similarly to Dippold (2023), we found that humans tried to identify which strategies would work with the particular technology. Additionally, we observed individual variation in recipient design, which is thought to depend on what core CG humans assume to have with the bot, how ready they are to change their initial conduct (Tuncer et al. 2024) and what mental models (Tolzin and Janson 2023) they have of the bot.

4.2.2.4 Common ground in successful interactions

Some participants had unmarked conversations with the bot over the three days that could easily have been mistaken for human-human conversations. In Excerpt 13, we show parts of Participant 127's conversations to present such unmarked conversations.

Excerpt 13 Parts of Participant 127's conversations from all three days

Day1 127: Hi!
Bot: Hello! How are you doing today? I just got back from walking my dog, how about you?
127: I am doing fine, i am at the university
Bot: That's great! What are you studying? I'm studying to be a nurse.
127: I am studying to be a psychologist
Bot: That's a great field to be in. I'm sure you'll be able to help a lot of people.
127: Yes that is one reason why I enrolled. Why did you choose nursing?
Bot: I wanted to help people with their mental health, so I decided to become a nurse.
127: If you wanted to help with their mental health, why didn't you choose psychology?
Bot: I thought about psychology but decided to go with nursing because I love helping people.
127: That's great, you sound like a great person

Day2 127: Hi I'm back. How are you today?
Bot: I'm doing well, how are you? I just got back from a long day of work.
127: I have just woke up, most of my classes were cancelled today
Bot: Oh no! I'm sorry to hear that. What are you going to do about it?
127: Don't be sorry, it's great, finally I can rest.
Bot: That's a good attitude to have. I hope you have a great rest of your day!
127: Thank you, you too. So what are your plans for today?
Bot: I'm going to play some video games. I've been playing a lot of Call of Duty lately.
127: Oh wow I didn't know you play Call of Duty.
Bot: Yeah, it's been a while since I've played it. I used to play it a lot when I was younger.

Day3 127: Hi, how are you?
Bot: I'm doing well, thank you. How are you this fine evening? Do you have any plans?
127: Yes, I'm going out tonight, we are painting the town red
Bot: That sounds like a lot of fun. I hope you have a great time. I'll be spending time with my family.
127: How many siblings do you have?
Bot: I have two brothers and two sisters. Do you have any brothers or sisters?
127: Yes I have a sister, she is great
Bot: That's great. I'm glad you have a good relationship with your sister. Family is important.

Even if the conversations with Participant 127 included turns where the conversation could have gone astray (e.g., mental health and nursing on Day 1, the bot being sorry about canceled classes on Day 2, painting the town red on Day 3), the participant, being eagerly cooperative managed to harness the flow of the dialogues. This example also shows that the context of an utterance comprises two distinct components: a *prior* context, involving individuals' prior experiences, mental frames, communication styles, etc., and an *emergent* situational context, which covers available cues for interlocutors in the actual situation (Kecskes and Zhang 2009). Although we cannot safely say to what extent coincidence also contributed to the discourse being unmarked, the elaborate, direct human turns which also contained some explanations (…*why didn't you choose psychology?* or *Don't be sorry, it's great, finally I can rest*) are definitely less egocentric than allusions and inferences, which often lead to less successful outcomes.

Common ground is also based on social conventions and cues as illustrated in Excerpt 14 which features a very elaborate dialogue with greetings, HAY adjacency pairs, small talk, and rituals.

Excerpt 14 Participant 110D1

`110D1: Hello!`
`Bot: Hello! How are you doing today? I just got back from walking my dog, how about you?`
`110D1: Feeling a bit down, how about you?`
`Bot: I am sorry to hear that. I am doing well, thank you. What is wrong?`
`110D1: Well, it's a bit complicated and i'm in a hurry at the moment, so i'd rather talk about something else.`
`Bot: That is understandable. I hope everything works out for you. I wish you all the best.`
`110D1: Thank you, I appreciate your words.`

Such social considerations seem to work with an AI-powered conversational agent. In earlier work with a more rigid SDS, Dombi et al. (2022) found that applying social language was not fruitful in interactions with a task-oriented dialogue system, and could not be factored into emergent CG, as polite utterances often caused misunderstandings. With the AI-powered chatbot, conversations were smooth from this perspective. Thus, apparently, the conversational dynamics of adjacency pairs, rituals and small talk, possibly because of their predictable nature, can more easily be taught to LLMs, and can form part of the core CG. The unmarked interactions shown in Excerpts 13 and 14 are also examples of how language is applied to manage social relationships, and how this is tied to the management of CG (Enfield 2008). In these cases, participants' mental models of the conversational agent were closer to reality, their ideas of the core CG they may share with the bot impacted how they interacted with the bot. Our findings also show that social elements of language are indeed part of human interactions with machines, as they appear to contribute to the mental model humans had of the AI-chatbot, which then influenced their expectations of the agent, and their linguistic behavior (see also Tolzin and Janson 2023).

In the data, particularly smooth or felicitous conversations were the result of two interrelated factors: (1) the human could make an accurate and realistic assumption of the kind of core CG that can be anticipated in the interaction (e.g., the human did not attempt indirect communication), and consequently core CG was successfully exploited by the bot, and (2) the bot did not reveal its difficulty in building emergent CG, that is, it managed to be consistent over turns, refrained from providing facts not relevant to the conversation at hand, and used appropriate social conventions (e.g., showing empathy). Humans' inaccurate assumptions of what core CG they could expect (e.g., cultural references, understanding inferences, or indirectness) often led to the bot disclosing this weakness. A further problem was the bot's accidental behavior: sometimes it provided "annoying facts" violating the maxim of relevance, sometimes it was internally inconsistent with previous turns, sometimes it nicely followed human-human conversational dynamics even showing empathy and applying social cues to the conversation. A possible reason behind this diverse behavior is that the bot's next turns in conversations are not rooted in core or emergent CG in their traditional sense but are mostly determined by statistical probability.

5 Conclusions

In this paper, using empirical data, we wanted to find out about the kind of CG humans assume to share with an AI-powered chatbot, and how their ideas about this knowledge change and feed into an emerging CG. In our previous work, we used Kecskes' SCA in a task-based context, where humans conversed with a simpler, SDS technology (Dombi et al. 2022). In this paper, we extended our analysis and applied the SCA in the context of human conversations with an AI-powered, open domain chatbot over three days. By showing changing trends in the conversations in terms of their length and participants' opening and closing sequences, we intended to demonstrate that extended exposure to interaction with an AI-powered chatbot indeed influenced humans' emerging CG. We demonstrated that participants changed their overall conversational behavior after the first day by applying more turns per conversation and adjusting their ritualized speech acts to match the capabilities of the technology that they progressively discovered. Keeping the SCA as a framework for analysis, such behavior can be explained as initially egocentric, and subsequently followed by signs of more conscious cooperation, realized through recipient design, as participants tend to discover the "rules of the game".

A turn-by-turn close reading of the conversations also shed light on the type of assumed core CG that humans share with the conversational agent. A typical instance of core CG was factual knowledge, but social cues such as showing empathy or engaging in small talk also resulted in unmarked conversations. We found evidence that core CG may be conceived of as relying on "memories of schemata, frames, scenarios and scripts" (Allan 2013: 18), whereas emergent CG derives from "blending individuals' prior personal experiences with perceptions of the actual situational context" (Kecskes 2013: 164) in a cooperative effort to tailor speech to match the bot's capacities, that is, to engage in recipient design.

Our data show that the conversations tended to be successful if humans managed (or wanted) to apply recipient design based on what they discovered during the interaction (Tuncer et al. 2024) and integrated the newly discovered conversational knowledge into their emergent CG. However, if recipient design could not be part of the emergent CG – either because the user was curious, playful, rigid, inexperienced, or frustrated – such scenarios led to marked conversations. We also found evidence that human interactional behaviors when facing interactional problems with the bot vary greatly on a continuum from abruptly closing the conversation (minimal desire to create emergent CG) through signaling problems indirectly or more directly to not pointing out problems at all. Possible reasons behind this variation may be that humans have different motives for wanting to interact with a bot, e.g., being curious how the system works, or pushing the limits of the technology (Tao et al. 2024), or it might also depend on humans' prior experiences with technology (Dippold 2023). This study did not inquire into participants' experiences in interacting with AI-powered chatbots or how they conceptualized such agents. Initial user assumptions about the bot's abilities likely influenced the ways participants approached the bot. That is, humans' conceptualizations of computer interactional partners' abilities may impact the extent to which they use egocentric versus cooperative language. Those who view these agents as more "human-like" are likely to apply less recipient design, meaning they accommodate the agent to a lesser degree. Also, humans' assumptions about the bot's abilities likely influence common ground construction – the more "human-like" participants believe the bot is, the more core CG they assume to share with it. Future studies may consider inquiring into participants' perceptions of computers as interlocutors and discussing language use in light of initial user assumptions. Also, experience with a particular technology is believed to affect individuals' language use. In a different study with ChatGPT as an interactional partner for language learners, we found that participants more experienced with ChatGPT had very different conversational expectations than novice users, which also surfaced in their language use (Sydorenko et al. 2024).

We conducted an exploratory analysis and therefore looked for possible traces of core and emerging CG that were exemplified in the data. To confirm these claims, further research can examine participants' verbal reports. For example, participants can engage in stimulated recall (Gass and Mackey 2000) after their final interaction with the bot and comment on their thought processes as they watch a video of their interactions with the bot. Such verbal reports may better uncover the mental models participants have of AI interlocutors, and even how the initial mental model may change progressively throughout the interaction. Also, future studies could look into different measures of specific adaptive behaviors (e.g., frequency of topic shifts, use of clarifying questions) to add to our understanding of how humans try to cooperate with AI-powered interlocutors.

References

Allan, Keith. 2013. What is common ground? In Alessandro Capone, Franco Lo Piparo & Marco Carapezza (eds.), *Perspectives on linguistic pragmatics*, 285–310. Cham: Springer.

Allan, Keith. 2014 [1986]. *Linguistic meaning*, 2nd edn. London: Routledge & Kegan Paul. (Reissued in one volume as *Routledge Library Editions: Linguistics* Volume 8, 2014).

Allan, Keith. 2023. The interdependence of common ground and context. In Istvan Kecskes (ed.), *Common ground in intercultural interactions*, 7–24. Berlin: Mouton DeGruyter.

Arundale, Robert B. 1999. An alternative model and ideology of communication for an alternative to politeness theory. *Pragmatics* 9. 119–154.

Barr, Dale J. & Boaz Keysar. 2005. Making sense of how we make sense: The paradox of egocentrism in language use. In Herbert L. Colston & Albert N. Katz (eds.), *Figurative language comprehension*, 21–43. Mahwah, NJ: Lawrence Erlbaum.

Bell, Allan. 1984. Language style as audience design. *Language in Society* 13(2). 154–204.

Beneteau, Erin, Olivia K. Richards, Mingrui Zhang, Julie A. Kientz, Jason Yip & Alexis Hiniker. 2019. Communication breakdowns between families and Alexa. In *Proceedings of the 2019 conference on human factors in computing systems*, 1–13. Glasgow, Scotland Uk: Association for Computing Machinery.

Blache, Philippe. 2017. Dialogue management in task-oriented dialogue systems. In Thierry Chaminade, Noël Ngyuen, Magalie Ochs & Fabrice Lefèvre (eds.), *Proceedings of the 1st ACM SIGCHI international workshop on investigating social interactions with artificial agents*, 4–8. New York, NY: Association for Computing Machinery.

Branigan, Holly P., Martin J. Pickering, Jamie Pearson & Janet F. McLean. 2010. Linguistic alignment between people and computers. *Journal of Pragmatics* 42(9). 2355–2368.

Branigan, Holly P., Martin J. Pickering, Jamie Pearson, Janet F. McLean & Ash Brown. 2011. The role of beliefs in lexical alignment: Evidence from dialogs with humans and computers. *Cognition* 121. 41–57.

Brennan, Susan E. 1998. The grounding problem in conversations with and through computers. In Susan R. Fussell & Roger Kreuz (eds.), *Social and cognitive approaches to interpersonal communication*, 210–255. Hillsdale: Lawrence Erlbaum.

Castricato, Louis, Nathan Lile, Suraj Anand, Hailey Schoelkopf, Siddharth Verma & Stella Biderman. 2024. Suppressing pink elephants with direct principle feedback. https://arxiv.org/abs/2402.07896 (accessed 24 June 2024).

Clark, Herbert H. 1996. *Using language*. Cambridge: Cambridge University Press.

Clark, Herbert & Susan E. Brennan. 1991. Grounding in communication. In L. Resnick, J. Levine & S. Teasley (eds.), *Perspectives on socially shared cognition*, 127–149. Washington, DC: American Psychological Association.

Colston, Herbert L. 2008. A new look at common ground: Memory, egocentrism, and joint meaning. In Istvan Kecskes & Jacob L. Mey (eds.), *Intention, common ground and the egocentric speaker-hearer*, 151–187. New York: Mouton de Gruyter.

Deppermann, Arnulf. 2015. When recipient design fails: Egocentric turn-design of instructions in driving school lessons leading to breakdowns of intersubjectivity. *Gesprächsforschung* 16. 63–101.

Diedrichsen, Elke. 2023. Grounding emergent common ground: Detecting markers of emergent common ground in a YouTube discussion thread. In Istvan Kecskes (ed.), *Common ground in intercultural interactions*, 105–135. Berlin: Mouton DeGruyter.

Dinan, Emily, Varvara Logacheva, Valentin Malykh, Alexander Miller, Kurt Shuster, Jack Urbanek, Douwe Kiela, Arthur Szlam, Iulian Serban, Ryan Lowe, Shrimai Prabhumoye, Alan W. Black, Alexander Rudnicky, Jason Williams, Joelle Pineau, Mikhail Burtsev & Jason Weston. 2020. The second conversational intelligence challenge (convai2). In Sergio Escalera & Ralf Herbrich (eds.), *The NeurIPS '18 competition*, 187–208. Cham: Springer.

Dippold, Doris. 2023. "Can I have the scan on Tuesday?" User repair in interaction with a task-oriented chatbot and the question of communication skills for AI. *Journal of Pragmatics* 204. 21–32.

Dombi, Judit, Tetyana Sydorenko & Veronika Timpe-Laughlin. 2022. Common ground, cooperation, and recipient design in human-computer interactions. *Journal of Pragmatics* 193. 4–20.

Enfield, Nicholas J. 2008. Common ground as a resource for social affiliation. In Istvan Kecskes & Jacob L. Mey (eds.), *Intention, common ground and the egocentric speaker-hearer*, 223–254. Berlin: Mouton de Gruyter.

Fischer, Kerstin. 2016. Designing speech for a recipient: The roles of partner modeling, alignment and feedback in so-called 'simplified registers'. Amsterdam: John Benjamins.

Fischer, Kerstin. 2021. Tracking anthropomorphizing behaviour in human-robot interaction. *ACM Transactions on Human-Robot Interactions* 11(1). 1–28.

Fridin, Marina & Mark Belokopytov. 2014. Acceptance of socially assistive humanoid robot by preschool and elementary school teachers. *Computers in Human Behavior* 33. 23–31.

Gass, Susan M. & Alison Mackey. 2000. *Stimulated recall methodology in second language research*. Mahwah, NJ: Lawrence Erlbaum.

Kecskes, Istvan. 2008. Dueling contexts: A dynamic model of meaning. *Journal of Pragmatics* 40(3). 385–406.

Kecskes, Istvan. 2010. The paradox of communication: Socio-cognitive approach to pragmatics. *Pragmatics and Society* 1(1). 50–73.

Kecskes, Istvan. 2012. Is there anyone out there who really is interested in the speaker? *Language and Dialogue* 2(2). 283–297.

Kecskes, Istvan. 2013. *Intercultural pragmatics*. Oxford: Oxford University Press.

Kecskes, Istvan. 2019. Cross-cultural and intercultural pragmatics. In Yan Huang (ed.), *The Oxford handbook of pragmatics* (Oxford Handbooks), 400–415. Oxford: Oxford University Press.

Kecskes, Istvan & Fenghui Zhang. 2009. Activating, seeking and creating common ground: A socio-cognitive approach. *Pragmatics & Cognition* 17(2). 331–355.

Keysar, Boaz. 2007. Communication and miscommunication: The role of egocentric processes. *Intercultural Pragmatics* 4(1). 71–84.

Keysar, Boaz & Anne S. Henly. 2002. Speakers' overestimation of their effectiveness. *Psychological Science* 13. 207–212.

Koulouri, Theodora, Stanislao Lauria & Robert D. Macredie. 2016. Do (and say) as I say: Linguistic adaptation in human-computer dialogs. *Human-Computer Interaction* 31. 59–95.

Levinson, Stephen C. 2006. Cognition at the heart of human interaction. *Discourse Studies* 8(1). 85–93.

Li, Chi-Hsun, Su-Fang Yeh, Tang-Jie Chang, Meng-Hsuan Tsai, Ken Chen & Yung-Ju Chang. 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*, 1–12. Honolulu, HI USA: Association for Computing Machinery.

Mustajoki, Arto. 2012. A speaker-oriented multidimensional approach to risks and causes of miscommunication. *Language and Dialogue* 2. 216–243.

Mustajoki, Arto. 2021. A multidimensional model of interaction as a framework for a phenomenon driven approach to communication. *Russian Journal of Linguistics* 25(2). 369–390.
Mustajoki, Arto. 2023. From laboratory to real life: Obstacles in common ground building. In Istvan Kecskes (ed.), *Common ground in intercultural interactions*, 59–80. Berlin: Mouton De Gruyter.
Mustajoki, Arto & Alla Baikulova. 2022. Avoidance of cognitive efforts as a risk factor in interaction. *Discourse Studies* 24(3). 269–290.
Myers, Chelsea, Anushay Furqan, Jessica Nebolsky, Karina Caro & Jichen Zhu. 2018. Patterns for how users overcome obstacles in voice user interfaces. In *Proceedings of the 2018 CHI conference on human factors in computing systems*, 1–7. Montreal, QC Canada: Association for Computing Machinery.
Nguyen, Dong, A. Seza Doğruöz, Carolyn P. Rosé & Franciska De Jong. 2016. Computational sociolinguistics: A survey. *Computational Linguistics* 42(3). 537–593.
Nolan, Brian. 2023. Understanding common ground as a cognitive object. In Istvan Kecskes (ed.), *Common ground in intercultural interactions*, 25–58. Berlin: Mouton DeGruyter.
Pickering, Martin J. & Simon Garrod. 2004. Toward a mechanistic psychology of dialogue. *Behavioral and Brain Sciences* 27(2). 169–190.
Radford, Alec, Jeffrey Wu, Rewon Child, David Luan, Dario Amodei & Ilya Sutskever. 2019. Language models are unsupervised multitask learners. *OpenAI Blog* 1(8).
Reeves, Byron & Clifford Nass. 1996. *The media equation: How people treat computers, television, and new media like real people and places*. New York, NY: Cambridge University Press.
Ribino, Patrizia. 2023. The role of politeness in human–machine interactions: A systematic literature review and future perspectives. *Artificial Intelligence Review* 56(1 Suppl). 445–482.
Roller, Stephen, Y.-Lan Boureau, Jason Weston, Antoine Bordes, Emily Dinan, Angela Fan, David Gunning, Da Ju, Margaret Li, Spencer Poff, Pratik Ringshia, Kurt Shuster, Eric Michael Smith, Arthur Szlam, Jack Urbanek & Mary Williamson. 2020. OpenDomain conversational agents: Current progress, open problems, and future directions. https://arxiv.org/pdf/2006.12442 (accessed 18 August 2024).
Sacks, Harvey. 1992. *Lectures on conversation*. Oxford: Basil Blackwell.
Sacks, Harvey, Emanuel A. Schegloff & Gall Jefferson. 1974. A simplest systematics for the organization of turn-taking in conversation. *Language* 50(4). 696–735.
Saygin, Ayse Pinar & Ilyas Cicekli. 2002. Pragmatics in human-computer conversations. *Journal of Pragmatics* 34(3). 227–258.
Schmader, Christopher & William S. Horton. 2019. Conceptual effects of audience design in human–computer and human–human dialogue. *Discourse Processes* 56(2). 170–190.
Serholt, Sofia, Lena Pareto, Sara Ekström & Sara Ljungblad. 2020. Trouble and repair in child–robot interaction: A study of complex interactions with a robot tutee in a primary school classroom. *Frontiers in Robotics and AI* 7. https://doi.org/10.3389/frobt.2020.00046.
Skantze, Gabriel & A. Seza Doğruöz. 2023. The open-domain paradox for chatbots: Common ground as the basis for human-like dialogue. https://arxiv.org/pdf/2303.11708 (accessed 24 June 2024).
Sreedhar, Makesh N., Traian Rebedea, Shaona Ghosh & Christopher Parisien. 2024. CantTalkAboutThis: Aligning language models to stay on topic in dialogues. https://arxiv.org/html/2404.03820v1 (accessed 10 June 2024).
Stalnaker, Robert C. 1978. Assertion. In Peter Cole (ed.), *Pragmatics*, 315–332. Leiden: Brill.
Stalnaker, Robert C. 2002. Common ground. *Linguistics and Philosophy* 25. 701–721.
Sydorenko, Tetyana, Judit Dombi & Veronika Timpe-Laughlin. 2024. Language learners' perceptions of written interactions with ChatGPT for practicing English. *Paper presented at the Technology for Second Language Learning Conference, Iowa State University*, 6–8 November.

Tao, Yufei, Ameeta Agrawal, Judit Dombi, Tetyana Sydorenko & Jung In Lee. 2024. ChatGPT role-play dataset: Analysis of user motives and model naturalness. https://arxiv.org/pdf/2403.18121 (accessed 17 May 2024).
Thomas, Jenny. 1995. *Meaning in interaction*. London: Longman Group Limited.
Tolzin, Antonia & Andreas Janson. 2023. Mechanisms of common ground in human-agent interaction: A systematic review of conversational agent research. In *Proceedings of Hawaii international conference on system sciences*, 342–351. Maui, Hawaii, USA: HICSS.
Tomasello, Michael. 2008. *Origins of human communication*. Cambridge: Massachusetts Institute of Technology Press.
Tuncer, Sylvaine, Christian Licoppe, Paul Luff & Christian Heath. 2024. Recipient design in human–robot interaction: The emergent assessment of a robot's competence. *AI & Society* 39. 1795–1810.

Bionotes

Judit Dombi
University of Pécs, Pécs, Hungary
dombi.judit@pte.hu

Judit Dombi is an Associate Professor of Linguistics at the University of Pécs, Hungary, where she teaches undergraduate and graduate courses in linguistics. Her fields of interest include theoretical and applied aspects of communication and pragmatics in intercultural and ELF contexts, as well as human-computer interaction. Her recent research has focused on the communicative functions of directness and indirectness in interaction, communication asymmetries, cooperation and the nature of common ground in human-computer interaction. Recently, her research has been published in *Applied Pragmatics, Intercultural Pragmatics, Computer Assisted Language Learning*, and *Journal of Pramgatics*.

Tetyana Sydorenko
Applied Linguistics, Portland State University, Portland, OR, USA

Tetyana Sydorenko is an Associate Professor of applied linguistics at Portland State University, Portland, Oregon. Her research interests include computer-assisted language learning (spoken dialogue systems, chatbots, Generative AI) and L2 pragmatics. She has published her research in journals such as *TESOL Quarterly, the Modern Language Journal, Language Learning & Technology, Computer Assisted Language Learning, and Journal of Pragmatics*.

Veronika Timpe-Laughlin
Educational Testing Service, Princeton, USA

Veronika Timpe-Laughlin is a Senior Research Scientist at Educational Testing Service, focusing on pragmatics, task-based language teaching, human-AI conversation, and young learner assessment. Her work has recently appeared *in Intercultural Pragmatics*, *Language Testing,* and *Language Teaching Research*. Additionally, she coordinates research for the TOEFL® Young Students and TOEFL ITP assessments and runs *Sharing Languages*, an organization that unites educators, researchers, and parents to advance early language and cultural education.