



**University of
Zurich**^{UZH}

**Zurich Open Repository and
Archive**

University of Zurich
University Library
Strickhofstrasse 39
CH-8057 Zurich
www.zora.uzh.ch

Year: 2022

Form Follows Function: Designing For Tensions Of Conversational Agents In Service Encounters

Bucher, Andreas ; Stähelin, Dario Luca ; Dolata, Mateusz ; Schwabe, Gerhard

Posted at the Zurich Open Repository and Archive, University of Zurich
ZORA URL: <https://doi.org/10.5167/uzh-227919>
Conference or Workshop Item
Accepted Version

Originally published at:

Bucher, Andreas; Stähelin, Dario Luca; Dolata, Mateusz; Schwabe, Gerhard (2022). Form Follows Function: Designing For Tensions Of Conversational Agents In Service Encounters. In: 30th European Conference on Information Systems (ECIS), Timisoara, Romania, 18 June 2022 - 24 June 2022, AISeL.

6-18-2022

Form Follows Function: Designing For Tensions Of Conversational Agents In Service Encounters

Andreas Bucher
University of Zurich, bucher@ifi.uzh.ch

Dario Staehelin
University of Zurich, staehelin@ifi.uzh.ch

Mateusz Dolata
University of Zurich, dolata@ifi.uzh.ch

Gerhard Schwabe
University of Zurich, schwabe@ifi.uzh.ch

Follow this and additional works at: https://aisel.aisnet.org/ecis2022_rp

Recommended Citation

Bucher, Andreas; Staehelin, Dario; Dolata, Mateusz; and Schwabe, Gerhard, "Form Follows Function: Designing For Tensions Of Conversational Agents In Service Encounters" (2022). *ECIS 2022 Research Papers*. 158.

https://aisel.aisnet.org/ecis2022_rp/158

This material is brought to you by the ECIS 2022 Proceedings at AIS Electronic Library (AISeL). It has been accepted for inclusion in ECIS 2022 Research Papers by an authorized administrator of AIS Electronic Library (AISeL). For more information, please contact elibrary@aisnet.org.

FORM FOLLOWS FUNCTION: DESIGNING FOR TENSIONS OF CONVERSATIONAL AGENTS IN SERVICE ENCOUNTERS

Research Paper

Bucher, Andreas, University of Zurich, Zurich, Switzerland, bucher@ifi.uzh.ch

Stahelin, Dario, University of Zurich, Zurich, Switzerland, stahelin@ifi.uzh.ch

Dolata, Mateusz, University of Zurich, Zurich, Switzerland, dolata@ifi.uzh.ch

Schwabe, Gerhard, University of Zurich, Zurich, Switzerland, schwabe@ifi.uzh.ch

Abstract

The proliferation of conversational agents (CAs) promises efficiency and quality improvements while enabling a more seamless integration of technology into service encounters. However, it remains unclear how CAs should be designed to provide the optimal experience for the key users: clients and frontline employees. Based on qualitative research with those key users, this study delivers a vision of an adaptable CA. It proposes a differentiated approach toward the design of CA: there is no "one-size-fits-all" design regarding the level of social presence, autonomy, or agency. The analysis reveals three tensions in user expectations leading to inconsistent design requirements for CAs. To resolve those tensions, CAs should be adapted to the changing context of a service encounter considering the appropriate level of autonomy, task complexity, interpersonal intimacy, and social role of the CA. The study contributes three design principles emphasizing the importance of the context for which a CA is designed.

Keywords: Conversational Agent, Social Presence, Financial Advice, Advisory Service Encounter

1 Introduction

Conversational agents (CAs) have gained popularity as consumer electronics products and their abilities are constantly increasing. CAs use improved speech recognition and natural language processing (NLP) to interact with users similar to human-to-human interaction and offer them access to conventional applications or artificial intelligence abilities. CAs not only have a variety of names such as chatbot, dialog system, or digital agent (Maedche et al. 2019), they also differ in terms of their characteristics and social cues, including their appearance, interaction modalities or behavior (Feine et al. 2019; Zierau et al. 2020). Research reports on the increasing use of such CAs in services and professional contexts, like financial advisory services (Dolata, Kilic, et al. 2019; Morana et al. 2020). As financial advisory services become increasingly complex due to changing legal regulations and customer requirements, CAs present a possible solution to support the advisor during the service encounter and to enhance the client experience (Dolata et al. 2019b). Contrary to current research, which has primarily focused on CAs that substitute contact with a human agent by handling simple, standardized requests (Gnewuch et al. 2017), CAs can allow for collaboration between advisors, clients, and CAs within the institutional framing of service encounters. While information systems (IS) research has recently attended to the potentials of employing CAs as a teammate (Seeber et al. 2020a) the contributions remain generic. The guidance on how to design CAs for use in human-to-human service encounters remains missing. Concurrently, more generic literature on CAs and their use in collaborative situations make inconclusive suggestions. Whereas such discourses as sociable robotics (Breazeal 2002) or computers-are-social-actors (Biolca et al. 2003; Nass and Moon 2000) suggest that CAs should ultimately exhibit high social presence and

resemble human beings, human-centered artificial intelligence (HAI) (Shneiderman 2020a), and phenomena like uncanny valley phenomenon (Mori et al. 2012) might suggest that lower social presence and limited agency might be a more promising direction. Advisory service encounters are characterized by an inherent asymmetry of information and interactional rights. Consequently, every designer investigating the potential use of CAs in advisory service encounters faces a complex, multidimensional design space and contradictory suggestions on how to move through this space. This study addresses this problem while exploring alternative designs of CAs for financial advisory services and addresses the research question:

RQ: *How should we design the appearance and behavior of a conversational agent for advisory service encounters?*

The exploration is embedded in a larger research project aimed at developing a CA for financial advisory services. Therein, the study provides detailed insights for the subsequent and iterative development of a software prototype of a CA based on comprehensive qualitative research with representatives of two Swiss banks and a survey with 49 subjects. First, it identifies three tensions that drive the informants' expectations: some informants seem to present two contradictory visions of the CA in a single interview. Second, the study resolves those tensions by suggesting that a CA adapts to the task and situation at hand during the advisory service encounter. Third, it derives design principles (DPs) to inform the design. The outcomes are discussed in the context of the literature showing that the tensions can be resolved if one relaxes the assumption that the CA needs to keep its social presence, agency, and look constant during the service. The research renders practical and scientific implications by revealing the potentials and challenges of creating CAs for human-to-human service encounters.

2 Related Work

2.1 Supporting advisory service encounters with IT

Advisory service encounters, or simply service encounters, are a unique form of collaboration as they involve two parties differing in their knowledge, incentives, goals, and interaction rights. Especially in commercial advisory services, like financial advisory services, the objectives and incentives might lie far apart. The advisee participates in a service encounter to find a solution to their problem or identify an offering that meets their needs. The advisor represents the service provider that uses the service encounter to market their solutions (Jungermann 1999; Oehler et al. 2010; Oehler and Kohlert 2009). While the service literature frames advisory service encounters as a marketing channel, communication science sees them as instances of institutional talk (Dolata et al. 2019b; Dolata and Schwabe 2017a; Heritage 2005). Thus, the parties involved in the conversation do not only enact their specific knowledge or service exchange interests but act according to their institutional identity. Concurrently, advisory service encounters often address very complex topics and require an intimate interaction between the advisor and the advisee with lots of sensitivity as to what can be asked or requested and what goes beyond the expectations of the conversation partner (Kilic et al. 2017). This poses challenges for the development of information technology (IT) for advisory service encounters. IT might emphasize or make implicit dynamics explicit (Kilic et al. 2017), simply disturb the natural flow of conversation (Kilic et al. 2016), or even impair non-verbal communication (Heinrich, Kilic, Aschoff, et al. 2014). Thus, despite making advice more efficient and unique, IT should not only cover the required functions of the encounter but should also integrate well into the natural flow of interaction between the advisor and advisee. Nevertheless, it remains open how this can be achieved. Also, research on supporting the collaboration between the advisor and the advisee remains a niche and rendered solutions primarily oriented at more effective visualization but only rarely oriented at the overall experience for the client and the advisor (Dolata et al. 2019a; Dolata and Schwabe 2016; Kilic et al. 2017).

We claim that using a CA might be a step in the right direction. A CA able to understand the conversation flow might identify the right moment for making an input. It might also refrain from showing and visualizing some information, which was reported to generate confusion (Kilic et al. 2016, 2017), and instead, make a verbal contribution that seems less persistent and explicit. It could also improve the

efficiency of the overall process by taking over documentation tasks and helping the participants remember core decisions taken during the service. Finally, combined with effective but seamless visualizations (Dolata et al. 2020a), a CA could enhance the overall perception of the service. Yet, it remains open how to design CAs that complement existing services rather than replacing the advisor. Very nascent research has emerged in IS (Dolata et al. 2019b; Seeber et al. 2020b), but clear guidance is still missing leaving practitioners and researchers lost in the complexity of advisory service encounters.

2.2 Conversational agents – smart tools or teammates?

The idea of developing a computer program that can elaborate with humans in natural conversations and that supports and complements human users in various situations has been of interest since the development of early implementations of CAs (Almog et al. 1989; Franchi and Güzeldere 1995). In contrast to other IS/IT artifacts, CAs do not only react passively and deterministic to user input, but also interact actively and probabilistic (Baird and Maruping 2021). To create a positive and meaningful user experience, the human user, including his behavior, beliefs, and psychological needs should be the center of developing CAs (Clark et al. 2019; Yang and Aurisicchio 2021). This approach to designing CAs is closely related to research on HAI. In the last years, HAI as a field of study has emerged from the widespread application of AI in industry and society (Xu 2019). HAI puts the human at the center of the design of AI systems to create AI systems that amplify, augment, and enhance human performance instead of replacing them (Bond et al. 2019). Generally, HAI addresses common aspects of AI research, including challenges related to AI safety, trustworthiness, or reliability, and explores technical questions related to the perception and understanding of natural language, behavior, and human mental models (Bond et al. 2019; Riedl 2019; Shneiderman 2020a). A fundamental aspect of HAI is the notion of autonomy and its influence on the way humans operate and control AI systems. As the primary objective of HAI is to design AI systems with the human in the center, it decouples autonomy and control to make it possible to simultaneously achieve high levels of human control and high levels of computer automation (Shneiderman 2020b). However, HAI uses a relatively abstract notion of AI and does not provide clear guidance on specific design considerations. Thus, it remains open how fundamental concepts of HAI can be achieved in complex, social situations involving conversations with multiple interlocutors and institutional or cultural contexts.

Research often draws on social presence theory to understand people's social perceptions, feelings, and behavior toward CAs (Gnewuch et al. 2017; Luria et al. 2019; Schmid et al. 2022). Social presence is the ability of a technology to transport social cues to create a sense of "being with another" (Biocca et al. 2003). Early works on social presence focused on subjective qualities of a communication medium that foster social presence in human-to-human interaction (Short et al. 1976). The concept of "being with another" was later also introduced to human-to-technology interaction (Biocca et al. 2003). Similarly, Nass and Moon (2000) showed that humans mindlessly react with social responses, like applying social rules, expectations, or stereotypes, to social cues of computers or virtual avatars. Various studies have further explored the effects of social technology cues on the perception of social presence when interacting with websites, CAs, or other forms of interfaces and technology (von der Pütten et al. 2010; Washington State University et al. 2009). Therein, it has been investigated and shown that the degree of social presence positively influences different aspects of human-to-technology interaction, such as trust-building, usefulness, or joy-of-use (Clark et al. 2019). Research on CAs found that higher degrees of social presence can support human-to-CA interaction in various fields, including education, customer service encounters, or human-AI teaming, as people perceive the interaction with the CA as more natural and enjoyable (Gnewuch et al. 2017; Winkler et al. 2020). In contrast, CAs with higher social presence may also increase discomfort in users, leading to reduced self-disclosure (Schuetzler et al. 2018).

The social presence of technology can be manipulated in a variety of ways, including the use of social cues, personality traits (Biocca et al. 2003; von der Pütten et al. 2010), or other anthropomorphic features (e.g., human-like embodiment) (Schuetzler et al. 2018). As social presence is often linked to the ability of a technology to transport these cues, it is also influenced by the richness of the communication medium itself (Daft and Lengel 1986; Hassanein and Head 2007). Gnewuch et al. (2017) developed several design principles focusing on conversational attributes such as the quality of the CAs responses, the quantity of shared information, the flexibility of conversational flows, or the ability to express social

cues. Feine et al. (2019) picked up the notion of social cues, who proposed a taxonomy of social cues for CAs, including verbal, visual, auditory, and invisible cues intended to issue social signals and trigger social reactions in the users. Similarly, Zierau et al. (2020) developed a taxonomy for the design of CAs from a user-experience-based perspective, which also exhibits aspects related to the behavior and appearance of the CA in human-to-CA interaction, including its embodiment, interaction style, interaction mode, or locus of control. Most research, including the mentioned taxonomies and studies, is focused on simple tasks, like customer service encounters that require a CA with a single social presence. However, in complex interactions, like in collaborative advisory sessions in the financial or medical domain, a single CA must perform and switch between multiple tasks and roles. For instance, a CA for financial advice may fulfill the role of the risk profiler, the portfolio manager, or the minute taker. The idea of a flexible social presence has been the subject of previous research, including studies on a dynamic embodiment of CAs (Duffy et al. 2003; Koay et al. 2016), on re-embodiment, and co-embodiment (Luria et al. 2019). Herein, Luria et al. (2019) showed that people prefer a single CA having multiple social presences for different tasks over multiple CAs with a single social presence for each task. Nonetheless, research on the flexibility of social presence has not discussed the mapping of task characteristics, like its complexity or intimacy, to aspects of social presence. Furthermore, the implementation of CAs with a continuum of social presence in complex conversations like service encounters remains unexplored.

3 Methods

The study is situated in a larger Design Science Research (DSR) project. The project develops a CA for bank advisory service encounters in Switzerland. It involves two universities, as well as two technology partners and two regional banks. The project aims at reducing the cognitive burdens of financial advice in an increasingly complex setting. Besides changing regulatory requirements, the advisors need to recommend relevant products while creating and maintaining a personal relationship with the client. The goal of the first project phase was to define a use case and develop prototypes for the CA. Herein, this study aims to explore the design considerations for a CA accepted by advisors and advisees and seamlessly integrated into the service encounter. This lays the foundation for the subsequent and iterative development of a software prototype of a CA for financial advisory services.

3.1 Data collection

This study is based on two different types of qualitative data to study the user's perspective on CAs in service encounters. As a first step, we conducted 10 exploratory interviews with experienced advisors and banking experts from the two participating banks in June 2020. All of the interviewees identified as male and were between 30 and 50 years old. We aimed at gathering rich insights into existing practices and reflections about changes in investment advice. After conducting seven interviews, we reached theoretical saturation (Saunders et al. 2018) because of informational redundancy and recurring statements. The advisors were incentivized and encouraged by their employers to participate in the interviews. The interviews were conducted online (i.e., Microsoft Teams, Skype, or Zoom) due to the COVID-pandemic and lasted about 90 minutes. The advisor interviews were divided into three parts: (1) general experience as client advisors, (2) advisors' current practices and their use of information technology in the preparation, encounter, and post-processing, (3) their vision of a service encounter supported by a CA. This study refers to the last part of the interview. This part used the narrative interview technique to avoid introducing bias in the advisors' accounts (Allen 2017). Even as the advisors did not interact with a prototype themselves, the gathered data is valuable in the context of this study as it provided the foundation for developing the prototypes that were later evaluated from the client's perspective.

Subsequent to the advisor interviews and to assess the client's perspective, 49 students evaluated CA-supported service encounters using a video-based survey from the client's point-of-view in March 2021. Most of the students were master's students, of whom 40 identified as male. The evaluation of CA-supported service encounters was part of their homework on dyadic collaboration in a computer-supported cooperative work (CSCW) course. To explore the most suitable social presence for a CA in service encounters, we designed two prototypes, Box and Pius, that differ in their social presence according to the taxonomy of Feine et al. (2019). Following Gibbs's reflective cycle (Gibbs 1988), the students

described the actions and their feelings during the encounter, and evaluated what worked and what did not work in the encounter. Further, the students compared Box and Pius regarding their impact on the encounter, the advisor, and the client and chose a favorite. We conducted the study using students as young people have an increasing interest in financial products, like stocks, and thus have seen the largest increase in new investors in the last years (Balonier et al. 2022; Kritikos et al. 2022). Also, prior research showed that students can emphasize well with financial advisory services and thus are valuable research subjects (Heinrich, Kilic, and Schwabe 2014). Hence, we believe that our students are appropriate participants for our study.

The two prototypes are designed to support the advisor in a variety of tasks (e.g., collecting the customer data, updating the investment portfolio or providing additional information and visualizations) and to improve the client's experience. They are distinguished based on their social cues and interaction style to manipulate social presence. Following, Box was designed to be as little socially perceptible as possible. A cylindrical Bluetooth speaker was chosen as the physical representation to resemble a geometric form similar to Google Nest. For the study, the speaker did not receive a name nor a voice to avoid humanization. For the sake of simplicity, this prototype is referred to as "Box" in this paper. Box was designed as a mono-modal CA that can display visualizations on the table. Interaction with Box was designed as "command & control", meaning that this CA could not process the conversational context and would only interact on specific commands similar to Amazon's Alexa (e.g., "Show me Amazon's stock over the last 6 months"). On the other hand, Pius was designed to achieve a very high social presence that should, ideally, create the sense of an additional person in the service encounter. Consequently, Pius has a name and a human-like avatar shown on a separate display on the table. It is designed as a man with brown hair wearing jeans and a shirt with a tie. Additionally, Pius can speak Swiss German dialect, mimic facial expressions (e.g., joy by smiling or unease by moving the eyebrows), and use gestures (e.g., waving or putting down a pen). The interaction with Pius was designed to fit naturally into a conversation between humans meaning that the CA interacts autonomously with both human users. In summary, Pius is a human-like CA that can use the same communication channels as the advisor and initiate interaction autonomously.

To provide an experience as realistic as possible, we recorded service encounters based on the workshops and interviews conducted during the contextual inquiry together with a financial advisor from one of the banking partners. The following paragraphs provide an overview of the four different phases, the targeted outcome, and the CA's task for each phase. The advisor welcomes the client in the *Welcome-Phase* and outlines the structure of the encounter. The CA either is introduced by the advisor (Box) or introduces itself (Pius). Next, the advisor asks the client to tell him about herself and the reason for attending the advisory session (e.g., date of birth, financial goals, and personal interests). The CA records the conversation and fills the corresponding fields on the screen with data. When the client asks not to record a specific answer, the recording is stopped and re-started by the advisor (Box) or automatically by the CA (Pius). Once all data is collected, the advisor asks the CA to save it. In the *Strategy-Phase*, the client fills out a regulatory required risk profile under the advisor's guidance, which is the basis for the investment strategy. For the discussion of the investment strategy, the CA displays the strategy and additional views (e.g., diversification by regions). Pius also explains the visualizations shortly, whereas the advisor provides all explanations in the encounter with Box. Finally, the client and advisor decide on the appropriate strategy for the investment portfolio. In the following *Investment-Phase*, the advisor and advisee further discuss and refine the investment strategy and the corresponding portfolio. Both CAs support this phase by providing visualizations and flexibly adapting the portfolio based on the client's interests. Again, Pius enriches the visualizations by providing additional explanations. Furthermore, they monitor the risk level and intervene either visually (Box) or auditorily (Pius) in case that the transactions exceed the client's risk profile. Finally, the advisor summarizes the discussion in the *Closure-Phase*. Only Pius says goodbye to the client.

3.2 Data analysis

The collected data was analyzed in three rounds. Each round consisted of data coding, a workshop between all co-authors to ensure the validity of the data analysis and the consultation of relevant literature. First, the transcribed interviews and responses to the survey were coded following a bottom-up approach

(Saldaña 2013) and focused on three pre-defined areas of interest: the form of the CA, interaction modalities, and tasks carried out by the CA. The open coding approach allowed for the deconstruction of the data into small, topically related units. The goal of the first round of coding was to elicit requirements for the designs of the final prototype. The analysis led to the observation that many statements contradict each other and that some contradictions even occur within the same interviews or responses. The results of the first round of coding were discussed in a workshop between the co-authors. The researchers defined 14 codes and 39 sub-codes based on design elements (especially Feine et al. 2019; Zierau et al. 2020) to understand the contradictions better.

The goal of the second round of analysis was to elaborate on the contradicting requirements. For this, the data was reviewed and coded to identify opposing accounts or conflicting statements leading to opposing implications for the design of a CA. The extracted statements were then grouped into categories that emerged during the coding, such as social presence, system capabilities, and control mechanisms. In the second workshop between the co-authors, the grouped statements were discussed and elaborated in the context of the literature. This workshop yielded three tensions regarding the desired behavior and appearance of the CA. Each of those tensions focuses on one of the initially defined areas of interest. In our context, tensions are conflicts in opposing ideas. Organizational literature views tensions as contradictory requirements that can only be resolved by compromises (Smith et al. 2017). Finally, the researchers conducted a third round of analysis to confirm the identified tensions. Accordingly, they focused on statements revealing the origin of the tensions and design principles that potentially resolve the identified tensions. In the third workshop, the co-authors reviewed the tensions presented in the results section and defined three design principles to solve them.

4 Results

The analysis uncovered that all advisors and students have a positive attitude toward CAs and recognize the added value for the service encounter. Still, they differ in framing what type of agent they want, how it should behave and look like, and how much influence they wish to exert over it. Sometimes, the same person expresses two contradictory wishes concerning those aspects. Those contradictions pointed us to the three most prominent tensions. Table 1 displays these tensions by summarizing and contrasting individual statements of advisors and students. In the following, we will refer to specific statements by providing the corresponding identifiers as listed in the table (e.g., C1).

4.1 Control vs. Autonomy Tension

The advisors and students express contradicting wishes regarding the control (i.e., sovereignty) over the interaction with the CA. In this study, control considers the authority of advisors over the overall structure of the service encounter and the interaction between the CA and humans (i.e., advisor and client). On the other hand, autonomy is the capability of the CA to interact with the user(s) independently.

Most advisors have a very clear opinion on the control of the interaction when imagining a CA-supported service encounter: All of them seem to prefer a user-driven interaction (C1). The advisors consider themselves as the primary user controlling the CA. However, some also see the client as an active user and mention that they would welcome the client to interact with the CA directly (e.g., by asking for information). We interpret that the advisors view CAs as a potential threat to the interaction in the service encounter. They value the highly personal interaction with the clients and often see their relationship as the main asset for a successful client relationship. The advisors are afraid that an autonomous CA participating in the service encounter could jeopardize the relationship between the human actors. Hence, the advisors prefer to be in control over the CA's thinking and actions. They imagine a command-based interaction to prevent the CA from interfering.

Interestingly, the advisors diverge from their opinion as soon as they embrace the idea of a CA-supported interaction. The advisors refer to a more autonomous CA for specific tasks (e.g., providing the latest news on stocks without request) and describe situations in which autonomous thinking and action of a CA would be appreciated (AT1). The advisors expand the CA's autonomy once they embrace the idea

of a CA-supported encounter that augments the discussion instead of challenging their authority. They imagine tasks beyond information search as they permit the CA to think and act more independently. However, Advisor 1 (**AT1**) mentions that he "*can show it directly to the client*," meaning that he still wants to decide whether to introduce the information directly or later in the discussion. Apparently, control and autonomy are not necessarily opposites. An appropriate combination could facilitate an active human-CA collaboration with potentially better overall performance.

The students draw a very similar image, especially in their reflection on the two encounters. In terms of control vs. autonomy, most students favored Box over Pius. This is mainly the case because Box is controlled by explicit commands and does not interact autonomously. Many students state that they felt the advisor sometimes lost control over the conversation in the encounter with Pius (**C2**). While Pius' autonomous interaction seems to disrupt the interaction between the client and advisor, it also felt more natural and fluid to the students. Many students state that Box's lack of independence was not a satisfying interaction. The students require the CA to register implicit needs to allow for a fluid and natural discussion enhanced with additional input at the right time without a disrupting request (**AT2**). As a result, the students mostly favor either Box for not actively interrupting the discussion, or Pius, for supporting the discussion due to its autonomous detection of implicit requests. Still, they are torn between control and autonomy. Additionally, the students expand on the role of the client (**AT3**), which the advisors only rarely mention. From the students' perspective, a CA can potentially contribute to a better customer experience if the client is allowed or encouraged to manipulate the system.

4.2 Dyad vs. Triad Tension

The analysis revealed paradoxical claims regarding the social cues of the CA, including its embodiment, visual representation, communication channels, and style. Advisors and students still hold on to the dyadic setting of the encounter, meaning that the CA should exhibit only a limited number of social cues and thus a lower social presence. However, both parties also express a positive attitude toward a present CA with many social cues, including the ability to communicate via different channels. Therein, many students perceived Pius as a third actor in the encounter, enhancing the discussion with its input.

In the interviews, the advisors express their favor for retaining a dyad in multiple ways. Regarding the embodiment of the CA, they advocate a discreet design and limit the CA to listening in the background and displaying information on a screen as the only output possibility (**D1**). Similarly, few advisors think about facial expressions or multi-modal communication combining visual and auditory output. More than half of the interviewed advisors even expect that a CA capable of human-like communication would distract from the personal interaction between the client and the advisor (**D2**). Although many advisors imagine the CA only with an artificial shape like a box or even without a physical representation, they experience limitations due to the limited social cues and social presence. Hence, they quickly attribute the CA with additional social cues. For example, many advisors would welcome it if the CA provided the latest information about a stock. Most advisors describe that the CA could display the stock chart and explain the latest developments. Some even see the CA engaging in inter-human communication such as greeting rituals and asking questions (**T1**). Even as many advisors see the advantage of a socially perceptible and present CA in specific tasks, they are reluctant to attribute the CA with a human-like representation that can mimic facial expressions or emotions.

The students agree mainly with the advisors. For them, a perceptible CA enhances the transparency and explainability of the processes performed by the CA. The students expect the CA to provide explanations for its actions. Therefore, they prefer Pius over Box as Pius can communicate with a human-like range of communication channels. They state that this made the interaction more pleasant as they expect the CA to react to interactions similar to inter-human interaction. Apparently, the missing "human-like" interaction bothered some of the students in the encounter with Box (**T2**). Many students welcome the strong social presence of Pius as they like the way the CA interacts and often describe a shift from a dyadic to a triadic setting. However, the students also raise privacy and transparency concerns about the interaction with a CA that has a high number of social cues. For example, during the Welcome-Phase, when the client asked to stop the recording, the students raise no privacy concerns in the encounter with Box (**D3**). Some even state that they feel confident that nothing is recorded as an "Off"-Sign was visible.

| | Advisors | Students (clients) |
|--------------------------------|---|---|
| 1) Control vs. Autonomy | | |
| Control | The advisors prefer control to allow for user-driven interaction: C1: <i>[The CA] must not interrupt the flow of the conversation. "Ping, new information". No, it has to be something that's running in the background so that I can call it up, me or the client, but not disturb the flow of the conversation. It shouldn't be dominant. IT should never dominate. (Advisor 1)</i> | Many students believe that an autonomous CA might cause the advisor to lose control over the conversation: C2: <i>At some point, I felt like Pius is doing all the work while the advisor seemed a bit useless. The advisor didn't come across as a knowledgeable person as it felt like he couldn't contribute anything additional to the consultation. As the advisor, I would have felt stupid and annoyed by the fact that Pius can just interrupt the talking. (Student 1)</i> |
| Autonomy | The advisors envision independent and autonomous actions of the CA for tasks that go beyond information search: AT1: <i>When I'm talking to a client, and I have to concentrate on how to explain to the [CA] that I would like to have the chart of the Apple stock, it is disruptive to the conversation. On the opposite, if it would register that in the background and show me the chart on a screen that I and the client can see, and I can show it directly to the client, that would be helpful. (Advisor 1)</i> | The students perceive that an autonomous CA fits more naturally into the conversation: AT2: <i>I think the main failing of [Box] is that it does not integrate naturally into the flow of the conversation. Other than in the data collection phase, the advisor needed to give explicit inputs to the assistant, which interrupted the flow of the conversation. The transitions between conversation and input were hard and abrupt. (Student 2)</i> The students believe that an autonomous CA empowers clients in the decision making: AT3: <i>[Pius] recognized which interests to remove as [the client] was talking about it, without having to wait on a command from the advisor, making the client feel more empowered in the decision-making. (Student 3)</i> |
| 2) Dyad vs. Triad | | |
| Dyad | The advisors advocate a discrete design with limited social cues and without human-like communication to retain a dyad between advisor and client: D1: <i>[the CA should be] modern, in the background, not disturbing, not interfering in the conversation. It should only listen. (Advisor 2)</i> D2: <i>I would imagine that in the future, I would only have a touchscreen table. And that the support would light up directly on it. Not as a person next to me; I don't need that. [...]. The more things we have on the table, the more likely [the client] gets distracted and the less focused he is on himself. (Advisor 3)</i> | Some students raise privacy concerns about the interaction with a CA with a high number of social cues: D3: <i>When the CA started to record, I was worried that [the client's] privacy would be hurt [...]. I positively noted the off-sign projected onto the paper when the recording was stopped. It ensured me that the recording was off and made me feel [the client's] privacy is protected. (Student 3, encounter with Box)</i> D4: <i>I still got the feeling that the conversation is being recorded even if the off-sign was shown. I could imagine that the client felt the same because Pius was now not a machine running in the background that can be switched on and off, but like a person part of the conversation and therefore like a person always listening. (Student 3, encounter with Pius)</i> |
| Triad | Most advisors envision a CA that is able to use a rich set of social cues and communication channels: T1: <i>[the CA] can introduce itself: "Hello, Ms. Müller, I'm the digital assistant, and today I'm supporting the advisor in this and that." This way, the client knows that someone here is not made of blood and flesh. If it is possible, [the CA] can also ask the client questions. (Advisor 4)</i> | Many students perceive the interaction with a CA that supports human-like communication as more pleasant: T2: <i>To me, it did not feel like [Box] naturally fit into the session. It is custom that a digital assistant that can respond to auditory prompts can also give auditory feedback (e.g., Alexa, Siri, etc.). However, this [CA] could only communicate through one channel. (Student 4)</i> |
| 3) Assistant vs. Expert | | |
| Assistant | The advisors prefer a CA in the role of an assistant: AS1: <i>It would be great if the system prepared certain things. If the schedule is already made. If specific work is being taken care of. It can do specific routine work and automate it. (Advisor 2)</i> | The students believe that a CA in the role of an expert might harm the advisor's competence: AS2: <i>The [CA] felt more like it was part of the conversation. But, it was not so clear what the institutional role of Pius was. He introduced himself as a helper and note-taker, but in the course of the session, he explained a lot of the financial concepts and answered questions of the client. This made the advisor seem less capable. (Student 5)</i> |
| Expert | For most advisors, a CA in the role of an expert appears to be better suited for complex matters: E1: <i>Beginning from the [strategy phase], almost in all places. Be it selecting an account type, all the way to the core-satellite approach (strategy), where it can pull things in, and I just stand next to it. From the [strategy phase] on, I can imagine a collaboration everywhere. (Advisor 5)</i> | The students believe that a team of advisor and CA might result in a more personalized experience: E2: <i>[The client] seemed comfortable during the session and appreciated both [the advisor's] and [the CAs] advice. She trusted in their consultation and believed that they had her best interest in mind. While she is no expert in investments, it felt like she had a solid grasp on the concepts and that the conversation was on an appropriate level of detail: This made me feel like her wishes, and personal interests were respected. (Student 6)</i> |

Table 1: Summary of individual statements of advisors and students corresponding to the identified tensions: 1) Control vs. Autonomy, 2) Dyad vs. Triad, and 3) Assistant vs. Expert

In contrast, the students question whether the recording stopped in the encounter with Pius (**D4**). In the same situation, students also mention Pius' high number of social cues and social presence to create an atmosphere in which they would not like to disclose personal details. They feel like the client was being questioned by two human beings, which intimidated her.

In conclusion, Box, which exhibits a low number of social cues and low social presence, puts the client and advisor in the spotlight by supporting their collaboration with appropriate visualizations without disturbing the relationship building. Yet, this complicates the interaction in some phases of the encounter and misses to establish transparency over the performed processes. As a socially present CA, Pius, on the other side, integrates better into the encounter and responds adequately in its interaction. The encounter is perceived as a triadic collaboration that enhances client-centricity. However, it creates an (inappropriately) high level of intimacy or might even intimidate the human users in situations where a dyad is better fitting.

4.3 Assistant vs. Expert Tension

The *Assistant vs. Expert Tension* concerns the role of the CA. It imposes the question of who the CA is and who it represents. Even though not asked directly, the advisors claim the CA immediately and implicitly as theirs. While this is superficially linked to the *Dyad vs. Triad Tension*, the origin of this claim lies deeper. IT in advisory service encounters is mainly seen as a tool to increase the efficiency of the advisor. Hence, more than half of the interviewed advisors focus on the potential support a CA could provide during the client interaction. The potential role includes almost exclusively time-intense tasks that require little to no advanced knowledge. For example, the CA should create minutes of the service encounter and fill out required regulatory forms. Many advisors also imagine the CA to register and prepare trades, so they only need to check and execute them in the post-processing of the encounter or want a CA to remind them in case that they forgot to address a pre-defined topic.

The general notion of the advisors is clearly toward an assistant or, more concrete, a tool that supports them in routine tasks (**AS1**). Apparently, the advisors feel threatened by the introduction of a CA. It might outperform them in certain areas or even make them redundant. Hence, they limit its scope to the role of an obedient secretary. However, the individual advisors contradict themselves again: many believe good advice-giving includes supporting clients individually and holistically. This includes identifying the client's implicit needs for other products (e.g., mortgages) or legal assistance (e.g., inheritance). As the encounter's complexity increases, many advisors admit not feeling confident in all areas. They try to identify the client's needs and direct them to appropriate specialists (e.g., a mortgage advisor). This is where almost all advisors see the potential for a CA. They expand its role from an assistant to an expert in investing, regulatory matters, and in the bank's products (**E1**). The statements of Advisor 2 (**AS1**) and Advisor 5 (**E1**) reveal the core of this tension. On the one hand, Advisor 2 only thinks about time-consuming tasks that he could delegate. The advisors believe that they are not really in need of a CA. On the other hand, Advisor 5 hands over the core of the service encounter to the CA. These paradoxical statements can be found in the majority of advisor interviews.

The students raise very similar questions as they reflect on the institutional role of the CA. They perceive both service encounters as very informative and competent. However, in the encounter with Box, the students attribute the competence solely to the advisor rather than to Box. In the encounter with Pius, perceived competence is divided between the advisor and CA, which, in some situations, leads to questioning the role distribution (**AS2**). While only a few students mention a negative impact of the CA (as in **AS2**), the majority welcome the idea of an intelligent CA. The students value the input from the CA and see the benefits of additional information and explanations. By providing information that the advisor does not have at hand, the CA carries out tasks that are not in conflict with the core competencies of the advisor. Thus, the advisor and CA are seen as a team complementing each other, resulting in a more personalized experience for the client (**E2**).

Overall, the students reflect on the effect of CAs on the role of the advisor and the impact on the service encounter itself. In the encounter with Pius, the students often talk about a change from a former dyad to a triad with the CA as a second advisor contributing to the discussion with its expertise. For some students, this leads to a shift of the expert role from the advisor to the CA, making him seem less

competent. This shift in competence is mainly mentioned when high personal interaction is required (e.g., explaining a strategy recommended to the client). In general, however, the students regard the advisor and CA as a highly competent team. In conclusion, the *Assistant vs. Expert Tension* arises from contradictory expectations regarding the role of the CA. While a CA in the role of a secretary does not question the advisor's competence, there is a slight improvement to the client's experience. A CA as an expert can enrich the encounter with beneficial information and eliminate the need for additional appointments to see other specialists. In turn, it might challenge the advisor's competence.

5 Discussion

Integrating a CA into advisory service encounters as seamlessly as possible while increasing the efficiency and uniqueness of these complex dyadic conversations is a challenge for practitioners and researchers as little guidance is available for the design of CAs for such settings. Prior studies have explored the effects of specific design choices, e.g., the personality or embodiment of a CA, on the interaction between *one* human and CA, and have proposed design guidelines and principles for the development of CAs (Feine et al. 2019; Gnewuch et al. 2017; Zierau et al. 2020). However, the integration of CAs into *dyadic* collaboration and its implications for the design of these CAs is less understood. Our study attends to answer the research question how a CA should look like and behave in an advisory service encounter. In our qualitative study with advisors and clients, we uncovered tensions that suggest revising the commonly used paradigm of designing a CA with a static social presence, i.e., keeping the autonomy, social role, or social cues static. The observed three tensions - *control vs. autonomy*, *dyads vs. triads*, and *assistant vs. expert* - come down to the same connotation that a CA with a static social presence is inappropriate to account for the complexity of advisory service encounters. For instance, the first tension *dyad vs. triad* primarily relates to the question of whether a CA should stay passively in the background, like a tool, or actively participate in the conversation as a third interlocutor. Existing taxonomies for designing CAs illustrate this general notion as they do not account for changing designs of CAs (Feine et al. 2019; Zierau et al. 2020). In the context of our results, these design principles appear to be insufficient to resolve the tensions.

Instead, research on dynamic embodiment (Duffy et al. 2003; Koay et al. 2016) and re- and co-embodiment of CAs (Luria et al. 2019) seems more suitable as it suggests a flexible social presence. However, as this line of research is still nascent and has not been expanded to advisory service encounters, it does not guide the handling of the observed tensions. In line with the notion of a flexible social presence (Luria et al. 2019), we conclude that addressing the tensions requires a break with the commonly used paradigm of designing CAs with a static social presence. Instead, the social presence of a CA should be flexible and context aware. In this sense, and to answer our research question, we suggest designing CAs for advisory service encounters and dyadic configurations based on the principle "*form follows function*" initially coined by architect Louis H. Sullivan (1896). In our context, *form* refers to those attributes and characteristics of a CA that establish its social presence, including the choice of interaction modes, its level of autonomy, and its specific role in the conversation. *Follow* relates to the adaptability and flexibility of these attributes resulting in different levels of social presence. Lastly, *function* considers the contexts in which a CA is embedded and the tasks and actions it supports. As a result, we derived three design principles (DP) that support designers and research to avoid the identified tensions and solve design challenges when integrating a CA in advisory service encounters. The design principles are inspired by Gregor et al. (2020) and follow the common strategy to manage tensions by resolving them over time (Poole and van de Ven 1989).

DPI: *Decouple control and autonomy of a CA for advisory service encounters to independently adapt their levels to the function of the CA:* Our results show that advisors and clients struggle with how much control the user should exert over the CA compared to the level of autonomy that the agent should exhibit. On the one side, too much control over the CA restricts its capabilities and hinders the natural conversation flow. On the other side, too much autonomy evokes the fear that the CA may act against the advisors' will, e.g., by providing opposing information or that the advisors may lose their face in the client's presence. This fear mainly originates from the notion that control over a CA and its autonomy are opposing characteristics. However, following the line of thought of Shneiderman (2020), control and autonomy are non-exclusive characteristics that complement each other and can lead to increased

performance of human-AI collaboration. Hence, decoupling control and autonomy allows for highly controllable and autonomous CAs.

Following the principle "*form follows function*", designers should adapt the CA's level of control and autonomy to the specific situation or task when designing a CA for advisory service encounters. For example, as some advisors mentioned, an autonomous CA could be supportive if it actively participates in the discussion about the investment portfolio by autonomously providing additional information or recommendations. However, in other, more intimate situations, a non-autonomous CA who only listens to the advisor's instructions appears more suitable. Nonetheless, mechanisms that allow the advisor to control and correct the output of the CA could still ensure high control in both situations. When defining the level of control and autonomy, they must be aware of its implications on the CA's social presence. Such design requires careful considerations regarding the fit between the level of autonomy and the level of social presence as the level of autonomy influences the social presence of a technology (Slater 2009). Nonetheless, we expect that the tension between control and autonomy can be resolved by decoupling and adapting the CA's levels of control and autonomy to specific situations and tasks.

DP2: *Situationally adjust the social cues of a CA for advisory service encounters to account for different levels of intimacy and task complexity:* In our study, advisors and clients were torn between a CA that does not change the dyadic nature of an advisory session by staying passively in the background and a CA that actively participates in the conversation resulting in a triadic collaboration. From another angle, this tension relates to the notion of a CA with only a few social cues (e.g., a CA that can interact with its users solely through visualizations, Box) or a highly present CA with a high number of social cues (e.g., a CA that additionally can communicate in natural language, and exhibits bodily movements and other social behavior, Pius). Whereas Box was perceived as a supporting tool by advisors and clients, Pius was seen as an additional participant that can interact with the users resulting in a more natural discussion. Initially, most advisors favored a CA that behaves passively and displays information only visually on the screen. However, the advisors also acknowledged that a perceptible CA could be supportive in certain situations (e.g., when providing helpful information about investment products). In such situations, a CA is allowed to interact with the client by asking her questions. A CA that can offer explanations about its actions was also important to clients, who missed the transparency provided by Pius when evaluating Box.

Nonetheless, Pius also implicitly created expectations in the clients. For example, in situations requiring high confidentiality and intimacy, the clients did not trust the CA when Pius was advised to stop listening. In contrast, our data show that clients did not have these reservations regarding Box in the same situation. The dichotomy between a CA that retains the dyadic setting of the advisory service encounter and a CA that is perceived as an additional participant, thus creating a triad, exists not only on the advisor's but also on the clients' side. Most research has focused on simple 1:1 communication between humans and a CA (Gnewuch et al. 2017, Luria et al. 2019). Hence, the question, whether a CA can influence the interaction mode did not occur (e.g., by changing a conversation from a dyad to a triad). Keeping the social cues and the CA's social presence static through the entire advisory session appears to be insufficient.

Instead, we propose that designers match the social cues of a CA to specific situations of the service encounter. On the one hand, designers have to consider the level of intimacy and self-disclosure required in certain situations when selecting the CA's social cues for service encounters. Schuetzler et al. (2018) show that people disclose less personal information when conversing with a CA with a high number of social cues and thus show a higher level of social presence. Similarly, our results show that the clients were more concerned about sharing their personal details in the encounter with Pius. On the other hand, CA designers need to account for the complexity of the task when selecting the social cues of a CA. Consequently, they should consider how suitable the social cues of a CA are to transport the content of information given the equivocality or uncertainty of a task (Daft and Lengel 1986). This suggests that tasks with high equivocality or uncertainty are better supported by a CA that possesses a rich set of social cues. This argumentation is in line with our results, which show that a highly social CA that actively engages in the conversation would be better suited for explaining complicated investment proposals than a non-social CA.

DP3: *Adapt the social roles of a CA for advisory service encounters based on its specific functions to align the CA's perception and abilities with the requirements of a specific task or context:* Lastly, our data shows that integrating a CA into a dyadic setting creates a conflict between a CA in the role of an assistant and an expert. Whereas the advisors were afraid that a CA expert harms their authority and competence, they also believe that such a CA is valuable in specific situations. Similarly, the clients perceived a CA in an expert's role as reliable when offering portfolio recommendations and explanations. Nonetheless, when integrating a CA expert in the service encounter, the clients felt the competence was distributed among the advisor and the CA, leaving the impression that the advisor is less capable. The perception that the CA's specific role affects certain aspects of the conversation, like the perception of the advisor's capabilities, is not surprising and has been subject to previous research. According to Hinds et al. (2004), the presentation of a CA's role affects the user's willingness to rely on the CA. For instance, people take on more responsibility for a task if a robotic agent is more machine-like and is assigned to the role of a subordinate (Hinds et al. 2004).

The role of a CA in dyadic settings should thus be considered quite carefully in the design. Hence, solving the tension regarding the different roles of a CA is just as important as solving the previous tensions. With regard to the principle "*form follows function*", we propose that designers should adapt the social role of a CA depending on the requirements of a specific task. According to the research on dynamic and re-embodiment of CAs (Duffy et al. 2003; Luria et al. 2019), one CA with multiple roles is more natural and efficient than a situation-specific interaction with multiple CAs with static social roles. For example, the CA in our advisory service encounter could be assigned to the role of a subordinate during tasks that require only one expert (e.g., relationship-building at the beginning of an encounter). In situations where the CA acts as an investment expert, the CA would be perceptible as such. This reduces the ambiguities regarding the CA's role. However, as social roles are tightly linked to the perception of being a social actor (Steuer 1995), designers need to be aware of the implications of a social role on the CA's social presence. As a result, designers have to adapt the role of the CA to a specific situation and have to ensure that the social cues of the CA fit its social role.

6 Conclusion, limitations & future research

When designing a CA for collaborative situations, including advisory service encounters, one gets easily overwhelmed by a multidimensional design space and contradictory suggestions. Therein, we identified three tensions related to the autonomy, social cues, and social role of a CA in advisory service encounters that originate from keeping the CA's social presence static. To relax the contradictory tendencies of these tensions, we suggest breaking with the commonly used paradigm of designing a CA with a static social presence. Instead, we propose that the design of a CA in advisory service encounters should be based on the principle "*form follows function*," allowing for a flexible social presence, including aspects like the autonomy, social cues, and social role of a CA, which can be adapted to specific situations and tasks.

This work builds on very nascent research of complementing advisory service encounters with CAs and is a first attempt to understand the implications that arise for the design of such CAs. Therefore, our study contributes not only to the body of knowledge about the design of CAs, but also adds to the discourse on advisory service encounters and machines-as-teammates. Therein, designers benefit from a more nuanced perspective on CAs. Instead of attempting to make one ultimate decision about the design of the CA, they can dissect the problem into several decisions to be taken for specific tasks or phases of the service encounter. They are encouraged to reconsider and reconcile their views on autonomy, social presence, and roles of a CA. Banks and other service providers obtain insights into the potentials of CAs in service encounters and learn about the opinion of employees and clients. By pointing toward the potentials and challenges of CAs in this context, our study contributes to research on the application of IT in advisory services (Dolata et al. 2020a). Finally, our research extends the discourse on machines-as-teammates in IS and HCI (Seeber et al. 2020a) by providing a systematic and qualitatively rich comparison of several approaches toward the design of CAs for use in human-to-human service encounters. Researchers attending to this discourse also receive an inspiration to reconsider and reconcile their views upon the common ground of adaptivity.

Because of the exploratory nature and the early stage of our research, we have to rely on qualitative data from the interviews and survey instead of testing different design aspects and propositions for a CA in a real advisory service encounter. Furthermore, our study did not consider privacy and data protection requirements, but rather focussed on the design aspects of the CA and assumed that data is readily accessible. Future research should look into two directions: (1) The design principles can be further grounded in theory. For example, it may not be the 'task' determining the autonomy, social presence, and role of the CA (as implied, e.g., by Media Richness theory (Daft and Lengel 1986)), but rather the generic communication process (following the insights of the theory of media synchronicity (Dennis et al. 2008)) or the generic patterns of collaboration (Briggs et al. 2016) that determine them. (2) The form of the CA in each situation may not need to be hardcoded into the software but may have some flexibility grounded in the CAs 'understanding' of the situation and the advisor's preference. Such adaptability can range from small, predetermined changes in autonomy and social presence to dynamically 'learned' large-scale adaptation of the CAs form.

References

- Allen, M. 2017. *The SAGE Encyclopedia of Communication Research Methods*, 2455 Teller Road, Thousand Oaks California 91320: SAGE Publications, Inc.
- Almog, J., Perry, J., Wettstein, H. K., and Kaplan, D. (eds.). 1989. *Themes from Kaplan*, New York: Oxford University Press.
- Baird, A., and Maruping, L. M. 2021. "The Next Generation of Research on IS Use: A Theoretical Framework of Delegation to and from Agentic IS Artifacts," *MIS Quarterly* (45:1), *MIS Quarterly*, pp. 315–341.
- Balonier, S., Di Dio, D., and Fey, G. 2022. *Deutschland Und Die Aktie - Weiter Auf Hohem Niveau*, Deutsches Aktieninstitut.
- Biocca, F., Harms, C., and Burgoon, J. K. 2003. "Toward a More Robust Theory and Measure of Social Presence: Review and Suggested Criteria," *Presence: Teleoperators and Virtual Environments* (12:5), pp. 456–480.
- Bond, R. R., Mulvenna, M., and Wang, H. 2019. *Human Centered Artificial Intelligence: Weaving UX into Algorithmic Decision Making*, p. 8.
- Breazeal, C. L. 2002. *Designing Sociable Robots*, Intelligent Robots and Autonomous Agents, Cambridge, Mass: MIT Press.
- Clark, L., Pantidi, N., Cooney, O., Doyle, P., Garaialde, D., Edwards, J., Spillane, B., Gilmartin, E., Murad, C., Munteanu, C., Wade, V., and Cowan, B. R. 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*, Glasgow Scotland Uk: ACM, May 2, pp. 1–12.
- Daft, R. L., and Lengel, R. H. 1986. "Organizational Information Requirements, Media Richness and Structural Design," *Management Science* (32:5), pp. 554–571.
- Dennis, Fuller, and Valacich. 2008. "Media, Tasks, and Communication Processes: A Theory of Media Synchronicity," *MIS Quarterly* (32:3), p. 575.
- Dolata, M., Agotai, D., Schubiger, S., and Schwabe, G. 2019. "Pen-and-Paper Rituals in Service Interaction: Combining High-Touch and High-Tech in Financial Advisory Encounters," *Proc. ACM Hum.-Comput. Interact.* (3:CSCW), 224:1-224:24.
- Dolata, M., Agotai, D., Schubiger, S., and Schwabe, G. 2020. "Advisory Service Support That Works: Enhancing Service Quality with a Mixed-Reality System," *Proceedings of the ACM on Human-Computer Interaction* (4:CSCW2), 120:1-120:22.
- Dolata, M., Kilic, M., and Schwabe, G. 2019. "When a Computer Speaks Institutional Talk: Exploring Challenges and Potentials of Virtual Assistants in Face-to-Face Advisory Services," in *Proc. Hawaii Intl. Conf. on System Sciences*, January 8, 2019.
- Dolata, M., and Schwabe, G. 2016. "More Interactivity with IT Support in Advisory Service Encounters?," in *Proc. Conf. Mensch Und Computer*.

- Dolata, M., and Schwabe, G. 2017. "Paper Practices in Institutional Talk: How Financial Advisors Impress Their Clients," *Computer Supported Cooperative Work (CSCW)*, pp. 769–805.
- Dolata, M., Steigler, S., Nüesch, F., Schock, U., Agotai, D., Schubiger, S., Kilic, M., and Schwabe, G. 2019. "Welcome, Computer! How Do Participants Introduce a Collaborative Application during Face-to-Face Interaction?," in *Proc. Intl. Conf. on Human-Computer Interaction - INTERACT*.
- Duffy, B. R., O'Hare, G. M. P., Martin, A. N., Bradley, J. F., and Schon, B. 2003. "Agent Chameleons: Agent Minds and Bodies," in *Proceedings 11th IEEE International Workshop on Program Comprehension*, New Brunswick, NJ, USA: IEEE Comput. Soc, pp. 118–125.
- Ehsan, U., and Riedl, M. O. 2020. "Human-Centered Explainable AI: Towards a Reflective Sociotechnical Approach," in *HCI International 2020 - Late Breaking Papers: Multimodality and Intelligence* (Vol. 12424), Lecture Notes in Computer Science, C. Stephanidis, M. Kurosu, H. Degen, and L. Reinerman-Jones (eds.), Cham: Springer International Publishing, pp. 449–466.
- Feine, J., Gnewuch, U., Morana, S., and Maedche, A. 2019. "A Taxonomy of Social Cues for Conversational Agents," *International Journal of Human-Computer Studies* (132), pp. 138–161.
- Franchi, S., and Güzeldere, G. 1995. *Constructions of the Mind Artificial Intelligence and the Humanities.*, Stanford Humanities Review, California Stanford University Press.
- Gibbs, G. 1988. "Learning by Doing: A Guide to Teaching and Learning Methods," *Further Education Unit*.
- Gnewuch, U., Morana, S., and Maedche, A. 2017. "Towards Designing Cooperative and Social Conversational Agents for Customer Service," in *Proceedings of the International Conference on Information Systems (ICIS) 2017*, South Korea, p. 15.
- Gregor, S., Chandra Kruse, L., and Seidel, S. 2020. "Research Perspectives: The Anatomy of a Design Principle," *Journal of the Association for Information Systems* (21), pp. 1622–1652.
- Hassanein, K., and Head, M. 2007. "Manipulating Perceived Social Presence through the Web Interface and Its Impact on Attitude towards Online Shopping," *International Journal of Human-Computer Studies* (65:8), pp. 689–708.
- Heinrich, P., Kilic, M., Aschoff, F.-R., and Schwabe, G. 2014. "Enabling Relationship Building in Tablet-Supported Advisory Settings," in *Proc. Conf. Computer Supported Cooperative Work and Social Computing*, Baltimore, MD, USA: ACM Press, pp. 171–183.
- Heinrich, P., Kilic, M., and Schwabe, G. 2014. "Microworlds as the Locus of Consumer Education in Financial Advisory Services," in *ICIS 2014 Proceedings: IS Design Science*, Association for Information Systems.
- Heritage, J. 2005. "Conversation Analysis and Institutional Talk," *Handbook of Language and Social Interaction*, pp. 103–147.
- Hinds, P. J., Roberts, T. L., and Jones, H. 2004. "Whose Job Is It Anyway? A Study of Human-Robot Interaction in a Collaborative Task," *Human-Computer Interaction* (19), pp. 151–181.
- Jungermann, H. 1999. "Advice Giving and Taking," in *Proc. Hawaii Intl. Conf. System Sciences. HICCS-32*.
- Kilic, M., Dolata, M., and Schwabe, G. 2016. "How IT-Artifacts Disturb Advice Giving – Insights from Analyzing Implicit Communication," in *Proc. Hawaii Intl. Conf. System Sciences*, , January, pp. 878–887.
- Kilic, M., Dolata, M., and Schwabe, G. 2017. "Why Do You Ask All Those Questions? Supporting Client Profiling in Financial Service Encounters," in *Proc. Hawaii Intl. Conf. System Sciences*, Waikoloa Beach, HI, USA.
- Kilic, M., Heinrich, P., and Schwabe, G. 2015. "Coercing into Completeness in Financial Advisory Service Encounters," in *Proc. Intl. Conf. Computer-Supported Cooperative Work and Social Computing*, Vancouver, BC, Canada: ACM Press, pp. 1324–1335.
- Koay, K. L., Syrdal, D. S., Ho, W. C., and Dautenhahn, K. 2016. "Prototyping Realistic Long-Term Human-Robot Interaction for the Study of Agent Migration," in *2016 25th IEEE International Symposium on Robot and Human Interactive Communication (RO-MAN)*, New York, NY, USA: IEEE, August, pp. 809–816.
- Kritikos, A., Handrich, L., Priem, M., and Morales, O. 2022. *Hype or New Normal? Insights into the Motives and Behavior of a New Generation of Investors*, DIW Econ GmbH.

- Luria, M., Reig, S., Tan, X. Z., Steinfeld, A., Forlizzi, J., and Zimmerman, J. 2019. "Re-Embodiment and Co-Embodiment: Exploration of Social Presence for Robots and Conversational Agents," in *Proceedings of the 2019 on Designing Interactive Systems Conference*, San Diego CA USA: ACM, June 18, pp. 633–644.
- Maedche, A., Legner, C., Benlian, A., Berger, B., Gimpel, H., Hess, T., Hinz, O., Morana, S., and Söllner, M. 2019. "AI-Based Digital Assistants: Opportunities, Threats, and Research Perspectives," *Business & Information Systems Engineering* (61:4), pp. 535–544.
- Martin, A., O'Hare, G. M. P., Duffy, B. R., Schön, B., and Bradley, J. F. 2005. "Maintaining the Identity of Dynamically Embodied Agents," in *Intelligent Virtual Agents* (Vol. 3661), Lecture Notes in Computer Science, T. Panayiotopoulos, J. Gratch, R. Aylett, D. Ballin, P. Olivier, and T. Rist (eds.), Berlin, Heidelberg: Springer Berlin Heidelberg, pp. 454–465.
- Morana, S., Gnewuch, U., Jung, D., and Granig, C. 2020. *The Effect of Anthropomorphism on Investment Decision-Making with Robo-Advisor Chatbots*, presented at the European Conference on Information Systems (ECIS), Marrakesh, Morocco.
- Mori, M., MacDorman, K., and Kageki, N. 2012. "The Uncanny Valley [From the Field]," *IEEE Robotics & Automation Magazine* (19:2), pp. 98–100.
- Nass, C., and Moon, Y. 2000. "Machines and Mindlessness: Social Responses to Computers," *Journal of Social Issues* (56:1), pp. 81–103.
- Oehler, A., and Kohlert, D. 2009. "Financial Advice Giving and Taking—Where Are the Market's Self-Healing Powers and a Functioning Legal Framework When We Need Them?," *Journal of Consumer Policy* (32:2), pp. 91–116.
- Oehler, A., Kohlert, D., Jungermann, H., Reisch, L., and Micklitz, H.-W. 2010. "The Quality of Financial Investment Advice for Private Investors: Problems in the Advice Process and Potential Solutions," *Statement by the Scientific Advisory Council on Consumer and Food Policy at the Federal Ministry for Food*.
- Poole, M. S., and van de Ven, A. H. 1989. "Using Paradox to Build Management and Organization Theories," *The Academy of Management Review* (14:4), p. 562.
- von der Pütten, A. M., Krämer, N. C., Gratch, J., and Kang, S.-H. 2010. "'It Doesn't Matter What You Are!' Explaining Social Effects of Agents and Avatars," *Computers in Human Behavior* (26:6), pp. 1641–1650.
- Riedl, M. O. 2019. "Human-centered Artificial Intelligence and Machine Learning," *Human Behavior and Emerging Technologies* (1:1), pp. 33–36.
- Saldaña, J. 2013. *The Coding Manual for Qualitative Researchers*, (2nd ed.), Los Angeles: SAGE.
- Saunders, B., Sim, J., Kingstone, T., Baker, S., Waterfield, J., Bartlam, B., Burroughs, H., and Jinks, C. 2018. "Saturation in Qualitative Research: Exploring Its Conceptualization and Operationalization," *Quality & Quantity* (52:4), pp. 1893–1907.
- Schmid, D., Staehelin, D., Bucher, A., Dolata, M., and Schwabe, G. 2022. "Does Social Presence Increase Perceived Competence?: Evaluating Conversational Agents in Advice Giving Through a Video-Based Survey," *Proceedings of the ACM on Human-Computer Interaction* (6:GROUP), pp. 1–22.
- Schuetzler, R. M., Giboney, J. S., Grimes, G. M., and Nunamaker, J. F. 2018. "The Influence of Conversational Agent Embodiment and Conversational Relevance on Socially Desirable Responding," *Decision Support Systems* (114), pp. 94–102.
- Seeber, I., Bittner, E., Briggs, R. O., de Vreede, T., de Vreede, G.-J., Elkins, A., Maier, R., Merz, A. B., Oeste-Reiß, S., Randrup, N., Schwabe, G., and Söllner, M. 2020. "Machines as Teammates: A Research Agenda on AI in Team Collaboration," *Information & Management* (57:2), p. 103174.
- Seeber, I., Waizenegger, L., Seidel, S., Morana, S., Benbasat, I., and Lowry, P. B. 2020. "Collaborating with Technology-Based Autonomous Agents: Issues and Research Opportunities," *Internet Research* (30:1), Emerald Publishing Limited, pp. 1–18.
- Shneiderman, B. 2020a. "Design Lessons From AI's Two Grand Goals: Human Emulation and Useful Applications," *IEEE Transactions on Technology and Society* (1:2), pp. 73–82.

- Shneiderman, B. 2020b. “Bridging the Gap Between Ethics and Practice: Guidelines for Reliable, Safe, and Trustworthy Human-Centered AI Systems,” *ACM Transactions on Interactive Intelligent Systems* (10:4), pp. 1–31.
- Shneiderman, B. 2020c. “Human-Centered Artificial Intelligence: Reliable, Safe & Trustworthy,” *International Journal of Human–Computer Interaction* (36:6), pp. 495–504.
- Short, J., Williams, E., and Christie, B. 1976. *The Social Psychology of Telecommunications*, London ; New York: Wiley.
- Smith, W., Erez, M., Jarvenpaa, S., Lewis, M. W., and Tracey, P. 2017. “Adding Complexity to Theories of Paradox, Tensions, and Dualities of Innovation and Change: Introduction to Organization Studies Special Issue on Paradox, Tensions, and Dualities of Innovation and Change,” *Organization Studies* (38:3–4), pp. 303–317.
- Steuer, J. 1995. “Vividness and Source of Evaluation as Determinants of Social Responses toward Mediated Representations of Agency,” Doctoral Dissertation, Doctoral Dissertation, Stanford University.
- Sullivan, L. H. 1896. *The Tall Office Building Artistically Considered*, (Vol. 339), Lippincott’s Monthly Magazine, Philadelphia : J.B. Lippincott Co.
- Washington State University, Hess, T., Fuller, M., University of Massachusetts, Amherst, Campbell, D., and Millsaps College. 2009. “Designing Interfaces with Social Presence: Using Vividness and Extraversion to Create Social Recommendation Agents,” *Journal of the Association for Information Systems* (10:12), pp. 889–919.
- Winkler, R., Hobert, S., Salovaara, A., Söllner, M., and Leimeister, J. M. 2020. “Sara, the Lecturer: Improving Learning in Online Education with a Scaffolding-Based Conversational Agent,” in *Proceedings of the 2020 CHI Conference on Human Factors in Computing Systems*, Honolulu HI USA: ACM, April 21, pp. 1–14.
- Xu, W. 2019. “Toward Human-Centered AI: A Perspective from Human-Computer Interaction,” *Interactions* (26:4), pp. 42–46.
- Yang, X., and Aurisicchio, M. 2021. *Designing Conversational Agents: A Self-Determination Theory Approach*, presented at the 2021 ACM Conference on Human Factors in Computing Systems, Yokohama, Japan, p. 27.
- Zierau, N., Wambsganss, T., Janson, A., Schöbel, S., and Leimeister, J. M. 2020. *The Anatomy of User Experience with Conversational Agents: A Taxonomy and Propositions of Service Clues*, p. 18.