

Where is the Bot in our Team? Toward a Taxonomy of Design Option Combinations for Conversational Agents in Collaborative Work

Eva A. C. Bittner
University of Hamburg
bittner@informatik.uni-hamburg.de

Sarah Oeste-Reiß
Kassel University
oeste-reiss@uni-kassel.de

Jan Marco Leimeister
University of St. Gallen
JanMarco.Leimeister@unisg.ch

Abstract

With rapid progress in machine learning, language technologies and artificial intelligence, conversational agents (CAs) gain rising attention in research and practice as potential non-human teammates, facilitators or experts in collaborative work. However, designers of CAs in collaboration still struggle with a lack of comprehensive understanding of the vast variety of design options in the dynamic field. We address this gap with a taxonomy to help researchers and designers understand the design space and the interrelations of different design options and recognize useful design option combinations for their CAs. We present the iterative development of a taxonomy for the design of CAs grounded in state of the art literature and validated with domain experts. We identify recurring design option combinations and white spots from the classified objects that will inform further research and development efforts.

1. Introduction

Collaboration has proven itself as an advantageous form of work in a great variety of domains, especially when it comes to complex tasks that exceed the capabilities of any individual. With the rapid advances in e.g. machine learning, artificial intelligence and language technologies, new potentials open up to bring information technology in collaboration not only as support systems [6, 40], but as autonomously acting conversational agents (CAs) that e.g. guide users in a familiar, conversation based way or react dynamically to the progress of a collaborative work practice [33]. CAs are virtual characters with a certain degree of intelligence, “which converse with humans by using natural language” [37]. A CA can enhance existing collaborative work practices and open up potentials for new forms of collaboration with humans. The behavior of CA relies on complex and interrelated technological and design decisions. Developing CAs for collaboration requires a sophisticated understanding of knowledge from e.g. human-machine interaction, collaboration research, machine learning and artificial intelligence. Against that

background little is known about the design of collaborative work practices, in which a peer, expert or the facilitator is a CA. Thus, in our paper we address the following question: *How can knowledge on CAs for collaborative work practices be classified in a way that helps designers and researchers to better understand useful design option combinations and to develop CAs for collaborative work practices grounded in the diverse knowledge in the field?* We address this question by creating a taxonomy for the design of CAs in collaborative work. Taxonomies are a useful approach to understand and analyze complex domains. They provide structure and organization to the domain knowledge [24]. Thus, the aim of our paper is to develop a taxonomy that helps researchers and designers recognize useful design option combinations of CAs. The taxonomy should be used a) *by designers of CAs* as a guiding scheme for the informed design of CAs for human-agent collaboration; and b) *by researchers on CAs* and collaborative work practices to identify design patterns and white spots for further exploration and theorizing. Our paper is structured as follows: In section 2 we refer to the theoretical background. In section 3 we describe our methodology to develop and evaluate our taxonomy in 2 iterations. In section 4 we present the taxonomy and define its dimensions and characteristics. In section 5 we discuss implications for the taxonomy usage and extract recurring design option combinations before we conclude.

2. Theoretical Background

2.1 Conversational Agents

CAs are “computer programs that interact with humans using natural languages” [34] and their goal is to simulate human conversation [34]. Lieberman describes an agent as a computer program that can be considered as an assistant for users [21]. In different domains, specific types of CAs are also investigated under different terms such as e.g. collaborative agents [1, 31] or chatbots/chatterbots [16, 27, 44]. Although the idea of CAs is not new [9], advances in artificial intelligence, natural language processing and cognitive systems have leveraged their abilities and boosted their application

potential [7]. As people are most familiar to use natural language for human communication, this mode is also accepted in human computer interaction as a powerful way for interacting with computers by using spoken or written language or natural body language [46]. In their analysis of the current state of multiparticipant chat research, Uthus and Aha [42] identify the participation of artificial agents in chat conversations, such as bots in chat rooms or non-player characters in video games as an underexplored area of research. They particularly stress the importance of learning to design CAs that can participate in chat conversations, and analyzing conversations, with some non human users. As with improvement of text-to-speech and speech-to-text conversion, spoken conversation with agents has become more common and feasible, we extend the scope to agents with different natural language communication modes.

2.2 Collaboration Engineering (CE) Research

When it comes to designing CAs for collaboration, an understanding of collaborative work practices among humans is needed. A collaborative work practice is a series of re-usable collaborative activities performed by multiple teammates (practitioners) to achieve a goal [18]. CE provides insights about design choices that need to be taken to develop re-usable collaborative work practices. When aiming to design CAs in collaborative work, the way of thinking that is represented by in the Six-Layer-Model of Collaboration [5] provides insights of the anatomy of a collaborative work practice [43]. It gives an overview of the areas of design concerns that need to be addressed [5]. The result is a collaborative work practice that packages facilitation expertise so that non-collaboration experts can use it without training in tools or techniques. It guides participants by inter alia setting process restrictions, diagnosing communication issues, stimulating reasoning, providing information access, ensuring goal congruence, preventing from distraction [5]. In most cases, so far, there is a human facilitator involved that interacts with human participants. Recent research has made initial steps into a new direction by packaging the “facilitator in a box” [6] and shows, that systems can take over parts of facilitation tasks [6]. New CE research opens and discusses opportunities to design and use CAs in collaborative work practices [33]. We build on this work to explore the design space for CAs in collaboration.

3. Methodology

To develop and evaluate our taxonomy, we follow the method for taxonomy development suggested by [24]. A taxonomy T is defined as a set of n dimensions D_i ($i = 1, \dots, n$) each consisting of k_i ($k_i \geq 2$) exclusive

and collectively exhaustive characteristics C_{ij} ($j=1, \dots, k_i$) [24], i.e. each object has only one C_{ij} for each D_i . $T = \{D_i, i = 1, \dots, n \mid D_i = \{C_{ij}, j=1, \dots, k_i; k_i \geq 2\}\}$ Following, we refer to the method’s steps and describe how they guided us toward creating our taxonomy.

Step 1 meta-characteristic: The objective of this step is to identify the most comprehensive characteristic as the basis for deriving subsequent characteristics. It is aligned with the purpose of the taxonomy. Thus, we define the purpose of our taxonomy as to help researchers and designers recognize useful design option combinations of CAs. In line with this purpose, we define as the meta-characteristic for the taxonomy: High level (i.e. technology-independent) design options (i.e. characteristics that a designer of a CA can decide on) of relevance for enhancing the interaction between humans and CAs (i.e. characteristics should have potential to impact the interaction) in human-agent collaboration. We chose a technology-agnostic scope of the taxonomy in favor of a focus on the interaction-oriented functional characteristics of CAs and for the taxonomy to be of more durable use in a rapidly evolving technological environment. Analogously to e.g. the Six-Layer-Model of Collaboration [5], technology decisions (e.g. for rule-based or machine learning based CAs) should be made once the general conceptual design decisions are made.

Step 2 determine ending conditions: The objective is to recognize when to stop with the taxonomy development. There are so-called objective and subjective ending conditions. While the fulfillment of the objective ending conditions can be verified by us as researchers, the subjective ending conditions need to be examined during an evaluation. *Objective ending conditions* refer to the requirements of a taxonomy and can be tested by answering the following questions [24].

- All objects or a representative sample of objects have been examined?
- At least one object is classified under every characteristic of every dimension?
- In the last iteration, no new dimensions or characteristics were (i) added; (ii) merged or split?
- Every dimension is unique and not repeated?
- Every characteristic is unique within its dimension?
- Each cell (combination of characteristics) is unique and is not repeated?

Subjective ending conditions refer to qualitative attributes that make a taxonomy useful and rigorous. Those are represented by five criteria [24]: Concise, Robust, Comprehensive, Extendible, Explanatory

Iteration 1:

Step 3 approach: The objective is to decide the suitable approach (empirical-to-conceptual vs. conceptual-to-empirical) for the taxonomy development. We start with the empirical-to-conceptual approach suggested by [24]

to consider a broad empirical foundation to develop a taxonomy of this novel, dynamic field.

Step 4e identify (new) subset of objects: The objective is to classify a subset of objects. We aim to achieve a broad overview of the existing literature on the design of CAs. Thus, to select a subset of objects, we start with a literature review. We search four highly relevant databases (IEEEExplore, sciencedirect, ACM DL, EBSCO) to cover literature on collaboration research as well as studies from the human-computer-interaction field. The search string ("conversational agent" OR "chatbot") AND ("collaboration" OR "facilitation" OR "facilitator") was used to derive matches that deal with both CAs and collaboration, consistent with our taxonomy's scope. We obtained 190 hits in the initial search, which we reduced to 71 in a first round of screening titles and abstracts. Exclusion criteria were: (1) pure technological focus as we are looking for technology-neutral design options; (2) focus on collaboration other than on individual or group level, e.g. research collaboration; (3) pure focus on aspects of embodiment, e.g. visual representation of facial expressions with no consideration of collaborative communication. In round two, by excluding one duplicate, nine unavailable full texts and reading the remaining articles for their conceptual fit, we reach a final set of 31 papers. As three papers discuss more than one configuration of CAs, we include 36 objects in our study.

Step 5e identify characteristics and group objects: In the light of the meta-characteristic, the objective is to identify shared characteristics of the selected objects. 21 characteristics have been inductively derived from the analysis of the publications. For this process, we took into consideration the qualitative attributes of a taxonomy [24]. We selected the characteristics based on how well they reflect our meta-characteristic and contribute to the goal of "providing useful explanations of the nature of the objects under study or of future objects to help us understand the objects" [24].

Step 6e group characteristics into dimensions: The objective is to identify dimensions that serve as top categories for the identified characteristics. Based on the twenty-one characteristics we use informally manual grouping and form nine dimensions. Thus, our initial taxonomy T1 can be described as follows: T1 = {number of humans (individual, team, crowd), number of CAs (single, multiple), communication mode (text, speech, multi-modal), representation of CAs (disembodied, embodied), collaboration goal (yes, no), socio-emotional behavior (none/ low, high), sequentiality of CA architecture (yes, no), specifics of the domain (closed, open), role of CA (facilitator, peer, expert)}

Step 7 ending conditions met: The objective is to verify whether objective and subjective ending conditions are met. In our case we firstly verified, whether

the objective ending conditions were met. Since this was achieved, we secondly verified whether the subjective ending conditions were achieved. To check these, we asked six experts in a semi-structured expert evaluation (senior researchers in the field of human-machine collaboration) to evaluate the five conditions on a 7-point Likert scale and to provide suggestions for improving our taxonomy in open ended questions.

Table 1. Expert Evaluation Results

Ending Condition	N	Mean
Concise	6	5.83
Robust	6	5.50
Comprehensive	6	6.33
Extendible	6	6.67
Explanatory	6	6.00
<i>7-point Likert Scale</i>		

All subjective ending conditions are on a mean of respective above 5.50 (Table 1). According to [24] those results constitute an indicator that our taxonomy has high potential of being useful and rigorous. The following qualitative recommendations (Rec) were provided in the expert questionnaires:

Table 2 Revision Recommendations

<i>Revise (R) dimensions and its characteristics (Rec-R)</i>
Rec-R1: Add dimension "character of CA (active/passive)" and merge it with "representation of CA; socio-emotional behavior of CA; role of CA";
Rec-R2: Add dimension "building type of CA (rule-based/ML based)";
Rec-R3: Revise dimension "sequentiality of CA architecture" by i) splitting "yes/no" into "structured conversation process; tracking conversation; rule-based conversation; ad-hoc conversation; non-connected"; ii) explaining its difference to the dimension "goal"
Rec-R4: Revise dimension "role of CA" and more precisely rename characteristic "expert/assistant" to "expert" as assistant can be a facilitator as well.
<i>Clarify (C) issues and meaning (Rec-C)</i>
Rec-C1: Check whether "yes/ no" characteristics are not too simplified;
Rec-C2: Give more guidance to designers for enhancing human-CA-collaboration;
Rec-C3: Improve the delineation of the taxonomy to a literature analysis
Rec-C4: Give guidance on how the taxonomy can be used to derive design patterns.

Iteration 2

Step 3 approach: For the second iteration we used the conceptual-to-empirical approach since we gained significant domain knowledge from the experts and data about the objects. Table 3 illustrates the changes made when revising the taxonomy in iteration 2.

Step 4c Conceptualize new characteristics and dimensions: We ground our new characteristics and dimensions on our evaluation results (1st iteration: step 7).

Step 5c Examine objects for characteristics and dimensions: After revising the taxonomy, we re-examined the classification of all objects from the 1st iteration and

confirmed that all objects still fit. In addition, we added and tested an example application (see section 5).

Step 6c Create/ revise taxonomy: Adding and revising the new dimensions and characteristics creates the new taxonomy T2 displayed in section 4.2.

Table 3 Change History

Rec	Discussion of changes for Taxonomy V2
Rec-R1	Declined. The change would lead to merging dimensions. Role of CA is unique and central design choice and thus needs to stand alone.
Rec-R2	Declined. Focus of the taxonomy would lose its technology-agnostic conceptual design perspective of human-agent collaboration.
Rec-R3	Partially accepted. Recommended characteristics constitute different levels of abstraction, thus cannot be used in the same dimension. Description of dimension was revised and example for classification was added (section 4.1, 5).
Rec-R4	Accepted. Characteristic was changed to “expert”.
Rec-C1	Accepted. Description of dimension was revised (section 4.1) and example for classification was added (section 5).
Rec-C2	Accepted. Section 5 describes implications for using the taxonomy and inter alia discusses design option combinations.
Rec-C3	Accepted. Section 5 describes implications for using the taxonomy and inter alia discusses design option combinations.
Rec-C4	Accepted. Section 6 refers to future research and guides on how to derive design patterns.

Step 7 Ending conditions met: As we just renamed dimensions and characteristics we verified the objective ending conditions and were able to confirm them. Based on the already good to excellent evaluation in the 1st iteration and the refinements, we assess our taxonomy as concise, robust, comprehensive, extendible, explanatory.

Step 7 Ending conditions met: As we just renamed dimensions and characteristics we verified the objective ending conditions and were able to confirm them. Based on the already good to excellent evaluation of the taxonomy in the 1st iteration and the refinements, we assess our taxonomy as concise, robust, comprehensive, extendible, explanatory.

4. Taxonomy

In this section, we first define the dimensions and characteristics, and then present the revised taxonomy.

4.1 Dimensions and Characteristics

Role of CA (RO): In the sample, CAs take over different roles in the team or dya, which we grouped into three categories. First, CAs can act as *facilitators (ROF)* that guide users to reach a certain goal or execute a task. Contexts in which such CAs occur are inter alia tutoring and teaching [11, 36, 19], or facilitation of group interactions [10, 20]. Based on our analysis, most of these

CAs have a more complex model of the context of their field; the collaboration process steps to execute; the current and evolving state of the conversation; the group or user they support or the domain knowledge they contribute. Facilitation CAs have been found to mostly show proactive or directive behavior and often have an underlying process script or structural architecture for the task, e.g. [1, 31, 32]. The second category refers to *CA peers (ROP)*, who aim to merge into a human group [10, 26, 30] or become a sparring partner for an individual [15, 23, 29]. For these CAs, socio-emotional behavior seems to be an important dimension, e.g. [3, 10, 41]. That might indicate to design the CA to match well with the levels of proactivity and expertise of the human team members, e.g. the CA avoids dominating conversations in terms of utterances, which might push human participants into inactivity; the CA is as present as other peers in order to become an integral part of the team. With respect to expertise, a peer CA seems favorable, that is knowing enough in its fields of expertise to be approached by the human users, but equal enough to evoke their own thinking, contributions and critically challenging the peer CAs contributions. The third category covers *CA experts (ROE)* that have certain skills or fields of expertise that differ from those of their human teammates, but mostly act rather reactively upon request, e.g. [14, 22, 44]. As this constellation provides less interruption of the natural human conversation, the possibilities CA influence are limited. When the CAs role is reactive in a way that each of the CA’s activities or utterances is triggered by an utterance or action of human users, the required expertise of the CA in the form of an advanced model of the domain is usually relatively low. This expert role of CAs occurs e.g. in settings, where users engage in chitchat or question-answer dialogues, e.g. [12]. Often, these settings are characterized by spontaneous or evolving conversation rather than an overarching task to be addressed stepwise. Expert CAs react to e.g. user’s query, specific key words or a defined action with a single query like in a FAQ database. To do so, CAs don’t necessarily need an integrated understanding of the big picture and no ability to orchestrate the meta level collaboration process.

Domain (DO): We could identify no distinguishing set of characteristics of CA’s application domains. A reason might be the limited number of objects that spread across a large variety of domains and show differences concerning the task they support and the design they have. However, we found indication that the scope of the domain is related to design and technology choices. We characterize objects as *open (DOO)* or *closed domain (DOC)*. Closed domain CAs are designed for and restricted to a certain type of task or topic and access a bounded knowledge base, e.g. [3, 10, 17]. Open domain CAs often have chitchat purposes and lack

boundaries of a specified task or context information of a closed knowledge field, e.g. [8, 16, 27].

Number of Humans (NH): Design choices may be strongly impacted depending on whether one or more humans are involved in the interaction. Among the papers in our analysis, the vast majority cover scenarios where a CA interacts with one *individual human partner (NHI)*, e.g. [1, 15, 32]. This is a promising application domain, as CAs can be used to turn scenarios of individual work into collaborative work with all the potentials that come with it. The second characteristic we found refers to *two or more humans in a team (NHT)* that interact with the CA(s) simultaneously, e.g. [11, 20, 36]. An explanation is that such settings are much more complex to design, as the CA also needs to be capable of distinguishing utterances from different users and should ideally be able to grasp and steer social dynamics within the human group. Furthermore, while in a 1:1 setting the user only has the choice to interact with the CA or not/maybe abuse it (and thus might fail on the task), groups of humans can ignore the CA, i.e. the CA might be in competition with other humans for attention. The otherness of CAs can also evoke unplanned social dynamics (e.g. the team might unite against the CA or make fun of it; team members might tend to speak about the CA as if it were not present or abuse it). Individuals, who would have collaborated with the CA in a 1:1 setting, might behave differently, if being observed by other humans. While crowd collaboration is a field, gaining growing attention in research and practice, only one CA in our sample is used in a *crowd setting (NHC)* with an undefined, large number of human collaborators on a platform [4].

Number of CAs (NC): The vast majority of studies apply one *single CA (NCI)*, e.g. [12, 29, 41]. In some system designs, several CAs with different functionalities are indeed combined in a modular architecture, but usually appear as one CA toward the side of the human user(s). Only few objects in our sample mention the usage of *multiple CAs (NCM)* in one collaborative setting as an area of application for their systems, e.g. [10, 15, 45]. This prevalence of single CAs could either be due to a lack of practical necessity for multiple CAs or the nature of the scholarly studies so far. In line with the former, the technological feasibility to combine a multitude of social/emotional and task related skills into one CA that might even change in different situations and phases of the collaboration may make it unnecessary to confront the human user with different CAs. The latter might be a bias due to the mainly experimental or system design based nature of the objects under study, which aim to prove the functioning of agent systems in general and reduce complexity by starting with one CA in a lab situation. Among the studies within our analysis,

only [15] compare the impact of different team compositions, i.e. one CA vs. two CAs collaborating with one human participant in an experimental study. They find that participants in the single CA condition generated more ideas in a creative task than did participants in the double CA condition. Although this singular study certainly does not allow for general conclusions, it points to the necessity of close consideration of different team configurations in human-agent-collaboration design.

Communication Mode (CM): In the scope of our work (natural language based CAs), communication between human and CA can take place in *text (CMT)* or *speech (CMS)* form or in a combination of *multimodal (CMM)* communication channels. Due to the literature search focus on conversation based collaboration, we excluded papers that mainly focus on characteristics such as gaze, gestures or the visual design of an avatar, resulting in a prevalence of text or speech based CAs.

CA Representation (RE): In line with the communication mode and our search focus, many publications in the sample address *disembodied (RED)* CAs with no visual or physical representation, e.g. [2, 20, 23]. This characteristic restricts the possibilities of CA design to express a CAs' "personality" mainly to their conversational utterances. However, other objects use *embodied (REE)* CAs with a simple or more complex visual representation, e.g. [4, 14, 36]. Embodiments range from two dimensional avatars that are either static or employed with a simple animation, which are unrelated to the actual conversation, to more advanced types of embodiment, e.g. reactive avatars with adaptive gesture and facial expressions or physical representations in the form of a robot. The latter CAs may evoke complex reactions and perceptions among collaboration partners (see e.g. [29]), but also provide a broader range of design options for humanness.

Collaboration Goal Direction (CG): Among the analyzed CAs, we found such that are collaborating with users toward a common goal or guide them toward completing a certain task (*goal oriented (CGY)*), e.g. [11, 15, 20] and such that are *non goal-oriented (CGN)*, e.g. [8, 41, 45], whose value lies in the interaction itself.

Sequentiality of Process Structure (PS): Goal oriented CAs show functionalities that are directed toward reaching a certain goal or completing a task by rather enforcing a pre-defined task execution script (*sequential process structure (PSY)*), e.g. [3, 17, 23] or by enforcing framing conditions/restrictions for the task in a more adhoc or flexible execution (*no sequential process structure (PSN)*), e.g. [12, 16, 38].

Socio-emotional Behavior (SE): This dimension refers to behavior and utterances of CAs that aim to detect or express certain emotions, show affection, empathy toward users or evoke certain emotional reaction by

users. In the taxonomy, we distinguish between *none/low levels of socio-emotional behavior implemented in CAs (SEN)* and *explicit/high levels of consideration of these functionalities (SEY)*. Some functionalities are designed to make the CA appear more human-like with a unique personality and thus evoke affection for the CA. By giving the CA a perceptible personality, e.g. letting it express emotions like happiness or sadness or empathy [10], designers aim to build rapport and make users perceive CAs as similar to themselves. The underlying assumption that users accept and interact better with CAs that appear similar to humans needs to be tested in further studies. If appropriate emotional behavior can be implemented in CA design, designers might use psychological or sociological research to guide their design, to achieve better acceptance, trust, satisfaction etc. A recent experiment [45] suggests that deep learning technology under certain conditions can detect and react to emotional user requests better than baseline information retrieval systems and comparably well as human customer service. Results point out promising potential to make CAs capable of emotional behavior.

4.2 Display of Taxonomy

The 36 objects classified in the taxonomy (Table 4) are grouped by “roles of the CAs” and sorted in a way to extract design option combinations (in section 5).

4.3 Exemplary Classification

This example illustrates, how objects have been and can be classified. [3] present an „*animated, conversational computer agent designed to promote antipsychotic medication adherence among patients with schizophrenia*”. A **single CA** collaborates with a **single human user** toward the joint goal to promote medication adherence of the user and keep track of his/her health related behavior. The CA is **embodied** as an animated virtual character that is, in addition to chat communication, able to show **facial expressions and gestures**. In such, the CA architecture models a **closed domain** and knowledge record on the user. It contains a **process structure** of topics to cover in the conversations with the user and the CA uses **social chat** to build rapport. It is designed to be perceived as an **empathic peer** for the user to motivate him/her to share information honestly.

5. Implications for the Usage of the Taxonomy and Future Research

In this section, we provide guidance on how to read and use the taxonomy. It is important that the taxonomy is useful. According to [24], a useful taxonomy meets the ending conditions and can be used by others (designers and researchers of CAs). Thus, describing meaningful design option combinations will serve as a

foundation for being used. We identify “Role of CA” as key dimension for grouping objects as the context determines the role of the CA. Thus, it is one of the design concerns to be respected at the beginning and design option combinations differ more between roles than within a role. We describe combinations along three roles: facilitator, peer, expert. To ensure easy overview, we use a common structure in the subsection for each role, with summary of the objects, identified design option combinations, discussion of their implications, and white spots for potential future research. This structure is also a first step to design patterns of human-CA collaboration [25].

5.1 Facilitators

Descriptive and formal facts: In 13 out of 36 cases, CAs take over a facilitating role. Most CAs in this group occur in education or related domains, e.g. knowledge handling. Most interact with groups (9), while only 3 facilitate individual work and 1 is in a crowd setting. They are more often text-based (8) than speech-based (5) and more often disembodied (8) than embodied (5).

Design option combinations: We found the following interesting sets of recurring combinations:

- ROF-education: The majority of facilitators are used in education of related domains, e.g. knowledge acquisition or behavior change.
- ROF-DOC-CGY: In the sample, all facilitator CAs are focused on a closed domain and have a collaboration goal, but only seven follow a sequential structure.
- ROF-NH1-RED: All facilitator CAs with individual human counterparts in the sample are disembodied.
- ROF-CMT-RED: Text-based facilitator CAs are mostly disembodied.
- ROF-CMS-REE: Speech-based facilitators in the sample tend to have an embodiment.

Discussion of Design Option Combinations: The prevalence of facilitators in education might point to this domain being especially open to novel technology, e.g. intelligent tutoring systems, or that it has a high need for scalable learning process support. Other domains with similar characteristics, such as knowledge management or behavior change, might benefit from looking into this pioneer work. The finding that all facilitators in the sample have a collaboration goal is in line with the goal-oriented role of facilitation in Collaboration Engineering. However, not all of them follow a sequential process architecture, which would be a characteristic of traditional collaborative work practices. This points to an interesting potential, as CAs, due to their ability to flexibly react to the conversation flow, might allow for new, less rigid collaboration facilitation than traditional tools. The fact that most facilitators in the sample are text based and disembodied could either be explained by a low need for rich, human-like representation of facilitator CAs, as they don’t need to be

Table 4 Revised Taxonomy (T2)

Paper	Application Domain	CA Role	Role (RO)			Domain (DO)		No. of Humans (NH)			No. of CAs (NC)		Communication Mode (CM)			CA Representation (RE)		Collab. Goal Direction (CG)		Sequ. of Proc. Struct. (PS)		Socio-emot. Behav. (SE)	
			ROF - Facilitator	ROP - Peer	ROE - Expert	DOC - closed	DOO - open	NH1 - Individual	NHT - Team	NHC - Crowd	NC1 - Single	NCM - Multiple	CMT - Text	CMS - Speech	CMM - Multimodal	RED - disembodied	REE - embodied	CGY - Yes	CGN - No	PSY - Yes	PSN - No	SEY - Yes	SEN - No
Allen et al. (2002)[1]	probl. solving, behav. change	Medical Advisor	x			x		x		x		x		x		x		x				x	
Roda et al. (2003)[32]	knowl. sharing, behav. change	advisor	x			x		x		x		x		x		x		x			x		
Rich et al. (2002) [31]	education	Tutor	x			x		x		x		x		x		x		x				x	
Dyke et al. (2013) [11]	education	Tutor	x			x			x		x		x		x		x			x		x	
Kumar & Rosé (2010) [19]	education	Tutor	x			x		x		x		x		x		x		x			x		
Kumar & Rosé (2014) [20]	Decision making	Adminstr.	x			x		x		x		x		x		x		x			x		
Tegos et al. (2012b)[38]	education	Tutor	x			x		x		x		x		x		x		x			x		
Dohsaka et al. (2009) [10]	Gaming	Quizmaster	x			x			x		x		x		x		x			x		x	
Tegos et al. (2012a)[37]	education	Tutor	x			x		x		x		x		x		x		x			x		
Tegos et al. (2014a)[35]	education	Tutor	x			x		x		x		x		x		x		x			x		
Tegos et al. (2014b)[39]	education	Tutor	x			x		x		x		x		x		x		x			x		
Tegos et al. (2015)[36]	education	Tutor	x			x		x		x		x		x		x		x			x		
Bradesko et al. (2017) [4]	knowledge acquisition	Crowd Interrogator	x			x			x		x		x		x		x			x		x	
Mell et al. (2015) [23]	negotiation	negotiation partner		x		x			x		x		x		x		x			x		x	
Hayashi & Ono (2013) [15]	ideation	Peer feedback		x		x			x		x		x		x		x			x		x	
Hubal et al. (2008) [17]	education	skill assessment		x		x			x		x		x		x		x			x		x	
Hayashi & Ono (2013) [15]	ideation	Peer feedback		x		x			x		x		x		x		x			x		x	
v.d. Pütten et al. (2010) [29]	Inform. retrieval	active listener		x		x			x		x		x		x		x			x		x	
Bickmore et al. (2010) [3]	Health	Health Advisor		x		x			x				x		x		x			x		x	
Turunen et al. (2011) [41]	health and fitness	fitness companion		x		x			x				x		x		x			x		x	
Hayashi & Ono (2013) [15]	ideation	Peer feedback		x		x			x		x		x		x		x			x		x	
Xu & Lombard (2017) [45]	Decision making	Peer		x		x			x		x		x		x		x			x		x	
Hayashi & Ono (2013) [15]	ideation	Peer feedback		x		x			x		x		x		x		x			x		x	
Prendinger & Ishizuka (2002) [28]	Gaming	Black Jack advisor		x		x			x				x		x		x			x		x	
Dohsaka et al. (2009) [10]	Gaming	Peer		x		x			x		x		x		x		x			x		x	
Rehm (2008) [30]	Gaming	Opponent		x		x			x				x		x		x			x		x	
Portela & Granell-Canut (2017) [27]	Chitchat	General Bot		x		x			x		x		x		x		x			x		x	
Angeli & Brahnham (2008) [2]	Chitchat	General Bot		x		x			x		x		x		x		x			x		x	
Corti & Gillespie (2016) [8]	Chitchat	General Bot		x		x			x		x		x		x		x			x		x	
Hill et al. (2015) [16]	Chitchat	General Bot		x		x			x		x		x		x		x			x		x	
Corti & Gillespie (2016) [8]	Chitchat	General Bot		x		x			x				x		x		x			x		x	
Porcheron et al. (2017) [26]	Chitchat	General Bot		x		x			x				x		x		x			x		x	
Eisman et al. (2012) [12]	Q&A	Website Navigation			x	x			x				x		x		x			x		x	
Hasler et al. (2013) [14]	Survey	Interrogator			x	x			x				x		x		x			x		x	
Louvet et al. (2017) [22]	Inform. retrieval	assisted search			x	x			x				x		x		x			x		x	
Xu et al. (2017) [44]	customer service	customer agent			x	x			x				x		x		x			x		x	
Sum:			13	19	4	29	7	23	12	1	30	6	21	10	5	18	18	27	9	12	24	15	21

perceived by users as one of their kind. Alternatively, this finding might be driven by the fact that the analyzed CAs were research project, where the effort was put in the complex architecture of a CA rather than its embodiment and advanced communication mode. Both hypotheses deserve further investigation.

White spots and future research: In search for white spots in the taxonomy, we identified two exemption cases: [10] present the only CA facilitator with a second CA in the role of a peer. This is an interesting starting point for a stream of research that should explore the interaction dynamics and design potentials of CAs with different roles and among each other. The second unusual facilitator CA [4] is the only CA in the sample that is used in a crowd setting. With the rising importance of crowd collaboration and the scalability advantage of automated facilitation compared to human facilitation, we identify this as an exciting field for further exploration.

5.2 Peers

Descriptive and formal facts: Peers make up the largest group in the sample. In contrast to facilitators, there are also six open domain peer CAs. While still in the minority (5 out of 19), settings with multiple CAs are more prevalent than in the other roles. Peers are more often embodied (11) than disembodied (8), but no dominant form of representation has emerged. Peers use a mixed set of communication modes, with 9 being text-based, 5 speech based and 5 multimodal. Most peers (16) interact with an individual counterpart, while only 3 are designed for team settings. Peers are used in a variety of domains, but all gaming CAs and chitchat CAs in the sample are peers.

Design option combinations: We found the following interesting sets of recurring combinations:

- ROP-NH1: Peer CAs are predominantly used by one human partner.
- ROP-REE: Peer CAs are often embodied.
- ROP-NCM: Peers occur more often than any other role in multi-agent settings.
- ROP-gaming: All CAs in gaming in the sample are peers.
- ROP-chitchat: All CAs for chitchat in the sample are peers.
- ROP-DOO-NH1-NC1-CGN-PSN-chitchat: Three General Bots for chitchat have the exact same classification in all dimensions, three others only differ in communication mode, socio-emotional behavior and/or embodiment

Discussion of Design Option Combinations: The recurring ROP-NH1 combination strengthens the potential of CAs to turn an individual task into a collaborative one by inducing a non-human teammate. While

a peer in a team setting might need to compete for attention and acceptance with other team members, it is the only contact in 1:1 situations. The prevalence of ROP-REE might indicate the importance of humanness of peers compared to other roles. While a facilitator might have a certain authority due to its role, a peer might rely on being accepted by the human peers as one of their kind for its full effectiveness. This assumption deserves further investigation. The ROP-NCM combination provides a starting point for exploring different configurations of multi agent settings, either with several peer CAs to extend a team or a combination of CAs with different roles. The prevalence of peers in gaming and chitchat points out the variety of potential applications in these domains. Other domains like education or problem solving might want to look into adopting these approaches. The dominant ROP-DOO-NH1-NC1-CGN-PSN-chitchat in general bots may indicate that this type is either the most established one or it still lacks diversification.

White spots and future research: Future research should explore the effectiveness of different levels of socio-emotional behavior and “humanness” of peer CAs. This stream might also want to consider, whether peer CAs should disclose themselves as CAs or not.

5.3 Experts

Descriptive and formal facts: Experts are the smallest group and are all single CAs designed for interacting with individual humans via text. Three of four are designed for a closed domain, have a defined goal, non-sequential architecture and no socio-emotional behavior.

Design option combinations: We found the following interesting sets of recurring combinations:

- ROE-NC1-NH1-CMT: All expert CAs in the sample interact one on one with a human via text.
- ROE-NC1-NH1-CMT-DOC-CGY-PSN-SEN: Among these experts, a closed domain toward a known goal, but no sequential process structure and socio-emotional behavior are common

Discussion of Design Option Combinations: Combinations must be interpreted with caution due to the small number of objects. Experts seem to have less complex architectures (pairwise interaction, no sequentiality, no socio-emotionality) compared to others. As experts react to human requests and are addressed by them for expertise, they might not need the same advanced understanding of the collaboration process state or social relationships as e.g. facilitators.

White spots and future research: Due to the small number of experts in the sample, useful design option combinations and potential applications of this type of CAs remain to be explored further. For complex collaborative work practices, expert CAs might

enhance the abilities in a team that is facilitated by a human or another CA facilitator by complementing the skills of the human team members. In such, it should be explored how to extend the knowledge from one on one expert interaction to team settings.

6. Conclusion and Outlook

According to [13] and their understanding of theory in information systems our taxonomy resembles a theory for analyzing of type I that tells “what is” and thus makes a design science research contribution for both researchers and designers of CAs alike. Such types of theories are needed when just little is known about the phenomenon of interest as applies to the dynamic evolution of CAs. To validate the usefulness of the taxonomy a broader evaluation is needed, in which designers use the taxonomy to initially guide their design choices for CAs, further challenge the validity of the design option combinations or explore the suggested white spots. We also encourage researchers on CAs in collaboration to explore further ways of organizing knowledge on the domain, as different types of taxonomies might be needed for different meta characteristics than the one within scope of this paper, e.g. for a linkage of the conceptual/functional design of a CA and its technological implementation. In order to advance the discussed insights toward a “theory of explanation” [13], future research should aim to understand the science behind the observed design option combinations. This might lay the foundation to derive design patterns for CAs in collaboration.

7. References

[1] Allen, J., N. Blaylock, and G. Ferguson, “A problem solving model for collaborative agents”, First International joint conference on Autonomous agents and multiagent systems, 2002.

[2] Angeli, A.D. and S. Brahmam, “I hate you! Disinhibition with virtual partners”, *Interacting with Computers*, 20(3), 2008, pp. 302–310.

[3] Bickmore, T.W., K. Puskar, E.A. Schlenk, L.M. Pfeifer, and S.M. Sereika, “Maintaining reality: Relational agents for antipsychotic medication adherence”, *Interacting with Computers*, 22(4), 2010, pp. 276–288.

[4] Bradeško, L., M. Witbrock, J. Starc, Z. Herga, M. Grobelnik, and D. Mladenčić, “Curious Cat—Mobile, Context-Aware Conversational Crowdsourcing Knowledge Acquisition”, *ACM Transactions on Information Systems*, 35(4), 2017, pp. 1–46.

[5] Briggs, R.O., G.L. Kolfshoten, G.-J. de Vreede, C.C. Albrecht, S. Lukosch, and D.L. Dean, “A Six-Layer Model of Collaboration”, in *Collaboration Systems: Concept, Value, and Use*, J.F. Nunamaker Jr, R.O. Briggs, and N.C.R. Romano Jr, 2014, New York, USA.

[6] Briggs, R.O., G.L. Kolfshoten, G.-J. de Vreede, S. Lukosch, and C.C. Albrecht, “Facilitator-in-a-Box: Process Support Applications to Help Practitioners Realize the Potential of Collaboration Technology”, *Journal of Management Information Systems*, 29(4), 2013, pp. 159–194.

[7] Chaves, A.P. and M.A. Gerosa, “Single or Multiple Conversational Agents?”, *CHI Conference on Human Factors in Computing Systems*, 2018.

[8] Corti, K. and A. Gillespie, “Co-constructing intersubjectivity with artificial conversational agents: People are more likely to initiate repairs of misunderstandings with agents represented as human”, *Computers in Human Behavior*, 58, 2016, pp. 431–442.

[9] DALE, R., “The return of the chatbots”, *Natural Language Engineering*, 22(05), 2016, pp. 811–817.

[10] Dohsaka, K., R. Asai, R. Higashinaka, Y. Minami, and E. Maeda, “Effects of conversational agents on human communication in thought-evoking multi-party dialogues”, *SIGDIAL 2009 Conference*, 2009.

[11] Dyke, G., D. Adamson, I. Howley, and C.P. Rose, “Enhancing Scientific Reasoning and Discussion with Conversational Agents”, *IEEE Transactions on Learning Technologies*, 6(3), 2013, pp. 240–247.

[12] Eisman, E.M., V. López, and J.L. Castro, “A framework for designing closed domain virtual assistants”, *Expert Systems with Applications*, 39(3), 2012, pp. 3135–3144.

[13] Gregor, S., “The nature of theory in information systems”, *MIS Q*, 30(3), 2006, pp. 611–642.

[14] Hasler, B.S., P. Tuchman, and D. Friedman, “Virtual research assistants: Replacing human interviewers by automated avatars in virtual worlds”, *Computers in Human Behavior*, 29(4), 2013, pp. 1608–1616.

[15] Hayashi, Y. and K. Ono, “Embodied conversational agents as peer collaborators: Effects of multiplicity and modality”, *2013 IEEE RO-MAN*, 2013.

[16] Hill, J., W. Randolph Ford, and I.G. Farreras, “Real conversations with artificial intelligence: A comparison between human–human online conversations and human–chatbot conversations”, *Computers in Human Behavior*, 49, 2015, pp. 245–250.

[17] Hubal, R.C., D.H. Fishbein, M.S. Sheppard, M.J. Paschall, D.L. Eldreth, and C.T. Hyde, “How Do Varied Populations Interact with Embodied Conversational Agents? Findings from Inner-city Adolescents and Prisoners”, *Computers in Human Behavior*, 24(3), 2008, pp. 1104–1138.

[18] Kolfshoten, G.L., R.O. Briggs, G.-J. de Vreede, P.H.M. Jacobs, and J.H. Appelman, “A conceptual foundation of the thinkLet concept for Collaboration Engineering”, *International Journal of Human-Computer Studies*, 64(7), 2006, pp. 611–621.

[19] Kumar, R. and C.P. Rosé, “Engaging learning groups using social interaction strategies”, *Human Language Tech-*

- nologies: The 2010 Annual Conference of the North American Chapter of the Association for Computational Linguistics, 2010.
- [20] Kumar, R. and C.P. Rosé, “Triggering effective social support for online groups”, *ACM Transactions on Interactive Intelligent Systems*, 3(4), 2014, pp. 1–32.
- [21] Lieberman, H., “Autonomous interface agents”, *ACM SIGCHI Conference on Human factors in computing systems*, 1997.
- [22] Louvet, J.-B., G.D. Duplessis, N. Chaignaud, L. Vercoeur, and J.-P. Kotowicz, “Modeling a collaborative task with social commitments”, *Procedia Computer Science*, 112, 2017, pp. 377–386.
- [23] Mell, J., G. Lucas, and J. Gratch, “An Effective Conversation Tactic for Creating Value over Repeated Negotiations”, *International 2015 Conference on Autonomous Agents and Multiagent Systems*, 2015.
- [24] Nickerson, R.C., U. Varshney, and J. Muntermann, “A method for taxonomy development and its application in information systems”, *European Journal of Information Systems*, 22(3), 2013, pp. 336–359.
- [25] Petter, S., D. Khazanchi, and J.D. Murphy, “A Design Science Based Evaluation Framework for Patterns”, *The DATA BASE for Advances in Information Systems*, 41(3), 2010, pp. 9–26.
- [26] Porcheron, M., J.E. Fischer, and S. Sharples, “Do Animals Have Accents?” Talking with Agents in Multi-Party Conversation”, *2015 International Conference on Autonomous Agents and Multiagent Systems*, 2015.
- [27] Portela, M. and C. Granell-Canut, “A new friend in our smartphone? Observing Interactions with Chatbots in the Search of Emotional Engagement”, *XVIII International Conference on Human Computer Interaction*, 2017.
- [28] Prendinger, H. and M. Ishizuka, “SCREAM: Scripting Emotion-based Agent Minds”, *AAMAS’02*, 2002.
- [29] Pütten, A.M. von der, N.C. Krämer, J. Gratch, and S.-H. Kang, ““It doesn’t matter what you are!” Explaining social effects of agents and avatars”, *Computers in Human Behavior*, 26(6), 2010, pp. 1641–1650.
- [30] Rehm, M., ““She is just stupid”—Analyzing user-agent interactions in emotional game situations”, *Interacting with Computers*, 20(3), 2008, pp. 311–325.
- [31] Rich, C., N. Lesh, A. Garland, and J. Rickel, “A plug-in architecture for generating collaborative agent responses”, *AAMAS ’02*, 2002.
- [32] Roda, C., A. Angehrn, T. Nabeth, and L. Razmerita, “Using conversational agents to support the adoption of knowledge sharing practices”, *Interacting with Computers*, 15(1), 2003, pp. 57–89.
- [33] Seeber, I., E. Bittner, R.O. Briggs, G.-J. de Vreede, T. de Vreede, D. Druckenmiller, R. Maier, A.B. Merz, S. Oeste-Reiß, N. Randrup, G. Schwabe, and M. Söllner, “Machines as Teammates: A Collaboration Research Agenda”, *Hawaii International Conference on System Sciences (HICSS)*, 2018.
- [34] Shawar, B.A. and E. Atwell, “Chatbots: Are they Really Useful?”, *LDV Forum*, 2007, pp. 29–49.
- [35] Tegos, S., S. Demetriadis, and A. Karakostas, “Leve- raging Conversational Agents and Concept Maps to Scaffold Students’ Productive Talk”, *2014 International Conference on Intelligent Networking and Collaborative Systems*, 2014.
- [36] Tegos, S., S. Demetriadis, and A. Karakostas, “Promoting academically productive talk with conversational agent interventions in collaborative learning settings”, *Computers & Education*, 87, 2015, pp. 309–325.
- [37] Tegos, S., S. Demetriadis, and T. Tsiatsos, “Using a Conversational Agent for Promoting Collaborative Language Learning”, in *Fourth International Conference on Intelligent Networking and Collaborative Systems*, 2012.
- [38] Tegos, S., S.N. Demetriadis, and A. Karakostas, “Exploring the Impact of a Conversational Agent When Triggering Students’ Discourse in Online Collaboration”, *12th International Conference on Advanced Learning Technologies*, 2012.
- [39] Tegos, S., S.N. Demetriadis, and A. Karakostas, “Conversational Agent to Promote Students’ Productive Talk: The Effect of Solicited vs. Unsolicited Agent Intervention”, *14th International Conference on Advanced Learning Technologies*, 2014.
- [40] Traumer, F., S. Oeste-Reiß, and J.M. Leimeister, “Towards a Future Reallocation of Work between Humans and Machines – Taxonomy of Tasks and Interaction Types in the Context of Machine Learning”, *International Conference on Information Systems*, 2017.
- [41] Turunen, M., J. Hakulinen, O. Ståhl, B. Gambäck, P. Hansen, M.C. Rodríguez Gancedo, R.S. de La Cámara, C. Smith, D. Charlton, and M. Cavazza, “Multimodal and mobile conversational Health and Fitness Companions”, *Computer Speech & Language*, 25(2), 2011, pp. 192–209.
- [42] Uthus, D.C. and D.W. Aha, “Multiparticipant chat analysis: A survey”, *Artificial Intelligence*, 2013.
- [43] Vreede, G.-J. de and R.O. Briggs, “Collaboration Engineering: Designing Repeatable Processes for High-Value Collaborative Tasks”, *38th Hawaii International Conference on System Sciences*, 2005.
- [44] Xu, A., Z. Liu, Y. Guo, V. Sinha, and R. Akkiraju, “A New Chatbot for Customer Service on Social Media”, *CHI Conference on Human Factors in Computing Systems - CHI ’17*, 2017.
- [45] Xu, K. and M. Lombard, “Persuasive computing: Feeling peer pressure from multiple computer agents”, *Computers in Human Behavior*, 74, 2017, pp. 152–162.
- [46] Zadrozny, W., M. Budzikowska, J. Chai, N. Kambhatla, S. Levesque, and N. Nicolov, “Natural language dialogue for personalized interaction”, *Communications of the ACM*, 43(8), 2000, pp. 116–120.