Designing Conversational Agents: A Self-Determination Theory Approach

Xi Yang
Imperial College London, UK
x.yang15@imperial.ac.uk

Marco Aurisicchio
Imperial College London, UK
m.aurisicchio@imperial.ac.uk

ABSTRACT

Bringing positive experiences to users is one of the key goals when designing conversational agents (CAs). Yet we still lack an understanding of users' underlying needs to achieve positive experiences and how to support them in design. This research first applies Self-Determination Theory in an interview study to explore how users' needs of competence, autonomy and relatedness could be supported or undermined in CA experiences. Ten guidelines are then derived from the interview findings. The key findings demonstrate that: competence is affected by users' knowledge of the CA capabilities and effectiveness of the conversation; autonomy is influenced by flexibility of the conversation, personalisation of the experiences, and control over user data; regarding relatedness, users still have concerns over integrating social features into CAs. The guidelines recommend how to inform users about the system capabilities, design effective and socially appropriate conversations, and support increased system intelligence, customisation, and data transparency.

CCS CONCEPTS

• **Human-centered computing** → Human computer interaction (HCI).

KEYWORDS

Conversational Agents, Competence, Autonomy, Relatedness, Conversational User Experience

ACM Reference Format:
Xi Yang and Marco Aurisicchio. 2021. Designing Conversational Agents: A Self-Determination Theory Approach. In *CHI Conference on Human Factors in Computing Systems (CHI '21), May 08–13, 2021, Yokohama, Japan.* ACM, New York, NY, USA, 16 pages. https://doi.org/10.1145/3411764.3445445

1 INTRODUCTION

Conversational Agents (CAs) such as Alexa, Google Assistant (GA) and Siri have become progressively integrated into our lives, as evidenced by their increased use in homes and in sectors such as education, healthcare, e-commerce and business. As much as computers and smartphones have impacted our life and society in recent decades, CAs are joining the stream of technology revolution [67, 68], which could bring us both benefits and negative effects. On the one hand, CAs pose risks because the consequences of the technological revolution are not always positive, similar to how smartphones may become sources of distraction and poor-quality sleep [75], not to mention the potential ethical, privacy and security risks of AI products [16]. This reminds us of the need to design CAs responsibly to reduce their potential negative impact. On the other hand, CAs provide opportunities, for example, the hands-free interaction brought by voice control and the potential to act on behalf of users because of the intelligence. As CAs become more ubiquitous, they will inevitably shape our behaviours and everyday experiences. This indicates the need to understand how CAs could be designed to facilitate the creation of positive experiences.

In the literature, there is increased interest in designing products and technologies that lead to positive experiences and psychological benefits [13, 22, 23, 39]. However, CAs have not yet been considered in this body of work partly because they have only recently become popular. As such, it is still unclear how to design CAs to create positive and meaningful experiences. One route to achieve this is to use Self-Determination Theory (SDT), a theory of human motivation, which suggests that people become self-determined and motivated to grow when their psychological needs are fulfilled [71]. Prior work has shown that SDT [71] is effective at facilitating the design process [39, 43, 62]. The SDT approach suggests that considering the fulfilment of psychological needs (i.e. competence, autonomy and relatedness) in design could result in substantial benefits, including but not limited to enhanced user experiences and human flourishing [62]. However, this theory has not yet been applied to the design of CAs. It is unknown what users expect from CAs regarding the three needs and how such needs could be supported by design. It is worth noting that the three needs sit among other psychological needs which also play important roles in positive experiences. For example, Sheldon et al. have identified ten needs that often exist in a positive experience [79]. Agreeing with SDT, Sheldon et al. also recognise competence, autonomy and relatedness as the most salient needs among the ten [79]. Overall, SDT was selected to be applied in this study because: 1) it is centered on the three most basic human needs; 2) it interprets psychological needs as means of achieving self-motivation and well-being [71]; 3) it is effective at facilitating the design process [39, 43, 62], although not yet applied to CAs.

Additionally, in the literature, most research insights are presented in the form of design implications or recommendations (e.g. in [17, 29, 37, 54, 85]) which could limit the effective application of these findings in design practice (e.g. whether the findings are practical enough to support design work) [2]. Presenting research findings as heuristics or guidelines, on the other hand, has shown to be helpful in facilitating the design and evaluation of products.

CHI '21, May 08–13, 2021, Yokohama, Japan

ACM ISBN 978-1-4503-8096-6/21/05…$15.00
https://doi.org/10.1145/3411764.3445445

For examples, Nielsen's usability heuristics [59] and Amershi et al.'s AI design guidelines [2] have become useful resources for the design of interactive and AI-infused products. However, there are still no such guidelines for CAs in the academic community. Nevertheless, the growing number of articles on how to design CAs in public domains, such as blogs [4, 9] or organisations' best practices [1, 36, 56] imply that such guidelines are in demand.

To fill these gaps, this research aims to apply SDT to develop design insights and guidelines to help the design of positive and meaningful user experiences with CAs. This was achieved through two phases.

Phase 1: Apply SDT to inform CA design.

Achieved through in-depth interviews with 14 participants which unpacked what the three psychological needs–competence, autonomy and relatedness–mean to CA users and how to design for these needs (Section 3).

- RQ1: What are users' perceptions and expectations regarding competence, autonomy and relatedness?
- RQ2: What aspects of CAs support or hinder the fulfilment of these needs?

Phase 2: Derive guidelines from Phase 1 findings.

Achieved by consolidating and synthesising the findings from Phase 1 into ten actionable guidelines, following a rigorous process as outlined by Amershi et al. in [2] (Section 4).

- RQ3: How can CAs be designed to support users' psychological needs?

Applying SDT to drive CA research allowed us to create design guidelines that are grounded in human needs and to bring psychological benefits to CA users. This study shows both the applicability and benefits of applying established psychological theory such as SDT to the domain of CA research, paving the way for theoretical-based research on CAs.

2 RELATED WORKS

2.1 The Design of Conversational Agents

In the literature, conversational agents (CAs), sometimes also referred to as "virtual assistants" or "digital assistants", are considered as applications that provide assistance to the user by "answering questions in natural language, making recommendations and performing actions"[3, 40, 54, 85]. Some commonly seen examples are Alexa, Siri and Google Assistant. In the last five years, the rise of CAs has made it evident that they can potentially influence many aspects of our lives. For example, businesses frequently choose to integrate CAs into their enterprise platforms to improve customer service or to promote productivity and collaboration at work [5, 8, 12, 49]. Beyond business, students can participate in group discussions moderated by CAs [83], and families can play games and have fun with CAs [7]. This indicates an urgent need to understand how this technology affects our experiences and how to design CAs to promote positive effects and mitigate negative ones.

Despite fast advancements in recent years, there is a long tradition of study regarding the design of intelligent agents. In 1994, Norman has posited that the main challenges of designing intelligent agents are not technical but social, that is, how people interact with and perceive agents [60]. This argument is still relevant today. In his paper [60], Norman outlines four important areas: 1) the user's feeling of control, i.e. ensuring that the user is in control of the system and comfortable with the agent's automated actions; 2) the accuracy of expectations, i.e. accurately depicting what the capabilities of the agent to minimise false hopes; 3) privacy concerns, i.e. easing the disturbing thoughts of many individuals that intelligent agents may have access to their personal information and financial activities; and 4) human–agent interaction, i.e. providing satisfying interactions with users so that they feel comfortable, natural and in control.

More recently, with the proliferation of CAs, the body of work pertaining to the design of CAs has begun to grow (e.g. [10, 14, 17, 29, 37, 42, 48, 54, 83, 85]). Some of the work has centred on human–agent conversations. For example, Porcheron et al. [63] have suggested that CAs should deliver resourceful responses which include cues about what the system was doing, why it was being done, and what was likely to be done next [25]. Myers et al. [57] have summarised four types of obstacles that users encounter during voice interaction–Natural Language Processing (NLP) errors, unfamiliar intent, system errors and failed feedback–and found that users handle these obstacles differently. Users speak louder or slower when they encounter NLP errors, whereas they are more likely to quit, settle (i.e. accept the flaw) or restart when they encounter the other types of obstacle.

Others have considered the modalities of human–agent interaction. For example, the graphical interaction in mobile CA apps may differ from voice interaction with a smart speaker. Considering these various interactions requires a focus on diverse contexts of use, such as different places (e.g. in a car, at home, at a workplace and in public) or different devices (e.g. smart speakers, smart phones, smart watches and TVs) [85]. Cowan et al. [20] have suggested switching between modalities when the context changes. For example, screen-based interaction should be minimised when moving into contexts that require hands-free operation. Moreover, researchers have provided suggestions to improve the fluidity and seamlessness of interactions, especially when integrating with third-party services [20, 85]. This is because failing to do so often interrupts users' hands-free operation, thus causing negative emotions, such as frustration.

Design recommendations have also been made that focus on revealing system adequacy and limitations [54] to help users establish accurate perceptions of CA intelligence. Such perceptions are often affected by how the system presents itself and information to the user, for example, the user interfaces or the ways in which the system delivers information to users [15]. This suggests that intelligence and limitations of CA systems be conveyed through these interactions.

The design of CAs also benefits from learnings on a broader spectrum of research on AI and interaction design in general. Within the scope of human–AI interaction, Horvitz [41] has proposed design principles for integrating automation with direct manipulation interfaces, such as considering uncertainty about a user's goals, allowing for efficient invocation and termination, employing socially appropriate behaviours for agent–user interaction and continuing to learn through observation. In recent years, the number of studies regarding various AI design topics continues to increase,

for example, transparency and trust [50, 66, 69, 81], personalisation [24, 30, 34], predictability and accuracy [35]. Most recently, Amershi et al. have derived 18 AI guidelines, which aim to help the design and evaluation of AI-infused systems [2]. Such guidelines are timely and a necessary complement to the general UI guidelines such as Nielsen's ten usability heuristics [58] due to the unique characteristics of AI technologies. For example, automated systems may have unpredictable behaviours and be disruptive or untrustworthy. Therefore, the traditional UI guidelines may not be fully applicable to AI systems, and new guidelines are needed.

Despite major improvements in the technologies, there is still a significant discrepancy between users' expectations and actual experiences with CAs [54]. How this discrepancy can be remedied remains unexplored. Moreover, most of the CA design insights discussed above are not expressed as actionable guidelines, which could limit their usefulness to be applied in real-world design practice. The value of articulating guidelines is present in examples such as the usability heuristics [59] and AI guidelines [2]. However, these guidelines are not fully adequate for CAs as they are intentionally designed to be general rather than specific [2]. The conversational interface of CAs presents new challenges and opportunities for designers, necessitating the provision of guidance regarding the design of conversational experiences.

2.2 Self-Determination Theory (SDT) and its Implications for Design

The design of positive experiences is strengthened when it supports people's psychological needs. For example, Hassenzahl et al. have utilised a set of needs, derived from the work of Sheldon et al. [79], to support experience design [38, 39]. The notion of needs satisfaction suggests that designers need to understand user expectations in relation to the needs and design to meet the expectations. Among the ten needs stated in [79], competence, autonomy and relatedness are the most salient and fundamental ones. SDT is a theory of human motivation that is centered on these three needs [71]. It posits that the fulfilment of competence, autonomy and relatedness will motivate the self to seek optimal function and growth. Using SDT, this study explores how the three basic human needs can be fulfilled in CAs by design. SDT was selected because it positions psychological needs as basic motives rather than mere experiential requirements and it has been successfully used to guide the design of interactive applications to achieve positive experiences (e.g. in [39, 43, 62]).

In SDT, competence describes "people's inherent desire to be effective in dealing with the surrounding environment", which is the desire of feeling capable and effective [21, 82]. Autonomy describes "people's universal urge to act in accord with their integrated interest and values" [21]. Relatedness describes "the propensity to experience the sense of belongingness and connectedness to others" [6, 21]. The three needs have shown to be predictive and reliable mediators of motivation, engagement and well-being [72, 80].

In design, the three needs often act as mediators which can be tweaked to improve the user experience. Specifically, they can be used as inspirations or criteria to adjust and evaluate a design [13, 62]. For example, Jansen et al. used SDT to develop engaging digital coaches to help users achieve their fitness goals [43]. In their study, SDT was applied to prepare and analyse user interviews and gather insights on how to increase the three needs through the digital coach. SDT was also applied to evaluate wearable activity trackers to improve product satisfaction [46]. The evaluation asked whether the product increased the feeling of mastery (competence), provided meaningful choices (autonomy) and facilitated the connection between the user and others (relatedness). In this way, specific features of the product were measured against the satisfaction of the psychological needs and adjusted accordingly to improve the user experience [62].

Previous research has shown that the sense of competence can be enhanced by offering optimal challenges, positive feedback or opportunities for learning [62]. For example, in video games, competence can be reflected in the controls being "intuitive", meaning that the controls should be easily mastered and not interfere with one's gaming experience [73]. Competence is also reflected in the perceived difficulty or novelty of a game, suggesting that designs should offer new difficulty levels that promise opportunities for learning [52, 62].

Autonomy is affected by system capability, complexity, misrepresentation and fluidity in software systems [31–33]. Friedman explains that user autonomy can be undermined if the system does not provide the necessary capabilities to help users realise their goals [31]. A loss of autonomy may also happen when users are given inaccurate information about the system or that the system does not adapt to the users' change of goals over time. Moreover, autonomy can be improved by providing users with sufficient choices and options [73] or by enhanced personalisation [70]. Beyond the interface level, autonomy is also facilitated by technology augmenting users' abilities to pursue personal goals and values, such as assistive technologies, behaviour change technologies or productivity tools [62].

Relatedness is enhanced when a product facilitates social interaction (e.g. communication tools enable long-distance connection to loved ones) or plays a symbolic role (e.g. Coca-Cola reminds us of a childhood family dinner, thus connecting us closer to our family) [86]. The first way concerns the "communication" function of a product, which the majority of research has focused on. For example, it has been suggested that designs should be meaningful and satisfy genuine relatedness rather than the mere semblance of connection [78]. The second way concerns the "symbolic" function of a product. This function reminds people of their past memories, which evokes the feeling of connectedness. For example, a family car could be a symbol of love because of the many family trips together in the car [86].

Although SDT shows positive promise to help with the design, it has not yet been explored in the context of CAs. This suggests an opportunity that motivated this research.

3 PHASE 1: APPLY SDT TO INFORM CA DESIGN

This phase explores how CAs could support or hinder the fulfilment of competence, autonomy and relatedness.

3.1 Method

Semi-structured interviews were conducted in early 2019. All interviews were carried out online via Zoom and were audio-recorded.

Table 1: Participants' information

P#	Age	Gender	Occupation	Previous CA experience		
				CA type	Duration of use	Frequency of use
P1	28	F	PhD student	Google Assistant	1–3 yr.	Daily
P2	49	F	Clerical officer	Alexa	1–6 mo.	Every other day
P3	37	M	Support worker	Alexa	1–3 yr.	Daily
P4	24	M	Shipping analyst	Alexa	6–12 mo.	Daily
P5	34	F	Pharmacy assistant	Alexa	1–6 mo.	Daily
P6	34	F	Teacher	Alexa	6–12 mo.	Daily
P7	26	F	Services coordinator	Google Assistant	1–3 yr.	Every other day
P8	26	F	Data entry clerk	Alexa	1–6 mo.	Daily
P9	36	M	Copywriter	Google Assistant	1–3 yr.	Daily
P10	38	M	IT specialist	Google Assistant	1–3 yr.	Daily
P11	37	M	Marketing director	Google Assistant	1–3 yr.	Daily
P12	50	F	IT specialist	Cortana	1–6 mo.	Every other day
P13	25	M	Business analyst	Alexa	1–3 yr.	Daily
P14	43	F	Property administrator	Google Assistant	1–6 mo.	2–3 times a week

Each session took approximately 60 minutes. Fourteen participants were interviewed (10 with undergraduate degree, four with post-graduate degree; 11 from Europe, one from Asia, two anonymous), see Table 1. Participants were recruited through the "respondent.io" platform. They were selected based on general experience with CAs: the type of CA they use most frequently (if more than one); and the duration and frequency of use within the past week. Participants' consent was collected prior to each interview session. They were paid $15.00 each for taking part.

3.1.1 Study Design. In the interview session, after warming up, we asked the participant focused questions around competence, autonomy and relatedness. Questions were informed and adapted from Peters et al.'s work [62] where they describe how psychological needs can be examined in technology-based experiences. For example, the satisfaction of needs can be assessed by asking whether "the user feels capable of using the technology" (competence), whether "the technology provides the user with useful options and choices" (autonomy) and whether "the technology makes the user feel connected to other people" (relatedness). Therefore, in this study, a set of interview questions were used to investigate each of the three needs as follows. To explore competence, questions were asked around the user's capability of and effectiveness in using the agent (e.g. "To what extent do you feel capable of using the CA and why?", "To what extent do you feel you know enough about the capabilities and limitations of your CA?", "Reflecting on your conversations with the CA, how satisfied are you with the way you ask questions and the answers you get? Why?"). To explore autonomy, questions were asked around the user's feeling of control and flexibility (e.g. "How much control do you think you have over the agent?", "What does control mean to you in the context of using the CA?", "How much flexibility do you have in your conversations with the agent?"). Concerning relatedness, questions were asked around the feeling of connectedness to other people and belongingness to the community (e.g. "To what extent does the agent make you feel connected to other people?", "What are your thoughts on the agent connecting you with the communities that you are in?", "Do you have any expectations or concerns in terms of connecting with other people through the CA?"). It is also worth a note that follow-up questions were asked in response to participants' answers.

3.1.2 Data Analysis. The interview data were transcribed and coded following an iterative and inductive thematic analysis. The analysis involved two researchers. Researcher 1 performed the initial analysis and identified the first set of codes and themes guided by the six phases outlined by Braun and Clarke [11], while researcher 2 was invited to discuss and refine the codes and themes over several sessions throughout the analysis process until any disagreements were resolved and consensus was reached. The data were analysed and coded to interpret users' perceptions and expectations of each of the three needs as well as the aspects of the CA that may support or hinder the needs.

Taking competence as an example, we first analysed the factors affecting users' sense of competence, in other words, what made people feel more (or less) competent. The data showed that competence is affected by how much the users know about the CA capabilities. Participants felt less competent if they thought that they knew little about what the CA can do. We then analysed the reasons causing this feeling. The data showed that new CA capabilities are emerging rapidly so users are often left unaware of what the CA can do. Following this, we analysed the consequences there may be, beyond decreased competence, if the issues are not addressed. The data showed that knowing little about the CA could lead to decreased engagement with the CA. The same process was also applied to the analysis of autonomy. Relatedness was analysed differently because most participants reported that they felt relatedness was not supported by CAs yet. Hence, the data mostly involved discussion on the potential values of the CA to support relatedness and users' major concerns of implementing social features in CAs. The analysis was, therefore, structured around the CAs' values and users' concerns. After this process, the codes were condensed into larger, overarching themes. The themes were then reviewed by five other researchers with expertise in HCI in the authors' institution

and the feedback used to further refine them. Finally, connections were established between the themes, which led to the key findings presented below.

3.2 Competence in CAs

The findings show that the users' sense of competence was often affected by (a) how much they thought they had made full use of the CA capabilities and (b) how effectively they communicated with the CA.

A) Making full use of the CA capabilities

The feeling of competence was undermined when participants felt they did not know all the CA capabilities, in their own words, "*not using the CA to the fullest*": "*because I feel I got a lot more to discover in its functionality*" (P5) or "*haven't really looked into anything other than what I use it for*" (P4). In general, the participants felt they only knew and used the basic functions of the CA or the ones they were already familiar with: "*I know it can set alarms and play music, so I use these functions. I don't know what else I can do with it. I haven't been given this knowledge*" (P1).

This was likely caused by how the participants had discovered and learnt the CA capabilities. Often, the participants had discovered new skills or capabilities by surprise: "*For a long time, I didn't know it [Alexa] can give definitions of words. Once, I asked the definition of a word, and it responded with a really good answer. I was like, 'Wow, I should do this more often"'* (P13). Other times, the participants had discovered new uses online or from their friends: "*We played games with it [Alexa] over Christmas. We didn't even know the games existed. Someone at work told me, so we gave it a try*" (P5). The fact that new capabilities kept emerging made users feel that they were not fully aware of what the CA was capable of. Despite the resources available online, the participants reported having little interest in approaching these resources actively, as this was often perceived as a huge effort, and the motivation to learn was low: "*[I don't need to learn new skills] there's nothing I can't do by myself with my phone or laptop*" (P12).

The extent of the users' knowledge of the CA capabilities could affect their engagement with the CA. As a result of insufficient knowledge, some participants had resigned to the fact that they should only use the basics of the CA. Others had wished for a fast and easy way to gain an overview of the capabilities: "*I'd like a list of all the things it can do*" (P8). On the contrary, knowing more about what the CA can do encourages the user to use it more often: "*When I realised what I could actually do with the assistant, I thought I would have used it more*" (P14). Setting realistic expectations has been found crucial when introducing new capabilities to users. In the first few interactions with a CA, over-promising or failing to live up to the user's expectations may eventually decrease engagement: "*The versatility and understanding were what we expected. But now it's much less than what we thought it was. At first, I was asking for anything to see what would happen. But now I've just kind of resigned myself to the basic functions*" (P8).

The above findings suggest three implications. The first is to provide an overview of the CA capabilities, ideally highlighting what a user has and has not tried. The second points to in-context education, which involves introducing relevant capabilities to users from within the conversation when that capability may be needed; this reduces the effort required for learning from outside resources. The third is to help the user set accurate expectations about the CA's performance.

B) Communicating effectively with the CA

The participants felt less competent when effective communication was hindered. The data showed that the participants' conversations with the CA were somehow shaped and restricted by it. They had to ask questions in a certain way to obtain the answers they needed, and they hardly had any control over what responses they would receive from the CA. For example, the participants had learnt to only ask very specific and structured questions: "*We have to be very specific about the questions we ask*" (P8) or "*be careful about how you structure your questions*" (P11). This participant was aware of the fact that talking to a CA was not the same as talking to a human being: "*You can't ask a question like you're speaking to a normal person. If you don't get the right answer, you have to restructure your question*".

There are a few reasons that may hinder effective communication. It could be due to a lack of robustness in voice recognition or weakness in the interpretation of complex sentences (e.g. the ability to retrieve or keep context like in human conversations): "*It doesn't understand the context correctly. It just gives an unrelated answer*" (P11). Another reason could be that the responses from the CA are often too robotic, lacking conciseness and specificity: "*I once asked about the sweetness of a new type of blueberry. It responded with an excerpt from Wikipedia. But that's not what I wanted. A human wouldn't do that. A human would just tell me how sweet it is. What I wanted was something more concise, not a scientific description. The same goes for films*" (P1). This participant was expecting a more human-like response, a concise and direct answer rather than a search result.

Not addressing these issues may cause negative consequences. First, future engagement could decline. Users could be unwilling to use more advanced features: "*I often use it [the CA] for dictation. But for anything related to the long form of writing, I would rather do it myself because it just can't work*" (P7).

Second, users may have less confidence in the CA, further precluding the accomplishment of more advanced tasks, for example, shopping via the CA: "*If I ask it to buy a film ticket, I wouldn't just give it the details and presume it's going to do it. I would still feel I need to watch over what it's doing; double-check that it's gone to the right place*" (P7). P12 expressed similar concern for writing email: "*I won't trust it 100% to write and send an email*".

Third, it may affect young children who are learning how to communicate–children may mimic the way they talk to the CA in the real world. As mentioned above, the participants had learnt to speak to the CA in a certain way. Such way of speaking requires being specific, direct and not always polite: "*You need to be quite direct at it like, 'Show me this"'* (P7). In this way, they were found to achieve effectiveness of communication at the expense of politeness. In particular, P7 said that she had learnt such a way of speaking from the CA: "*I have learnt how to do it. There are words that it prefers".* This shows that the participants had adjusted themselves to the way of talking that worked. During this process, the system (i.e. the CA) was teaching, or at least shaping, a style which worried P7: "*If you had a child, that would be a problem. We don't want them to talk to people like that in the real world*".

The findings suggest the following implications. First, the robustness of voice recognition should be improved to eliminate the need for a slow and clear talk. Second, the automatic retrieval of contextual information should be improved to reduce the need for specificity for every query and increase the opportunity for a more natural way of speaking. Third, more human-like responses need to be provided; these should be concise and straightforward. Finally, CA designers should take responsibility for eliminating the negative impacts of externalities on users. CA designers need to appreciate the fact that CAs are used by all age groups, and they can potentially shape how people (especially young children) talk to each other.

3.3 Autonomy in CAs

The findings show that the users' sense of autonomy was often reflected by their (a) having control over the conversation, (b) having a personalised experience and (c) having control over their data.

A) Having control over the conversation

User autonomy was reflected by whether they had control over their conversation with the CA, specifically, concerning how they asked questions and what responses they received: "*Control to me means how I ask the CA a question, and how it responds to me*" (P14). The data showed that such a sense of control was limited, and the interaction was more driven by the CA rather than the user. To be more specific, the participants reported having insufficient flexibility in how they could ask questions, and they felt they had no control over how the CA responds.

The flexibility was restricted because the usage of the CA was often engineered by the manufacturer, leaving the user no other alternatives. Thus the users had to adapt to the CA's way of operation and had little control themselves: "*Companies decide how the CA works, not me. . .if you have a couple of alarms on and want to change the time of one, you have to cancel the alarm and set up a new one. It has to be done in this way because it's engineered in this way. I don't have much control over that. . .this is not really how I want her to work*" (P8). Natural conversation was also said to be limited. The users were not able to talk freely to the CA as they tend to do to a human being.

The perceived lack of control was often caused by the responses being unpredictable and inconsistent, and a lack of sufficient explanation for the CA's behaviour. Sometimes the participants felt frustrated about the unexpected results: "*Often it comes with a surprise. I think I need more control because I often cannot get what I want*" (P1); "*You ask the same thing three days in a row. It will have three different ways of doing it*" (P7). Additionally, lacking explanation for the CA's unpredictable action can lead to a feeling of powerlessness: "*It [Alexa] is so itself. It's doing whatever it wants. If we ask it to turn on the TV, it doesn't do it, and we don't know why*" (P6).

One way to address the above issues is to allow users to customise how they want to ask questions and what responses they wish to receive. The concept of "routine" is one useful example. Routine is introduced in Google Assistant, Alexa and Siri (named shortcut). The idea is to allow the user to set up a trigger phrase or command for a routine (e.g. "good morning") and define what actions to receive in response to that command. For example, when the user says, "good morning", the CA would start playing the news, followed by the weather forecast and the schedule of the day. The user can set up their favourite actions in specific orders or even customise what news resources they wish to include.

Some participants mentioned they had used the routine feature and found it useful in the following ways. First, since it gives the user options to customise responses, it increases the chance of obtaining expected and consistent answers: "*It's quite valuable as I see a lot of potential in getting consistent r*esults" (P1). Second, routine allows the user to adapt the use of the CA to their own lifestyle, thus using the CA in accordance with their own values, for example, to make a productive day: "*I'd use it for scenarios like starting my workout playlist. And when I start to work, it could know me well and play a different playlist"* (P12), or to have a relaxing evening: "*when I say, 'good evening', it could start my evening cooking routine with music*" (P8). However, the routine idea works best when there is a sequence of actions related to diverse capabilities. This idea could gain more value if also made applicable to a single capability, meaning that the user can customise the trigger word and responses for a particular action. This way, the control is pushed towards the user's end. For example, this is what P5 hoped about her daughter using the CA when doing her homework: "*She's learning synonyms, so she asks a lot like, 'What's another word for happy?' If she can define the question herself, say, 'It's homework time', then when she says 'Happy', Alexa will automatically suggest alternative words for happy*".

The above findings suggest three implications. First, provide users with options to customise their conversations with the CA, as shown by the routine feature. Second, consider adding explanations for the CA's responses, especially when things go wrong. Third, improve the robustness of the technologies (e.g. NLP) to better understand the human language.

B) Having a personalised experience

The sense of autonomy manifests itself in having personalised experiences, which is made possible by an intelligent CA. This means that, among others, the CA learns about the user over time and can provide responses consistent with their preferences: "*Each time I ask Alexa to play that song, I have to specify which version I would like to listen to. Since I'm always asking the same thing every single day, Alexa should know by now that I want this version of the song*" (P6).

However, in reality, CAs are still making independent decisions in most cases without considering the user's preferences and habits: "*But now when I say the name of that song, Alexa still plays whatever it wants*" (P6). This is not uncommon. Most participants thought their CA was not intelligent enough to provide personalised responses: "*Everything is input manually rather than automatically*" (P10).

A perceived balance between control and intelligence is worth noting. On the one hand, a user with more control could lower the CA's intelligence, as this means the CA does what it is told to. On the other hand, users wish for some form of personalisation which requires the CA becoming intelligent and making decisions on their behalf: "*I think having more control over it means lowering its intelligence level and letting it do what I want it to do. But my less control means a more intelligent CA hiding more details from me*" (P1). Control and personalisation may be in conflict with each other, especially in cases where people's personal preferences change over time: "*Sure it can notice some pattern. But humans are always doing*

things slightly off. I often drink coffee for breakfast, but sometimes I drink tea" (P7).

The findings suggest two implications. First, improve the intelligence level of the CA, for example, increasing the accuracy in detecting the user's intentions based on what the user asks. Second, learn about the user's habits and preferences from past conversations and use this learning to provide responses that are tailored to their needs; allow the user to give feedback or take control when necessary, for example, ask the user for confirmation when the CA is unsure about the answer.

C) Having control over data

The sense of autonomy is also largely affected by whether the user has control over their data. Most participants talked about their concerns: "*I don't feel like I've got any control over it. I'm sure it's listening in the background and spying on me*" (P5); "*I don't know whether she's listening all the time, even if she's not responding to any commands*" (P6).

A lack of control over data leads to diminished trust. Some participants, for example, had decided that they would not put a CA speaker in the house because they did not trust it enough: "*I don't have Google Home speaker in my house because I don't trust that I have total control of it*" (P7); "*It's the reason I don't have Google or Alexa at home because, you know, there's always this conspiracy that they're listening to you all the time*" (P12). Other participants would switch the power off if they were concerned about their privacy: "*If you don't want Alexa to be involved, you switch it off*" (P2). There were also participants who had accepted this situation and admitted that this was part of what they had bought: "*I don't like the idea that it's collecting data on me, listening to my conversations... But I guess it's just part of what I paid for*" (P4). Despite this, they expected some transparency: "*Ideally, I would like to see exactly what has been set up. At the moment, it gives you a history of what you asked. It would be good to see what data have been sent to Amazon to see how they deal with the data and how the algorithms work*" (P3).

To address this concern in design, the findings suggest one implication. The user should be given more control over their data, for example, by allowing the user to view what data was collected by the CA, where the data was stored, and giving the user access to edit and delete their data if they wish to.

3.4 Relatedness in CAs

When asked whether the CA had made them feel connected to other people, very few answers were in the affirmative. One common reason is that hardly any functionality supports communication between people, nor is connecting with other people the main motivation for using the CA: "*I don't expect this. I wouldn't use it [to connect with other people] because I have my phone. I don't think Alexa is built for that*" (P2).

Although social functions are not fully available in (or may not be designed for) current CA systems (until the date of this study in early 2019), the data revealed three roles that CAs may play in supporting relatedness. It is worth noting that the participants had mixed views on these roles. The following section presents these roles in detail, elaborating on both the CA's benefits and the users' concerns.

A) CA as a communication tool (messaging and calling between individuals)

CAs could be used as communication tools to send messages and make calls or video chats. For example, Alexa has a "Drop" feature that allows users to talk to their contacts who also have the Echo device. The value of CA resides in the voice control, which makes it easier and faster to launch the communication channel.

Other than this, the participants found little value in using the CA as a communication tool. Three reasons were cited. First, they would not consider switching from the tools they were already using: "*I'd rather text or call. I doubt I would have the desire to use Google Assistant to message or make a call. I don't think that's something I will change what I do now*" (P14). P5 did not consider this either because this was not why she had bought the CA: "*This is not the reason we got Alexa. Especially for family, I'd much rather have that personal connection*". Another reason concerns the robustness of the system: "*I'd not be quite confident about asking the CA to make phone calls. A lot of my friends don't have regular English names. It would probably not recognise their names correctly*" (P11). Furthermore, many smart speakers lack a visual interface, which makes the CA less helpful in messaging. Sending messages via voice is less interactive because it is difficult to send emojis and other visual elements: "*I probably won't use it for communication. I would use text messaging and the phone because it's easier to send funny messages and select images through text apps rather than voice*" (P10).

The findings suggest that these issues need to be addressed before and during the implementation of social-related functions. Specifically, it is worth considering (1) the distinguished values that CAs may bring compared with traditional communication tools, (2) ways to gain users' confidence in utilising social functions on CAs and (3) ways to eliminate the limitations brought by the lack of visual interface.

B) CA as a voice interface to social channels (interacting with groups and communities)

Although most participants were reluctant to use the CA for individual communication, they did see value in connecting to social channels via their CA when it was about receiving news updates from groups or communities. This means that the user interacts with their social channels through the CA (e.g. listening to community news and posting messages to group channels). The voice interface of the CA makes interactions such as posting and commenting easier, which is especially helpful when the user needs to stay active in their groups. For example, when at work, "*it can connect to my slack, so I can ask it to read or reply to my message*" (P13). This also benefits parents by allowing them to stay active in their children's school community: "*The school has a Facebook group. If using the assistant, I may communicate more in the group.*" (P14). Moreover, the CA has the potential to pre-process information before delivering it to the user, for example, filtering out redundant or irrelevant information. In the study, a mum of two mentioned that she often receives duplicated notifications from the school: "*We have the school app, emails and text messages. . . We parents do get a lot of the same information. If you have more than one child, you get the same information many times. . .*" (P14). CA as an interface to the school channel has the potential to remove the duplicates.

Importantly, it is worth noting the users' concerns. First, the users did not adequately trust the CA for posting their comments on social channels, largely due to a lack of confidence in its robustness: "*My worry would be if I was having a conversation in a house and the CA picked up something randomly in the background, transcribed it and put it on the social media, I'd be uncomfortable with that*" (P10). Since this regards the user social profile, users want to be 100% sure before any action is taken: "*I'd need to be 100% confident about what I say and 100% sure it goes on the right platform*" (P10).

Second, the users were concerned about accessing social channels in shared environments. Often the CA is shared by the family. The users did not think it was appropriate for their children to listen to their social interactions. Moreover, a social account usually belongs to one individual. It is hard to use one shared device for all adults in the family: "*It's a family device. I don't really want my daughter to access my social media. Plus, my husband and I have separate social media accounts. We've got different interests, so different things come on our news feeds, and not all of our friends are mutual*" (P5).

Third, the users were uncomfortable about sharing data. In the case of social accounts, not only are the users' data accessible to the CA, their friends' and family's data may be revealed as well. The participants found this intrusive and unsettling: "*No, I wouldn't like it at all. It's intrusive because that means the assistant can access my social accounts. That would allow the assistant to read and learn about my friends. They haven't given any permission for that*" (P6).

Lastly, there were concerns about the lack of a visual interface. This undermines the usefulness of interactions, for example, for showing pictures and other visual elements.

Therefore, it is worth considering the following in CA design: first, ask for user permission before launching its social-related features, as these features may appeal to some, but not all, users; second, address the above issues before the development of features related to communication, making sure that the values outweigh the concerns.

C) CA as a facilitator of social activities

CAs can also support relatedness by facilitating fun and engaging social activities. This allows family members or friends to bond with each other. For example, using the CA to play games with family members helps parents to bond with their children: "*The only thing that has something to do with connection is that we use it to play games. Together we try to get an answer to the game*" (P5). Another example is that talking to the CA was found to break the ice between people or to invoke an interesting conversation: "*That would be a social thing; like 'Showing what GA can do' could be the topic of a conversation with friends when they're visiting*" (P11).

Although CA may not have intentionally been designed as a facilitator of social activities, this finding highlights CAs' implicit values and opportunities associated with social activities alongside their primary functions. These benefits give more prominence to "designing experiences" than "designing products". In other words, the design of the CA should begin with contemplating the experiences it can create, followed by concentration on designing product functionality. As the data show, CAs have the potential to enrich social interactions and bring additional social benefits. It is, therefore, useful to be aware of such benefits and, if possible, turn them into explicit designs in the future.

4 PHASE 2: DERIVE DESIGN GUIDELINES

This phase translates the findings and implications from Phase 1 into design guidelines.

4.1 Method

First, we took the implications from Phase 1 and filtered out those that pertain more to technology advancement than design, such as "improve the robustness of voice recognition", "improve the intelligence of CA" and "improve NLP technologies", as they were out of the scope of this research. This step resulted in ten implications. Then we organised them into four categories based on the phases of user interaction identified in [2]: "initially", "during interaction", "when wrong" and "over time". Next, we translated the ten implications into actionable design guidelines following the guidance outlined in [2] such as "written as a rule of action; starting with a verb; accompanied by a one-sentence description that qualifies or clarifies any potential ambiguities; not contain conjunctions". Finally, the guidelines were discussed, reviewed and iterated by the authors of the paper until an agreement was reached.

Note that relatedness was not taken into account in the creation of the guidelines. There are two main reasons. First, unlike competence and autonomy, relatedness depends heavily on particular features of the product, that is, features that allow users to communicate or interact with one another. However, such features were not part of CAs when the study was conducted. Second, capabilities linked to relatedness were not always expected in CAs. The findings from Phase 1 showed that participants had little motivation to use CAs for social purposes. Considering relatedness differently than the other two needs is not uncommon. For example, in Peng et al.'s study [61], relatedness was excluded for similar reasons. Despite of this, and given that the development of social capabilities in CAs is still in its infancy, it is useful to discuss the possible challenges and opportunities for CAs to support relatedness. These insights (as presented in Sec. 3.4) could help the design of CAs' social functions and the development of associated guidelines in the future.

4.2 Design Guidelines for CAs

The ten guidelines are presented in Table 2 and discussed in detail below. The first six guidelines were developed from implications relevant to competence and the other four from implications relevant to autonomy. It is worth noting that, psychological needs often co-exist in human experiences with various degrees of salience [79]. This means that although competence (or autonomy) may be the dominant need that is fulfilled through application of a guideline, the other need could also be satisfied.

Initially G1–G3 concern educating users about the CA's capabilities.

G1. Provide a personalised overview of CA capabilities: help the user gain a full picture of the capabilities compared to what they already know. Given the fast-growing number of capabilities, users need a way to access an overview of what the CA can do, ideally highlighting what they are already familiar with, what is new and what is relevant to them. In current CAs, available capabilities are often shown in a summary page organised into categories such as "things to try", Figure 1 (a). A page like this only partly helps the user understand what they can do with the CA

Table 2: Design Guidelines for CAs

Initially	G1: Provide a personalised overview of CA capabilities. Help the user gain a full picture of the capabilities compared to what they already know.
	G2: Introduce new capabilities in-context. Make it convenient for the user to discover and access relevant capabilities.
	G3: Reveal how well the CA can perform when introducing new capabilities. Help the user set accurate expectations about the CA capabilities.
During interaction	G4: Learn about the conversational context to maintain the flow of a conversation. Help the user have effective communication with the CA.
	G5: Present responses in a concise and informative way. Make it easy for the user to retrieve information.
	G6: Talk politely. Encourage polite and socially appropriate conversation style.
When wrong	G7: Provide an explanation regarding why the CA cannot complete a task. Help the user understand the current system status.
Over time	G8: Learn about user habits over time from past interactions. Help the user obtain tailored services from the CA.
	G9: Provide users with options to customise the commands and responses. Allow the user to have more control of the conversation when needed.
	G10: Provide opportunities for user data management. Allow the user to view and manage their personal data.

due to the generic organisation and the lack of a personalised view. A personalised summary, on the other hand, could be more informative. For example, the categories could be reordered or tailored according to the user's usage history and preferences instead of using generalised categories.

G2. Introduce new capabilities in-context: make it convenient for the user to discover and access relevant capabilities when needed during the conversation and concurrently with the flow of actions. For example, the CA may recommend cooking-related capabilities while the user is asking about dinner ideas. In current CAs, new capabilities are often introduced through a "what's new" page, Figure 1 (b). Unlike in-context introduction, this approach requires the user to exit their ongoing conversation with the CA and find the "what's new" page independently. This research suggests considering when and how a new capability is introduced.

G3. Reveal how well the CA can perform when introducing new capabilities: help the user set accurate expectations about the CA capabilities. This implies that the CA should not exaggerate what it can actually do. For example, during the Christmas season, Alexa introduced the skill "call Santa", Figure 1 (c). However, the first few trials failed as Alexa responded with "I can't find the name 'Santa' in your contact list". Although it worked eventually, users' initial excitement had already been replaced by disappointment. Therefore, when introducing new capabilities, it is important to anticipate users' expectations and ensure that they are met.

During interaction: G4–G6 concern the conversations between the CA and the user.

G4. Learn about the conversational context to maintain the flow of a conversation: help the user have effective communication with the CA. For example, the CA could remember and replay information previously shared in the conversation, so the user does not need to repeat the same information. Maintaining the flow of conversation is key to effective communication, which requires the CA to understand the conversational context. For example, if a user asked to book a flight and later wanted to change the return date, the user should not have to repeat the flight information, Figure 2 (a). In conversation design, it becomes challenging when the conversation jumps between multiple sub-topics, which is common in human communication. A user, for example, may ask about a flight to Paris, get distracted by the attractions and then ask about the Louvre museum. After the "distraction", the user may go back to the "flight booking" topic. In this case, the CA should be able to pick up on previous information to complete the original task.

G5. Present responses in a concise and informative way: make it easy for the user to retrieve information. In cases where responses are in the form of search results, current CAs have different forms of presentation. Sometimes, the CA gives a short briefing with the search link. Other times, it provides a direct answer, Figure 2 (b). It has also been noticed that such forms of presentation are often inconsistent, Figure 2 (c), and there is no option enabling the user to choose the form that better suits their needs. When making such design choices, this guideline suggests considering which approach could help the user retrieve useful information more easily.

G6. Talk politely: encourage polite and socially appropriate conversation style. This guideline concerns the negative influence that CAs may have on the way we talk, especially on young children. It has been found that it is not uncommon for users to speak to CAs in a commanding fashion, for example, "Do this" or "Stop".

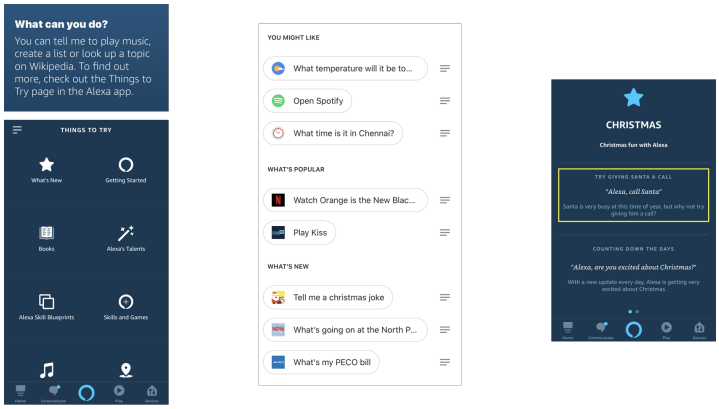

Figure 1: Examples of G1 to G3 (a) When asked "What can you do?", Alexa points the user to "things to try" page. (b) "Explore what's new" page on the GA app. (c) When asked "call Santa", Alexa is supposed to connect the user to Santa, but in the first few trials, users heard "I can't find the name 'Santa' in your contact list" and experienced disappointment.

In fact, users have been implicitly instructed by CAs to speak this way because CAs understand the command better when talking directly. Although this might not be an issue for adults, CAs are often shared among family members, and children are increasingly exposed to CAs or are CA users themselves. In this research, parents raised concerns about the potential negative externalities of CAs. This guideline suggests considering designs that discourage and eliminate such debatable influences, for example, by improving how the CAs understand and respond to complex and polite languages.

When wrong: G7 concerns how to deal with situations when the CA malfunctions.

G7. Provide an explanation regarding why the CA cannot complete a task: help the user understand the current system status. There are many reasons why CAs cannot complete certain tasks, including system settings or technical limitations. Regardless of the reason, it is beneficial to provide users with explanations regarding any malfunction. For example, explaining why the CA cannot complete a task, Figure 3 (a), helps the user make sense of

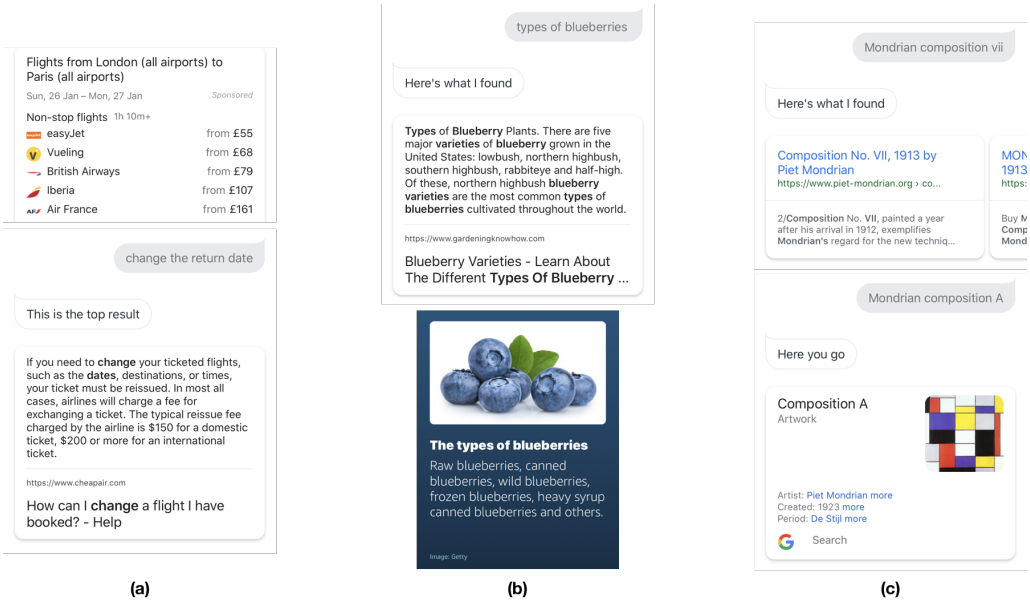

Figure 2: Examples of G4 to G6 (a) After a conversation with GA about "flights to Paris", the user wanted to change the return date; however, they had to repeat the conversation to do this. (b) When asked for "types of blueberries", GA replied with a search result (top); Alexa replied with a direct answer (bottom). (c) When asked about an artist's name and painting, sometimes GA replies with search results (top) and other times with a direct answer (bottom); the answers often come randomly and inconsistently.

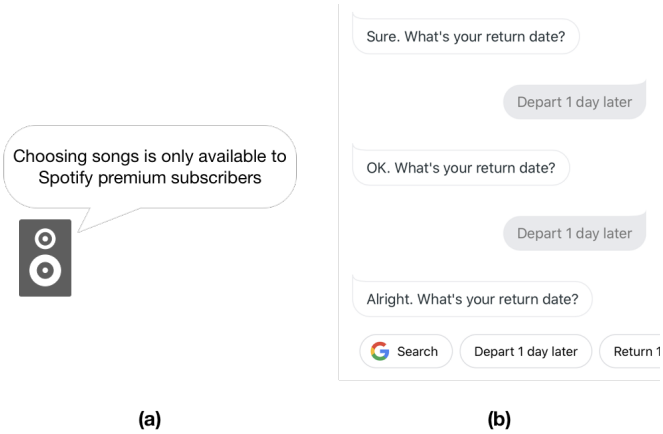

Figure 3: Examples of G7 (a) Google Home explained why it could not play the song. (b) Something went wrong with GA as it kept repeating the question "What's your return date?" and offered no explanation.

and resolve the current situation. A lack of explanation, conversely, may frustrate the user, Figure 3 (b).

Over time: G8–G10 concern the design and long-term use of CAs.

G8. Learn about user habits over time from past interactions: help the user obtain tailored services from the CA. In current CAs, some of the content can be personalised through service providers. For example, when a CA is asked to play a song, it is the media service (e.g. Spotify) that provides personalised music to the users based on data from that service platform, Figure 4 (a). However, the CA itself still provides little personalisation (except for the basics such as location). In this research, participants reported that they did not feel the CA was learning from their past conversations.

G9. Provide users with options to customise the commands and responses: allow the user to have more control of the conversation when needed. This means giving users the flexibility to customise their conversations with the CA, for example, defining how they ask questions and how they want the answers to be presented to them. In current CAs, the "routine" feature offers such opportunities, Figure 4 (b), although its primary intention is to combine multiple actions into one. This guideline suggests customisation within the action. For example, in most CAs, the user can now customise the news sources so that when the command "play news" is launched, the CA only plays news from the specified sources. However, so far as we know, the customisation is only limited to certain capabilities (e.g. "play news" as mentioned above). Enabling customisation of more capabilities could give users more flexibility and control of the conversation.

G10. Provide opportunities for user data management: allow the user to view and manage their personal data. The goal is to provide users with transparency on the data, for example, whether and when their data is collected and how it is used. Current CAs provide some form of data management. For example, users can find the privacy policy or activity history in the apps, Figure 4 (c). However, these could be designed to be more easily accessible, as in this research most participants did not know about the existence of such features.

5 DISCUSSION

Building on SDT, this research explores how to support competence, autonomy and relatedness in CA design, and develops ten SDT-based guidelines that articulate the key guiding principles relevant to competence and autonomy. The guidelines build connections between design and psychological needs satisfaction, and have the potential to enable designers to create CA experiences that bring psychological benefits to users. In this section, we first discuss the three needs. Then, we compare the CA guidelines proposed in this study to the AI guidelines by Amershi et al. [2] and the UI heuristics by Nielsen [58]. Each of the three sets focuses on a different design scope, i.e. CA design, AI design and interaction design, respectively. Through this comparison, we discuss what remains constant across the three scopes and what is new.

5.1 Design for Competence with CAs

Competence is not only related to how well the user can perform the current task [65, 73, 82]. It also concerns what the user knows about the system's capabilities. This is particularly relevant for CAs, due to the fact that many CA systems (e.g. Google Assistant, Alexa) play the role of a service platform where the user has access to external services. For example, through a CA, the user can listen to news from the BBC, play music from Spotify or order food from Uber Eats. Moreover, the number of services[1] and thus the capabilities of CAs is growing rapidly, for example in the United States the number increased by 147% within just one year from 2018 to 2019, to a total of 4,253 services [44]. The acceleration in CA capabilities makes it difficult for the user to gain a clear picture of what the CA can do. Even in 2016, when the number of services was much lower than today, users reported feeling overwhelmed by the unknown potential of CAs [54]. This research shows that such a feeling of ignorance could decrease a user's sense of competence. In response to this, this research suggests designs that educate users about the capabilities, for example through personalised overview, and designs that support the discoverability of the capabilities and provide convenient access to relevant ones.

5.2 Design for Autonomy with CAs

The sense of autonomy is found to decrease where the user has limited control over the conversations with CAs. For example, participants reported that their CAs often respond inconsistently and unpredictably, leaving them with little control of how the conversation will happen. On the other hand, users expect to have human-like conversations with CAs, as this is more natural and familiar to them [18, 20]. However, this inherently increases the chances of CAs behaving independently and unpredictably [63], thus increasing users' sense of lack of control. This contradiction may be addressed by differentiating between two conversational purposes:

[1] 'Services' may be referred to by other terms, depending on the CA: in the case of Google Assistant, they are called 'apps' or 'actions'; in the case of Alexa, they are called 'skills'.

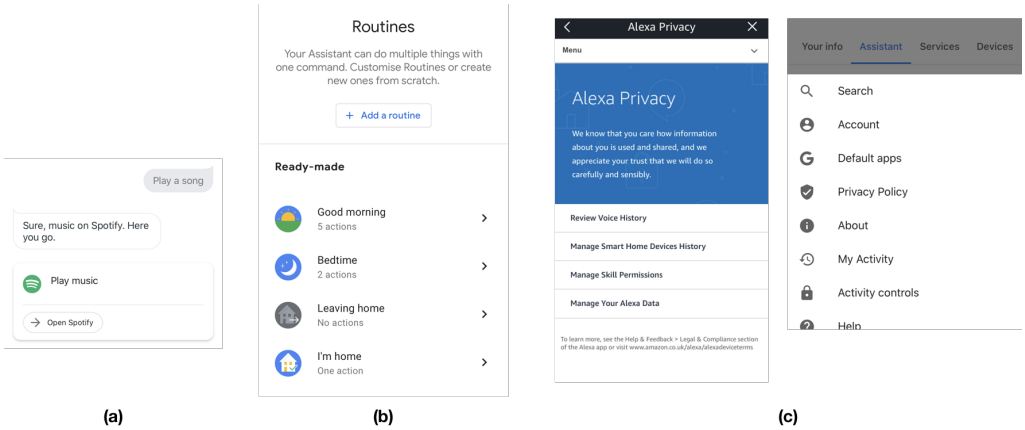

Figure 4: Examples of G8 to G10 (a) Music can be personalised through the service provider (e.g. Spotify) rather than the CA; screenshot from GA app. (b) Through the routine feature, users can customise the trigger phrase and actions to respond; screenshot from GA app. (c) Users can find privacy policies in CA apps; Alexa app settings page (left); GA app settings page (right). However, these pages were not easy to find.

transactional and social [27, 76]. In conversations with transactional purposes (e.g. seeking answers or requesting information), consistency and accuracy in CA responses are expected from users. Therefore, it is worth considering providing users with more control in such conversations. Conversely, in conversations with social purposes (e.g. perceiving the CA as a friend or a guru), unpredictability may add enjoyment to the user's experiences [51]. Therefore, such scenarios may benefit from more human-like conversation.

User autonomy is also affected by level of personalisation. It has been found that users expect personalisation to be applied to their behaviours and everyday life, for example, through CAs learning and adapting to users' daily habits and routines. This points to another key question in the design of CAs: how to design the human–agent relationship. In previous works, it has been shown that people were resistant to build relationships with CAs beyond the level of "strangers" or "casual acquaintances" [19]. However, findings from this paper contradict that argument, as a higher level of personalisation inevitably implies a closer relationship than that between strangers or acquaintances. In such cases, more information is shared and more common ground established, as described in the linguistic literature [44]. Importantly, given that improved personalisation is often based on the user sharing more data with the agent, this raises security risks and privacy concerns. It is, therefore, vital to consider how to balance the level of personalisation with potential risks and concerns.

5.3 Design for Relatedness with CAs

In the literature, there is little discussion around the topic of relatedness in CAs. Existing research has mostly explored the social roles of CAs with particular types of use cases, for example, supporting the elderly [74], playing with children [26, 55] and helping people with disabilities [64]. However, the social functions of CAs within people's everyday life and how relatedness can be supported have not been adequately explored. This paper contributes to this topic by identifying users' expectations and concerns with regard to the social roles of CAs in their daily lives.

The findings show that CA users' expectations of relatedness are mixed. This is connected to the readiness of the technology (e.g. not yet incorporating mature social-specific features), the shortcomings of (or lack of) visual UIs (e.g. not all CA devices have screens, which makes it inconvenient to use visual elements such as emojis, which are popular in social interactions [28]), the awkwardness of having social interactions in a shared environment (e.g. talking to a CA speaker about Facebook news feeds when families are around), as well as privacy concerns (e.g. the intrusiveness of a CA accessing information about a user's social circles).

Despite users' concerns and the high dependency on technological developments, the findings stress a design direction that does not introduce these extra concerns, i.e. the CA playing a role in facilitating social activities, for example, strengthening parent–child relationships by participating in family bonding activities. In recent years, this direction of design has started to emerge in studies that explore CAs providing companionship to individuals [53], supporting teamwork in schools [84], relieving employee anxiety in workplaces [47] and providing entertainment to family members [45, 77]. These studies have all touched upon the social effects of CAs from different perspectives. However, they have not explicitly linked these effects to the user's sense of relatedness. This paper has articulated the value of the roles that CAs may play, and proposed that if such social functions are successfully applied, CAs could indirectly help improve social relationships between people and thereby increase their sense of relatedness.

5.4 Comparing Guidelines for CA, AI and UI Design

Table 3 presents the CA guidelines along with related AI and UI guidelines. Comparing the CA guidelines to the AI and UI guidelines it can be seen how the design principles have evolved across application types. Table 3 reveals three groups: 1) guidelines that are relevant to all three fields (coloured green), 2) guidelines that mainly apply to intelligent applications – AI and CA applications

Table 3: Comparing CA guidelines to AI and UI guidelines

	CA guidelines (this research)	AI guidelines [2]	UI guidelines [58]	Comparison notes
Initially	G1: Provide a personalised overview of CA capabilities	Make clear what the system can do		G1 proposes to personalise how system capabilities are received, whereas the AI guideline calls for clarity on system capabilities
	G2: Introduce new capabilities in-context		Recognition rather than recall	G2 seeks to make new capabilities more discoverable by making them retrievable when needed, falling within the broad UI principle
	G3: Reveal how well the CA can perform when introducing new capabilities	Make clear how well the system can do what it can do		G3, acknowledging the evolving nature of CAs, stresses the importance of revealing the performance of new capabilities, a principle generalised in the AI guideline
During interaction	G4: Learn about the conversational context to maintain the flow of a conversation	Show contextually relevant information		G4 extends the AI guideline by asking to learn from the conversational context and using it to inform CA responses
	G5: Present responses in a concise and informative way		Match between system and the real world	G5, aligning with the broad UI principle, seeks to support natural human-agent interaction by making information retrieval intuitive
	G6: Talk politely	Match relevant social norms		G6, extending the AI guideline, stresses the conversational style of CAs
When wrong	G7: Provide an explanation regarding why the CA cannot complete a task	Make clear why the system did what it did	Help users recognise, diagnose and recover from errors	G7 makes the AI and UI guidelines CA-specific by suggesting to explain system malfunction through conversational UI
Overtime	G8: Learn about user habits over time from past interactions	Learn from user behaviour		G8 stresses the need to infer user habits from conversations, as generalised in the AI guideline
	G9: Provide users with options to customise the commands and responses G10: Provide opportunities for user data management	Provide global controls	User control and freedom	G9 and G10, adding to the AI and UI guidelines, provide specifications on what the controls are and how to provide them

(coloured blue), and 3) guidelines that are specific to CAs (coloured orange) though falling under the broad area of some UI guidelines.

In the first group (green), regardless of the application type (i.e. traditional UI, AI or CA application), the design principles related to handling system errors (G7) and providing user control (G9 and G10) are consistently relevant. In the UI guidelines, Nielsen calls for "help[ing] users recognise, diagnose, and recover from errors" [58]. When transferring this principle to AI applications, CAs included, the emphasis is on explaining why the system did what it did, as the AI guideline advises to, "make clear why the system did what it did" [2]. This is largely due to the unpredictability of AI applications. Therefore, transparency and visibility with regard to the system status (in this case, why the system malfunctioned) could help the user make sense of the situation, thus increasing their feeling of control [60].

The other principle that remains relevant is "user control and freedom" [58], although the actual types of control (e.g. what to control and how to control it) could vary largely and are often

dependent on the applications. For example, for AI applications, the control should allow the user to "customise what the AI system monitors and how it behaves" [2]. In the case of CAs, two types of control are identified. The first pertains to the control of conversations, specifically, how the user instructs the CA, and how the CA responds (G9). The second pertains to user data management, specifically the management of personal data collected by the CA over time (G10).

In the second group (blue), the guidelines are specifically applicable to intelligent applications (i.e. AI and CAs). They are related to revealing system capabilities and limitations (G1, G3), the effectiveness and social norms of human–agent interactions (G4, G6) and personalisation (G8). Despite the fact that the five guidelines share similarities with some of the AI guidelines, they are more specific about CAs than the AI guidelines and thus provide additional informative instructions regarding CA design. For example, G1 concerns an issue similar to the AI guideline "make clear what the system can do". However, instead of a general statement, G1 specifically emphasises the "personalised overview" of what the CA can do. This is important to CAs, as previous studies have consistently shown that many users are not aware of their CAs' capabilities [54, 85]. This could result in users feeling overwhelmed by unknown capabilities [54], and it could decrease users' sense of competence. The contrast between generalisation and specialisation applies also to the other four guidelines (G3, G4, G6, and G8). For example, G3 is aligned with the AI guideline "how well the system can do what it can do", but it stresses helping users set accurate expectations of the CA when introducing new capabilities to them.

Finally, in the third group (orange), the design principles, related to the discoverability of CA capabilities (G2) and the design of CA responses (G5), apply primarily to CAs as they have not been mentioned in previous AI guidelines and loosely fall under the umbrella of UI guidelines as indicated in Table 3. Both guidelines concern issues that are specific to CA products. For example, as the number of CA capabilities continues to grow, their discoverability becomes a challenge, which G2 aims to address. In the case of G5, it focuses on how to provide human-like responses to users, a challenge that typically occurs in conversations between users and CAs.

In addition to providing novel guiding principles and practical ways to implement the principles in CAs, this research also offers a theoretical foundation (i.e. SDT) to the guidelines. For example, G6 is founded on the need to support competence. Competence is influenced by the effectiveness of the conversation. In current CAs, effective conversation may encourage a commanding speaking style. As a result, people end up sacrificing politeness for effectiveness. However, G6 can also support autonomy, which is affected by the user's control of the conversation. Since current conversations are more led by the CA, user autonomy is limited. As a result, users often adapt to the CA's behaviour, which may not represent a polite manner. This shows that SDT-based guidelines can provide a more in-depth and holistic view of the phenomena (e.g. G6 is more than just using polite language) and suggest solutions contributing to psychological benefits.

5.5 Limitations and Future Work

The guidelines were developed from the perspective of SDT, and thus, they might not be comprehensive. An approach informed by a different theory or one that uses a synthesis of the literature, similar to that of Amershi et al. [2], might result in different guidelines. We consider the development of guidelines an ongoing process. Through this work, we hope to encourage more research like this to contribute practical guidelines for designers and developers in order to create intuitive and effective CAs. Additionally, although developed from empirical research, the guidelines have not yet been evaluated by design experts. To improve the practicability of the guidelines, an important next step is to investigate how they can be applied in design processes to support CA designers.

6 CONCLUSIONS

In this research, SDT has been applied to develop new understanding of what the three psychological needs (competence, autonomy and relatedness) entail in the CA experiences. Insights into users' perceptions and expectations on the three needs are obtained. These insights have enabled the development of informative suggestions for supporting the needs fulfilment in CA design. Specifically, the study describes how competence could be affected by users' knowledge of CA capabilities and the effectiveness of the conversation, and autonomy could be affected by users' control of the conversation and their data as well as personalisation. The study also discusses users' mixed expectations around relatedness, presenting the values and users' concerns over integrating social features into CAs. Finally, a set of CA design guidelines have been developed, which complement existing UI and AI guidelines. The guidelines offer suggestions to educate users on CA capabilities, design effective and natural conversations, consider the tone and politeness of CAs, respond resourcefully when CAs malfunction, personalise user conversations, and offer conversation customisation and data control options.

REFERENCES

[1] Amazon. 2020. Designing for Conversation. Retrieved April 3, 2020 from https://developer.amazon.com/ask-resources/guided/conversational-design-workshop#/

[2] Saleema Amershi, Kori Inkpen, Jaime Teevan, Ruth Kikin-Gil, Eric Horvitz, Dan Weld, Mihaela Vorvoreanu, Adam Fourney, Besmira Nushi, Penny Collisson, Jina Suh, Shamsi Iqbal, and Paul N. Bennett. 2019. Guidelines for Human-AI Interaction. In *Proceedings of the 2019 CHI Conference on Human Factors in Computing Systems - CHI '19*, 1–13. https://doi.org/10.1145/3290605.3300233

[3] C. Baber. 1993. Developing interactive speech technology, Interactive speech technology: human factors issues in the application of speech input/output to computers. Taylor & Francis, Inc., Bristol, PA.

[4] Ruben Babu. 2019. The conversational designer's guide to AI assistants. *Medium*. Retrieved December 4, 2019 from https://uxdesign.cc/the-conversational-designers-guide-to-ai-assistants-by-ruben-babu-6aa45de578f0

[5] Daniel Baier, Alexandra Rese, Maximilian Röglinger, D. Baier, A. Rese, and M. Röglinger. 2018. Conversational User Interfaces for Online Shops? A Categorization of Use Cases. In *ICIS*.

[6] Roy F. Baumeister and Mark R. Leary. 1995. The need to belong: desire for interpersonal attachments as a fundamental human motivation. *Psychological bulletin* 117, 3: 497.

[7] Diana Beirl, Nicola Yuill, and Yvonne Rogers. 2019. Using voice assistant skills in family life.

[8] Eva Bittner, Sarah Oeste-Reiß, and Jan Marco Leimeister. 2019. Where is the bot in our team? toward a taxonomy of design option combinations for conversational agents in collaborative work. In *Proceedings of the 52nd Hawaii International Conference on System Sciences*.

[9] Hillary Black. 2019. What is Conversation Design, and How to Design Your Chatbot. *Medium*. Retrieved December 4, 2019 from https://medium.com/swlh/

what-is-conversation-design-and-how-to-design-your-chatbot-3754f04ab1e7

[10] Michael Bonfert, Maximilian Spliethöver, Roman Arzaroli, Marvin Lange, Martin Hanci, and Robert Porzel. 2018. If you ask nicely: a digital assistant rebuking impolite voice commands. In *Proceedings of the 20th ACM International Conference on Multimodal Interaction*, 95–102.

[11] Virginia Braun and Victoria Clarke. 2006. Using thematic analysis in psychology. *Qualitative research in psychology* 3, 2: 77–101. Retrieved December 1, 2016 from http://www.tandfonline.com/doi/abs/10.1191/1478088706qp063oa

[12] Thomas M. Brill, Laura Munoz, and Richard J. Miller. 2019. Siri, Alexa, and other digital assistants: a study of customer satisfaction with artificial intelligence applications. *Journal of Marketing Management* 35, 15–16: 1401–1436.

[13] Rafael A. Calvo and Dorian Peters. 2014. *Positive computing: technology for well-being and human potential*. MIT Press.

[14] Heloisa Candello, Claudio Pinhanez, Mauro Pichiliani, Paulo Cavalin, Flavio Figueiredo, Marisa Vasconcelos, and Haylla Do Carmo. 2019. The Effect of Audiences on the User Experience with Conversational Interfaces in Physical Spaces. In *Proceedings of the 2019 CHI Conference on Human Factors in Computing Systems - CHI '19*, 1–13. https://doi.org/10.1145/3290605.3300320

[15] Justine Cassell. 2000. *Embodied conversational agents*. MIT press.

[16] Raja Chatila, Kay Firth-Butterflied, John C. Havens, and Konstantinos Karachalios. 2017. The IEEE global initiative for ethical considerations in artificial intelligence and autonomous systems [standards]. *IEEE Robotics & Automation Magazine* 24, 1: 110–110.

[17] Ana Paula Chaves and Marco Aurelio Gerosa. 2018. Single or Multiple Conversational Agents?: An Interactional Coherence Comparison. In *Proceedings of the 2018 CHI Conference on Human Factors in Computing Systems*, 191.

[18] Minji Cho, Sang-su Lee, and Kun-Pyo Lee. 2019. Once a Kind Friend is Now a Thing: Understanding How Conversational Agents at Home are Forgotten. In *Proceedings of the 2019 on Designing Interactive Systems Conference*, 1557–1569.

[19] Leigh Clark, Nadia Pantidi, Orla Cooney, Philip Doyle, Diego Garaialde, Justin Edwards, Brendan Spillane, Emer Gilmartin, Christine Murad, and Cosmin Munteanu. 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*, 1–12.

[20] Benjamin R. Cowan, Nadia Pantidi, David Coyle, Kellie Morrissey, Peter Clarke, Sara Al-Shehri, David Earley, and Natasha Bandeira. 2017. What can i help you with?: infrequent users' experiences of intelligent personal assistants. In *Proceedings of the 19th International Conference on Human-Computer Interaction with Mobile Devices and Services*, 43. Retrieved September 12, 2017 from http://dl.acm.org/citation.cfm?id=3098539

[21] Edward Deci and Maarten Vansteenkiste. 2004. Self-determination theory and basic need satisfaction: Understanding human development in positive psychology. *Ricerche di Psicologia* 27: 23–40.

[22] Pieter Desmet and Marc Hassenzahl. 2012. Towards Happiness: Possibility-Driven Design. In *Human-Computer Interaction: The Agency Perspective*, Marielba Zacarias and José Valente de Oliveira (eds.). Springer Berlin Heidelberg, 3–27. https://doi.org/10.1007/978-3-642-25691-2_1

[23] Pieter M.A. Desmet and Anna E. Pohlmeyer. 2013. Positive Design: An Introduction to Design for Subjective Well-Being. *International Journal of Design* 7, 3: 5–19. Retrieved March 9, 2016 from http://www.ijdesign.org/ojs/index.php/IJDesign/article/view/1666

[24] Tilman Deuschel and Ted Scully. 2016. On the Importance of Spatial Perception for the Design of Adaptive User Interfaces. In *2016 IEEE 10th International Conference on Self-Adaptive and Self-Organizing Systems (SASO)*, 70–79.

[25] Paul Dourish and Graham Button. 1998. On" technomethodology": Foundational relationships between ethnomethodology and system design. *Human-computer interaction* 13, 4: 395–432.

[26] Stefania Druga, Randi Williams, Hae Won Park, and Cynthia Breazeal. 2018. How smart are the smart toys? children and parents' agent interaction and intelligence attribution. In *Proceedings of the 17th ACM Conference on Interaction Design and Children*, 231–240.

[27] Suzanne Eggins and Diana Slade. 2005. *Analysing casual conversation*. Equinox Publishing Ltd.

[28] Ahmed Fadhil, Gianluca Schiavo, Yunlong Wang, and Bereket A. Yilma. 2018. The effect of emojis when interacting with conversational interface assisted health coaching system. In *Proceedings of the 12th EAI International Conference on Pervasive Computing Technologies for Healthcare*, 378–383.

[29] Ethan Fast, Binbin Chen, Julia Mendelsohn, Jonathan Bassen, and Michael S. Bernstein. 2018. Iris: A Conversational Agent for Complex Tasks. In *Proceedings of the 2018 CHI Conference on Human Factors in Computing Systems*, 473.

[30] Leah Findlater and Joanna McGrenere. 2004. A comparison of static, adaptive, and adaptable menus. In *Proceedings of the SIGCHI conference on Human factors in computing systems*, 89–96.

[31] Batya Friedman. 2010. *Value sensitive design*.

[32] Batya Friedman, Peter Kahn, and Alan Borning. 2002. Value sensitive design: Theory and methods. *University of Washington technical report*, 02–12.

[33] Batya Friedman and Helen Nissenbaum. 1996. User autonomy: who should control what and when? In *Conference Companion on Human Factors in Computing Systems*, 433.

[34] Krzysztof Z. Gajos, Mary Czerwinski, Mary Czerwinski, Desney S. Tan, and Daniel S. Weld. 2006. Exploring the design space for adaptive graphical user interfaces. In *Proceedings of the working conference on Advanced visual interfaces*, 201–208.

[35] Krzysztof Z. Gajos, Katherine Everitt, Desney S. Tan, Mary Czerwinski, and Daniel S. Weld. 2008. Predictability and accuracy in adaptive user interfaces. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*, 1271–1274.

[36] Google. 2019. Conversation design. *Google Developers*. Retrieved December 4, 2019 from https://developers.google.com/assistant/actions/design

[37] Arthur C. Graesser, Nia Dowell, and Danielle Clewley. 2017. Assessing Collaborative Problem Solving Through Conversational Agents. In *Innovative Assessment of Collaboration*. Springer, 65–80.

[38] Marc Hassenzahl. 2010. *Experience Design: Technology for All the Right Reasons*. Morgan & Claypool Publishers.

[39] Marc Hassenzahl, Kai Eckoldt, Sarah Diefenbach, Matthias Laschke, Eva Len, and Joonhwan Kim. 2013. Designing moments of meaning and pleasure. Experience design and happiness. *International Journal of Design* 7, 3. Retrieved February 22, 2016 from http://search.proquest.com/openview/f1d52b35adb0c2727cc8478e2d1bfbd8/1?pq-origsite=gscholar

[40] Johann Hauswald, Michael A. Laurenzano, Yunqi Zhang, Cheng Li, Austin Rovinski, Arjun Khurana, Ronald G. Dreslinski, Trevor Mudge, Vinicius Petrucci, and Lingjia Tang. 2015. Sirius: An open end-to-end voice and vision personal assistant and its implications for future warehouse scale computers. In *Proceedings of the Twentieth International Conference on Architectural Support for Programming Languages and Operating Systems*, 223–238.

[41] Eric Horvitz. 1999. Principles of mixed-initiative user interfaces. In *Proceedings of the SIGCHI conference on Human Factors in Computing Systems*, 159–166.

[42] Ting-Hao'Kenneth' Huang, Joseph Chee Chang, and Jeffrey P. Bigham. 2018. Evorus: A Crowd-powered Conversational Assistant Built to Automate Itself Over Time. *arXiv preprint arXiv:1801.02668*.

[43] Arne Jansen, Maarten Van Mechelen, and Karin Slegers. 2017. Personas and Behavioral Theories: A Case Study Using Self-Determination Theory to Construct Overweight Personas. In *Proceedings of the 2017 CHI Conference on Human Factors in Computing Systems*, 2127–2136. Retrieved June 25, 2017 from http://dl.acm.org/citation.cfm?id=3026003

[44] Barbara Johnstone. 2018. *Discourse analysis*. John Wiley & Sons.

[45] Jasmine Jones and Mark S. Ackerman. 2018. Co-constructing Family Memory: Understanding the Intergenerational Practices of Passing on Family Stories. In *Proceedings of the 2018 CHI Conference on Human Factors in Computing Systems*, 1–13.

[46] Evangelos Karapanos, Rúben Gouveia, Marc Hassenzahl, and Jodi Forlizzi. 2016. Wellbeing in the making: peoples' experiences with wearable activity trackers. *Psychology of well-being* 6, 1: 4.

[47] Everlyne Kimani, Kael Rowan, Daniel McDuff, Mary Czerwinski, and Gloria Mark. 2019. A Conversational Agent in Support of Productivity and Wellbeing at Work. In *2019 8th International Conference on Affective Computing and Intelligent Interaction (ACII)*, 1–7.

[48] Philipp Kirschthaler, Martin Porcheron, and Joel E. Fischer. 2020. What Can I Say? Effects of Discoverability in VUIs on Task Performance and User Experience. In *Proceedings of the 2nd Conference on Conversational User Interfaces*, 1–9.

[49] Jana Koehler. 2018. Business Process Innovation with Artificial Intelligence: Levering Benefits and Controlling Operational Risks. *European Business & Management* 4, 2: 55–66.

[50] Todd Kulesza, Margaret Burnett, Weng-Keen Wong, and Simone Stumpf. 2015. Principles of explanatory debugging to personalize interactive machine learning. In *Proceedings of the 20th International Conference on Intelligent User Interfaces*, 126–137.

[51] Jill Fain Lehman and Iolanda Leite. 2016. Turn-taking, children, and the unpredictability of fun. *AI Magazine* 37, 4: 55–62.

[52] J. Derek Lomas, Kenneth Koedinger, Nirmal Patel, Sharan Shodhan, Nikhil Poonwala, and Jodi L. Forlizzi. 2017. Is difficulty overrated?: The effects of choice, novelty and suspense on intrinsic motivation in educational games. In *Proceedings of the 2017 CHI conference on human factors in computing systems*, 1028–1039.

[53] Irene Lopatovska and Harriet Williams. 2018. Personification of the Amazon Alexa: BFF or a mindless companion. In *Proceedings of the 2018 Conference on Human Information Interaction & Retrieval*, 265–268.

[54] Ewa Luger and Abigail Sellen. 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*, 5286–5297. Retrieved August 28, 2017 from http://dl.acm.org/citation.cfm?id=2858288

[55] Emily McReynolds, Sarah Hubbard, Timothy Lau, Aditya Saraf, Maya Cakmak, and Franziska Roesner. 2017. Toys that listen: A study of parents, children, and internet-connected toys. In *Proceedings of the 2017 CHI Conference on Human Factors in Computing Systems*, 5197–5207.

[56] Microsoft. 2019. Principles of bot design. Retrieved December 4, 2019 from https://docs.microsoft.com/en-us/azure/bot-service/bot-service-design-principles

[57] Chelsea Myers, Anushay Furqan, Jessica Nebolsky, Karina Caro, and 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*, 6.

[58] Jakob Nielsen. 1994. 10 Heuristics for User Interface Design: Article by Jakob Nielsen. *Nielsen Norman Group*. Retrieved October 24, 2019 from https://www.nngroup.com/articles/ten-usability-heuristics/

[59] Jakob Nielsen. 1995. 10 usability heuristics for user interface design. *Nielsen Norman Group* 1, 1.

[60] Donald A. Norman. 1994. How might people interact with agents. *Communications of the ACM* 37, 7: 68–71.

[61] Wei Peng, Jih-Hsuan Lin, Karin A. Pfeiffer, and Brian Winn. 2012. Need Satisfaction Supportive Game Features as Motivational Determinants: An Experimental Study of a Self-Determination Theory Guided Exergame. *Media Psychology* 15, 2: 175–196. https://doi.org/10.1080/15213269.2012.673850

[62] Dorian Peters, Rafael A. Calvo, and Richard M. Ryan. 2018. Designing for Motivation, Engagement and Wellbeing in Digital Experience. *Frontiers in Psychology* 9. https://doi.org/10.3389/fpsyg.2018.00797

[63] Martin Porcheron, Joel E. Fischer, Stuart Reeves, and Sarah Sharples. 2018. Voice Interfaces in Everyday Life. In *Proceedings of the 2018 CHI Conference on Human Factors in Computing Systems*, 640.

[64] Alisha Pradhan, Kanika Mehta, and Leah Findlater. 2018. Accessibility Came by Accident: Use of Voice-Controlled Intelligent Personal Assistants by People with Disabilities. In *Proceedings of the 2018 CHI Conference on Human Factors in Computing Systems*, 459.

[65] Andrew K. Przybylski, C. Scott Rigby, and Richard M. Ryan. 2010. A motivational model of video game engagement. *Review of General Psychology* 14, 2: 154–166. https://doi.org/10.1037/a0019440

[66] Emilee Rader, Kelley Cotter, and Janghee Cho. 2018. Explanations as mechanisms for supporting algorithmic transparency. In *Proceedings of the 2018 CHI Conference on Human Factors in Computing Systems*, 103.

[67] Sid Reddy. 2018. Conversational Artificial Intelligence – The Final Phase Of The Information Revolution. *Hidden Insights*. Retrieved January 13, 2020 from https://blogs.sas.com/content/hiddeninsights/2019/05/03/hidden-insights-samuel-blanquet/

[68] Sid Reddy. 2018. Council Post: Conversational Artificial Intelligence – The Final Phase Of The Information Revolution. Retrieved January 13, 2020 from https://www.forbes.com/sites/forbestechcouncil/2018/08/17/conversational-artificial-intelligence-the-final-phase-of-the-information-revolution/#2a98cbd37de8

[69] Marco Tulio Ribeiro, Sameer Singh, and Carlos Guestrin. 2016. " Why should i trust you?" Explaining the predictions of any classifier. In *Proceedings of the 22nd ACM SIGKDD international conference on knowledge discovery and data mining*, 1135–1144.

[70] Scott Rigby and Richard M. Ryan. 2011. Glued to games: How video games draw us in and hold us spellbound: How video games draw us in and hold us spellbound. AbC-CLIo.

[71] Richard M. Ryan and Edward L. Deci. 2000. Self-determination theory and the facilitation of intrinsic motivation, social development, and well-being. *American psychologist* 55, 1: 68. Retrieved June 20, 2016 from http://psycnet.apa.org/journals/amp/55/1/68/

[72] Richard M. Ryan and Edward L. Deci. 2017. *Self-determination theory: Basic psychological needs in motivation, development, and wellness*. Guilford Publications.

[73] Richard M. Ryan, C. Scott Rigby, and Andrew Przybylski. 2006. The Motivational Pull of Video Games: A Self-Determination Theory Approach. *Motivation and Emotion* 30, 4: 344–360. https://doi.org/10.1007/s11031-006-9051-8

[74] Alessandra Maria Sabelli, Takayuki Kanda, and Norihiro Hagita. 2011. A conversational robot in an elderly care center: an ethnographic study. In *2011 6th ACM/IEEE International Conference on Human-Robot Interaction (HRI)*, 37–44.

[75] Muhammad Sarwar and Tariq Rahim Soomro. 2013. Impact of smartphone's on society. *European journal of scientific research* 98, 2: 216–226.

[76] Klaus P. Schneider. 1988. *Small talk: Analyzing phatic discourse*. Hitzeroth.

[77] Alex Sciuto, Arnita Saini, Jodi Forlizzi, and Jason I. Hong. 2018. " Hey Alexa, What's Up?" A Mixed-Methods Studies of In-Home Conversational Agent Usage. In *Proceedings of the 2018 Designing Interactive Systems Conference*, 857–868.

[78] Kennon M. Sheldon, Neetu Abad, and Christian Hinsch. 2011. A two-process view of Facebook use and relatedness need-satisfaction: Disconnection drives use, and connection rewards it. *Journal of Personality and Social Psychology* 100, 4: 766.

[79] Kennon M. Sheldon, Andrew J. Elliot, Youngmee Kim, and Tim Kasser. 2001. What is satisfying about satisfying events? Testing 10 candidate psychological needs. *Journal of personality and social psychology* 80, 2: 325. Retrieved May 23, 2016 from http://psycnet.apa.org/journals/psp/80/2/325/

[80] Maarten Vansteenkiste and Richard M. Ryan. 2013. On psychological growth and vulnerability: basic psychological need satisfaction and need frustration as a unifying principle. *Journal of psychotherapy integration* 23, 3: 263.

[81] Daniel S. Weld and Gagan Bansal. 2018. Intelligible artificial intelligence. *ArXiv e-prints, March 2018*.

[82] Robert W. White. 1959. Motivation reconsidered: The concept of competence. *Psychological review* 66, 5: 297.

[83] Rainer Winkler, Matthias Söllner, Maya Lisa Neuweiler, Flavia Conti Rossini, and Jan Marco Leimeister. 2019. Alexa, Can You Help Us Solve This Problem?: How Conversations With Smart Personal Assistant Tutors Increase Task Group Outcomes. In *Extended Abstracts of the 2019 CHI Conference on Human Factors in Computing Systems - CHI '19*, 1–6. https://doi.org/10.1145/3290607.3313090

[84] Ziang Xiao, Michelle X. Zhou, and Wat-Tat Fu. 2019. Who should be my teammates: Using a conversational agent to understand individuals and help teaming. In *Proceedings of the 24th International Conference on Intelligent User Interfaces*, 437–447.

[85] Xi Yang, Marco Aurisicchio, and Weston Baxter. 2019. Understanding Affective Experiences with Conversational Agents. In *Proceedings of the 2019 CHI Conference on Human Factors in Computing Systems - CHI '19*, 1–12. https://doi.org/10.1145/3290605.3300772

[86] Xi Yang, Marco Aurisicchio, James Mackrill, and Weston Baxter. 2017. On the products and experiences that make us happy. In *DS 87-8 Proceedings of the 21st International Conference on Engineering Design (ICED 17) Vol 8: Human Behaviour in Design, Vancouver, Canada, 21-25.08. 2017*, 499–508.