

# Is my agent good enough? Evaluating Embodied Conversational Agents with Long and Short-term interactions

Juliane B. S. dos Santos  
VHLAB

School of Tech., PUCRS  
Porto Alegre, Brazil

juliane.beatrycce@edu.pucrs.br

Paulo Ricardo Knob  
VHLAB

School of Tech., PUCRS  
Porto Alegre, Brazil

paulo.knob@edu.pucrs.br

Victor Putrich Scherer  
VHLAB

School of Tech., PUCRS  
Porto Alegre, Brazil

Victor.Putrich@edu.pucrs.br

Soraia Raupp Musse  
VHLAB

School of Tech., PUCRS  
Porto Alegre, Brazil

soraia.musse@pucrs.br

**Abstract**—The use of digital resources has been increasing in every instance of today’s society, being it in business or even ludic purposes. Despite such ever increasing use of technologies as interfaces, in all fields, it seems that it lacks the importance of users perception in this context. This work aims to present a case study about the evaluation of ECAs. We propose a Long-Term Interaction (LTI) to evaluate our conversational agent effectiveness through the user perception and compare it with Short-Term Interactions (STIs), performed by three users. Results show that many different aspects of users perception about the chosen ECA (i.e. Arthur) could be evaluated in our case study, in particular that LTI and STI are both important in order to have a better understanding of ECA impact in UX.

**Index Terms**—Embodied Conversational Agent, Virtual Agent, Long-term Interaction, User Experience.

## I. INTRODUCTION

Embodied Conversational Agents (ECAs) can help people with some of their tasks while transmitting the sensation of talking with a real human. In addition, many games nowadays use intelligent virtual agents to improve the immersion and reality of the gaming experience, being it as Non-Playable Characters (NPCs) or tutorial/guides. It is important that people who play these games feel comfortable and enjoy such interactions; thus, it is equally important to measure and evaluate ECAs based on the perception of people.

When talking about UX, it is important to consider the user interaction with the proposed solution, software, or object as part of the development process or, even, while testing the final product. This interaction can happen in different ways, like using Short Interaction Sequences (SIS) [1] or by exposing users to the final product and asking for feedback at the end of the interaction [1], [2].

In this work, we present a case study which evaluates an ECA based on users perception. First, we perform a Long-Term Interaction (LTI) and compare the perceptual results with Short-Term Interactions (STIs) performed by different users. The goal is to compare them and consolidate our recommendations for ECAs evaluation. The ECA used in our LTI, which occurred for 41 days, is Arthur, proposed by Knob et al. [3]. In order to register the LTI, an user diary was built

to log the daily interactions with Arthur, which we called Interaction Logbook. In the end, with the results achieved, we provide improvement suggestions for the chosen ECA that can be useful for any development team of Embodied Conversational Agents.

## II. RELATED WORK

Ruttikay et al. [4] listed four of possible ECA behaviors: Embodiment, Input Recognition System, Model of personality and Emotions. In their model, the authors divide the ECA into three main modules: Embodiment, which is responsible for how the ECA will look like; Mental Capacities, which defines the personality and behavior of such ECA; and Application Interface, responsible for the information processing. In the literature, it is possible to find many models proposed to build an ECA [3], [5], [6]. The work of Yalcin [7] aims to model empathetic behavior on Embodied Conversational Agents (ECAs). Their ECA has three stages: listening, where the agent captures input from the person it is talking to; thinking, where the agent process the information; and speaking, where the agent gives a proper response, both with words and gestural behavior.

According to Morville [8], seven factors influence User Experience (UX). Each of them is important to determine the success or failure of a project, being it a digital product or not. They are <sup>1</sup>: Useful, Usable, Findable, Credible, Desirable, Accessible and Valuable. Another concept that is vastly considered in UX is the “10 Usability Heuristics for User Interface Design”, created by Nielsen [9] They are <sup>2</sup>: Visibility of system status; Match between system and the real world; User control and freedom; Consistency and standards; Error prevention; Recognition rather than recall; Flexibility and efficiency of use; Aesthetic and minimalist design; Help users recognize, diagnose, and recover from errors; and Help and documentation.

In addition to the Usability Heuristics [9], the Nielsen’s Severity Ratings [10] are essential on the UX evaluation

<sup>1</sup>Please refer to Morville [8] for further details

<sup>2</sup>For further information, please refer to Nielsen [9] publication

because it helps to understand the priority of each change that may be needed in the system. The ratings are: 0 = I don't agree that this is a usability problem at all; 1 = Cosmetic problem only: need not be fixed unless extra time is available on project; 2 = Minor usability problem: fixing this should be given low priority; 3 = Major usability problem: important to fix, so should be given high priority; and 4 = Usability catastrophe: imperative to fix this before product can be released. Finally, there are some usability metrics used in the work of Ren et al. [11] and Tullis and Albert [12], that are based on the user perception such as Effectiveness, Efficiency and Satisfaction of the interaction with the product<sup>3</sup>.

Concerning ECAs evaluation, Bickmore et al. [13] aimed to discover how to maintain user engagement in interactions with virtual agents. The results achieved demonstrate that user engagement can be manipulated using relatively simple techniques. Babu et al. [14] evaluated Marve, a virtual receptionist placed at the entrance of their research laboratory. Results achieved indicate that the users enjoyed their interaction with Marve, also being able to perceive the agent as a human-like conversational partner.

### III. ARTHUR – THE ECA

In order to conduct our research we chose to work with Arthur, as proposed by Knob et al. [3]. The focus of Arthur is to provide a natural interaction with the final user by performing a consolidation of its Short-Term Memory and Long-Term Memory. Such process makes Arthur able to store and remember important pieces of the previous conversation, creating the feeling that the user is having a natural interaction, as if it was with another human [3].

1) *Smalltalking Module*: As defined by the Cambridge Dictionary, small talks have the basic definition that is a "conversation about things that are not important, often between people who do not know each other well"<sup>4</sup>. The main advantage of bringing this concept to Arthur is that it allows to build a more approximate relation between virtual agent and human, especially when it is for a Long-term Interaction [15]. In order to build our "Smalltalking" Module, we chose to create a simple conversational structure based on a Decision Tree [16]. Thus, our "Smalltalking" Module is divided into three parts: Topics, Dialogues and Dialog Tree. For each of the Dialogues defined inside a certain Topic, a Dialog Tree is built. This Dialog Tree is composed of branches and nodes which define what Arthur can speak to the user. The path through the tree is made by Arthur saying/asking something to the user and waiting for his/her answer. Such answer is used to find the next utterance at the tree, by analyzing the polarity of the sentence (as done by Arthur [3]) and comparing it with the polarity of the next nodes. Finally, since the "Smalltalking" Module is triggered only when the interaction seems to "cool down", a timer was defined. We empirically defined that if the user says nothing to Arthur for 30 seconds, but keeps there,

recognized by the Computer Vision module, Arthur randomly initiates a Smalltalk conversation.

### IV. THE CASE STUDY

In this section, we present the case study followed in this work. Firstly, we conduct the Interaction Evaluation Steps, which are divided into three stages: 1) To conduct a Long-term Interaction with Arthur; 2) To conduct Short-term Interactions with a few users; and 3) To compare the findings and elaborate suggestions. The idea behind conducting the Short-Term Interactions was, indeed, to compare two different ways of interaction and see how useful each of them can be in our study. Next, we describe the proposal for evaluation of users perception, which is divided into two parts: 1) Heuristic Evaluation, following Nielsen's 10 Usability Heuristics [9]; and 2) Severity Evaluation, following the Severity Ratings for Usability Problems [10].

#### A. The Long-term Interaction

Inspired by works like Kanda et al. [17], we evaluated Arthur using a Long-Term Interaction (LTI) with one user. This person was contacted by our team and agreed to participate in the case study presented here. It was a woman, age of 24, with high experience with technology, but no experience with embodied conversational agents. The main goal was to achieve measurable results on the perception of humans when interacting with Arthur. Thus, the LTI user's goal was to talk with Arthur everyday, as if Arthur was a friend of the user. The LTI included the following tasks: *i*) Perform daily interactions with the ECA model developed by Knob et al. [3], as presented in Section III; *ii*) Register the interaction in the Logbook described in Section IV-A1; *iii*) Evaluate the users perception based on the Effectiveness, Efficiency and Satisfaction metrics [11], [12]; and *iv*) Evaluate the ECA following Nielsen's Heuristics and Severity Ratings [9], [10]. Results are presented in Section V.

1) *Interaction Steps and Logbook*: Our Long-term Interaction had a total duration of 41 days. The approach chosen was to, each day, talk to Arthur as we talk to a person. The idea was to see if Arthur would be able to keep interacting in a natural way, as well as if he would be able to learn things with/about the user.

To provide a detailed vision of how the Long-term Interaction with Arthur occurred, a user diary was created, which we called Interaction Logbook. It contained a few questions which were created based on the features of Arthur, but following the Effectiveness, Efficiency and Satisfaction criteria [11]. In addition, a free text field was added to the questionnaire, so the user could write some details about her routine, feelings and so on. For each question, the researcher could select one of five options of a Likert scale: Very Unsatisfied, Unsatisfied, Regular, Satisfied, and Very Satisfied. After the LTI finished, the same questionnaire was applied to the users selected to perform the STIs. This way, we were able to compare the results of both LTI and STIs in terms of extracted

<sup>3</sup>Please refer to Ren et al. [11] and Tullis and Albert [12] for further details

<sup>4</sup><https://dictionary.cambridge.org/dictionary/english/small-talk?q=small+talk>

perceptual data from the users. These results are going to be presented in Section V.

### B. The Short-term Interactions

We tested our ECA with three participants: A woman with medium familiarity with computers, 52 years old; a man with high familiarity with computers, 22 years old; and a man with medium familiarity with computers, 60 years old. Neither of the three participants had any familiarity with conversational agents. These were selected because they represent a great variability. All three of them had a personal relation with the user who conducted the LTI, but not with the people who developed the project. All the three participants interacted with Arthur and answered the same questions that the participant who conducted the LTI for 41 days. Each interaction occurred for about 10-15 minutes, where all users were asked to interact with Arthur as they pleased, reacting to what they found interesting, what bothered them, or what was challenging for them. Nothing in special was asked: their only direction was to talk with Arthur as if they were talking with a new person they just met.

## V. RESULTS

In this section, we present the results obtained in our case study.

### A. LTI assessment

The results were observed for all the 3 questions/topics:

**Question 1 related to Effectiveness: It is defined as the accuracy and completeness with which users achieve specified goals in HCI. This should be responded considering Task completion, Accuracy and Recall [11], [12]:** For this question, the topics “Arthur’s effectiveness in learning something new” (Q1) and “Arthur’s effectiveness in showing something he learned” (Q2) were defined. For the most part of the LTI, these results were evaluated as Regular or Satisfactory, which means that Arthur was able to complete those actions. There were a few days where there was a bit harder to retrieve the information from Arthur or showing him something to learn, where the evaluation was Unsatisfied. The obtained average of Likert scale for this question is 3.7 (considering Very Unsatisfied as 1 and Very Satisfied as 5), with a standard deviation of 0.61. **Question 2 related to Efficiency: relates to the resources spent in relation to the accuracy and completeness with which the users achieve their goals. This should be responded considering Task completion time, Mental effort and Communication effort [11], [12]:** The topic “Effort applied while showing new things to Arthur” (Q1) was created to evaluate the way that the user can introduce new things to Arthur. In general, when the user says something like “Do you know my mother?”, Arthur tends to say “I do not know mother. Would you like to show me a picture?”. This topic was classified as Unsatisfied a few times, most of them for the first interaction days, as the user was not familiarized with the way Arthur works on this matter. After that, the interactions were evaluated as Regular, because

Arthur begins to know many terms, making the interaction much easier. In this topic the average value obtained in the Likert scale was 3.01, with a standard deviation of 0.24. **Question 3 related to Satisfaction: is defined as the degree to which user needs are satisfied when a product or system is used in a specified context of use. This should be responded considering ease-of-use, context-dependent questions, satisfaction before and during use, complexity control, physical discomfort of the interface, pleasure, the willing of use the chatbot again, and enjoyment and learnability [11], [12]:** Three topics were created to evaluate Arthur’s satisfaction: “Satisfaction with Arthur’s icebreakers” (Q1), “Satisfaction with Arthur’s answers to questions about him” (Q2) and “Satisfaction with Arthur’s answers to general questions” (Q3). In this case, the most chosen option was Unsatisfied due to some inconsistencies while interacting with Arthur, like the repetition of icebreakers and smalltalks. The average value for Likert scale is 2.2, with a standard deviation of 0.38.

Moreover, we can take a look at how the ratings changed over time. Again, considering Very Unsatisfied as 1 and Very Satisfied as 5, Figures 1 and 2 present how the answers varied for Effectiveness Q2 and Satisfaction Q3, which were identified as the two extremes.

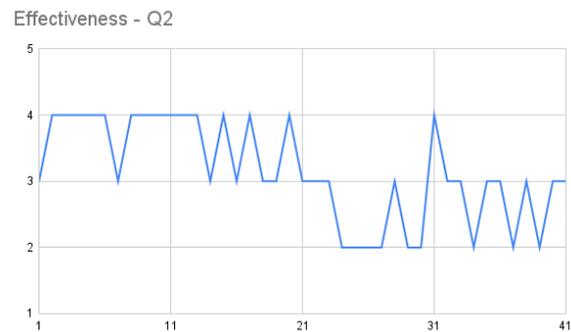


Fig. 1. Ratings changing for Effectiveness Q2 (showing something that Arthur learned). A great variance can be seen in the answers as interactions went by.

Figure 1 shows a great variance in the answers of the user concerning the second question of Effectiveness (Q2). It is possible to note that most answers were between Regular (3) and Satisfied (4) until about half of the interactions. Then, most answers were between Regular (3) and Unsatisfied (2), with one Satisfied (4). Based on this and the free text answers provided by the user, it seems that Arthur had little problem showing what he learned at the beginning. However, as he learned things and his memory became more complex, he began to find problems to show what he learned. We quote a part of one of the free text answers of the user: “...I showed him a picture of an anime, but when I try to retrieve it asking if he knows it, he shows a picture of himself, and if I only say the word anime, he only repeats it...”. We believe that, given the many relationships that Arthur has built in its memory between all that he learned, he had some problem deciding

which memory to retrieve when something was said to him, generating the discontentment.

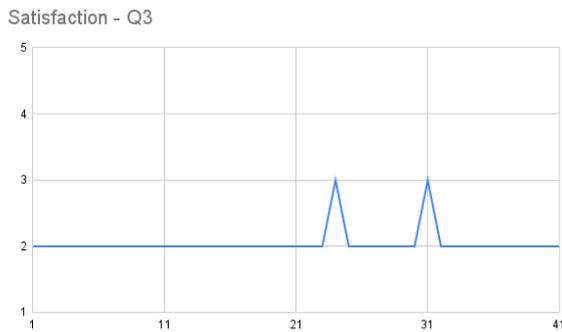


Fig. 2. Ratings changing for Satisfaction Q3 (Arthur's answers to general questions). Almost no variance can be seen in the answers as interactions went by, with only two answers breaking the pattern.

On the other hand, Figure 2 shows almost no variance in the answers of the user concerning the third question of Satisfaction (Q3). It is clear that most of the answers were Unsatisfied (2), with only two answers breaking the pattern as Regular (3). Once again, we quote a part of one of the free text answers of the user: "Arthur is kinda like saying random things. After saying Greetings, he started to say things like "Hum, show me some examples", or "I will check my agenda"...". It seems that Arthur found some problems when he needed to interact about something he had no knowledge about or was unsure on how to behave. Indeed, it could be also related with the situation reported in Figure 1: assuming that Arthur had some problem to decide which memory to retrieve, he would have problems to decide how to behave/answer as well.

### B. STI assessment

Although we are aware that a quantitative evaluation is not very useful having only three participants for the Short-Term Interactions, we believe it can be interesting to compare with the results achieved by the Long-Term Interaction. The three users were asked to interact with Arthur and, after that, answer the same questions used to compose the LTI Logbook. Some results are exposed as follows:

**Question 1, regarding Effectiveness:** When looking at Arthur's effectiveness in learning something new, all users were Satisfied or Very Satisfied, as they felt that Arthur was able to understand the things they said and showed to him. From the perspective of Arthur's effectiveness, the obtained average of Likert scale for this question is 3.8 (considering Very Unsatisfied as 1 and Very Satisfied as 5), with a standard deviation of 0.47. **Question 2, regarding Efficiency:** Although all users reported a difficulty on showing new things to Arthur (due to the time duration problem), all of them were Satisfied with this task. The obtained average of Likert scale for this question is 4, with a standard deviation of 0, since all three participants answered as being Satisfied.

**Question 3, regarding Satisfaction:** The satisfaction with Arthur had the most varied results in the STIs. The topic "Satisfaction with Arthur's answers to general questions" had the best performance (all three subjects were Satisfied). On the other hand, the topic "Satisfaction with Arthur's answers to questions about him" had answers varying from Unsatisfied to Very Satisfied. The obtained average of Likert scale for this question is 3.77, with a standard deviation of 0.57.

Comparing results of LTI and STI, in Question 1, it is possible to see that while the user conducting the LTI evaluated Arthur's effectiveness as Regular most of the time, the users which conducted the STIs evaluated Arthur's effectiveness as Regular, Satisfied and Very Satisfied. Regarding the comparison of Question 2 results, provided by both LTI and STIs users, it is possible to notice that while the user conducting the LTI evaluated Arthur's efficiency as Regular most of the time, all the users which conducted the STIs were Satisfied with the efficiency. Finally, with respect to the Question 3, provided by both LTI and STIs users, it is possible to say that the user conducting the LTI evaluated its satisfaction with Arthur as Unsatisfied most of the time, the users which conducted the STIs evaluated their satisfaction with Arthur as Satisfied most of the time, even having a Very Satisfied answer.

### C. Usability Evaluation

After having analysed LTI and STIs, we categorized every suggestion based on Nielsen's 10 Usability Heuristics for User Interface Design [9], as well as Severity Ratings for Usability Problems [10], as it was found to be one of the best models for evaluating interfaces. The evaluation goes as follows: **1st Heuristic – Visibility of system status:** Arthur does not show the latest user message in the conversation field, which can give the impression that it was not correctly sent. Some things can be done to solve this: the whole conversation can be shown in the text area; a notification sound can be triggered every time that the user sends a successful reply to Arthur; a visual notification that gives the impression that Arthur is thinking could be added. **Severity: 3.** **2nd and 3rd Heuristics – Match between system and the real world; User control and freedom:** These heuristics were not violated. **4th Heuristic – Consistency and standards:** The button used to send messages to Arthur is not intuitive and it is hard to understand (it is pictured as a button with only a "V" inside it). The ideal option, in this case, would be to use a more standard button, like putting "Send/Send Message" on the button or use a button similar to the ones available in social media, which tends to be more comfortable and more used by people. **Severity: 1.** **5th and 6th Heuristics – Error prevention and Recognition rather than recall:** These heuristics were not violated. **7th Heuristic – Flexibility and efficiency of use:** Adding a faster and easier way for the user to send the messages to Arthur would be a good way to increase its efficiency. Instead of having to move a hand to the mouse and then click on the button, just pressing the "Enter" button on the keyboard would be better and more natural for the user. **Severity: 2.** **8th Heuristic – Aesthetic and minimalist design:** Arthur's

color palette is not a problem, but some components could be changed to make Arthur more aesthetic. For example: changing the format, position and size of the text box; buttons positioning; positioning of the webcam. Also, one of the comments specifies specifically the distribution of components on the interface. **Severity: 0. 9th Heuristic – Help users recognize, diagnose, and recover from errors:** Even though rare, the errors displayed by Arthur could be better handled. **Severity: 3. 10th Heuristic – Help and documentation:** Arthur does not have a “help” tool or documentation available for its interface. It would be interesting to add this kind of information in case of an eventual need to know more about Arthur’s functionalities in the future. **Severity: 2.**

In short, it is possible to notice that most of the suggestions were classified as low priority (0, 1 and 2 in Severity Ratings scale), with only two suggestions being classified as a major problem (3 in Severity Ratings scale) and none as a Usability Catastrophe (4 in Severity Ratings scale).

## VI. FINAL CONSIDERATIONS

This work discusses a case study performed to evaluate ECAs using Long-term (LTI) and Short-term Interactions (STIs). With the results achieved, we evaluated the users perception while interacting with Arthur, considering the effectiveness, efficiency, and satisfaction concepts for ECAs evaluation, described in the studies from Castro et al. [11] and Tullis and Albert [12]. We included LTI and STIs in the methodology because, as we could see in our experiments, different periods of interaction can make users focus on different situations. By performing and comparing the LTI with the STIs, it was possible to notice that STI brought more satisfaction for the users at the end of the interactions and their qualitative comments focused more on concrete interface aspects. On the other hand, the LTI user “learned” how to use the ECA with the provided interface and focused more on communication aspects, such as ice breakers and small talks. So, we believe that both interactions are interesting and that is why we propose them both in our methodology to evaluate ECAs.

As future work, we are planning to use the findings of this work to improve Arthur, focusing on the naturalness of interaction and on the UX evaluation. After that, we also plan to conduct other LTI and STIs in order to compare with the results of this work and evaluate if the changes indeed improved the behavior and the effectiveness of Arthur. Moreover, it is possible that, with time, the participant who conducted the Long-Term Interaction (LTI) learned how to work around some aspects of Arthur which were causing issues. In this scenario, the later comfort with the ECA would be caused by such adaptation, not by improvements resulted from Arthur’s learning or behavior. Our future work also focuses on these new interactions as a way to measure such possibility.

Concerning Short-Term Interactions, it is important to discuss some limitations which could have influenced the results. It is not possible to know what were the expectations of

the three users about their interaction with Arthur, especially because none of them had any previous experience interacting with conversational agents. Therefore, it is possible that those three users perceived such interaction as an interesting experience and, even noticing some problems, it could also have been perceived as “part of the fun” of a new experience, which could have been reflected in the better evaluation perceived in the Results. In a future work, we plan to mitigate such possibility by conducting more STIs, including people which are already used with conversational agents.

## REFERENCES

- [1] B. Wimmer, B. Wöckl, M. Leitner, and M. Tscheligi, “Measuring the dynamics of user experience in short interaction sequences,” in *Proceedings of the 6th Nordic Conference on Human-Computer Interaction: Extending Boundaries*, 2010, pp. 825–828.
- [2] Y. Chen, A. Naveed, and R. Porzel, “Behavior and preference in minimal personality: A study on embodied conversational agents,” in *International conference on multimodal interfaces and the workshop on machine learning for multimodal interaction*, 2010, pp. 1–4.
- [3] P. Knob, W. S. Dias, N. Kuniechick, J. Moraes, and S. R. Musse, “Arthur: a new eca that uses memory to improve communication,” in *2021 IEEE 15th International Conference on Semantic Computing (ICSC)*. IEEE, 2021, pp. 163–170.
- [4] Z. Ruttikay, C. Dormann, and H. Noot, “Embodied conversational agents on a common ground: A framework for design and evaluation,” in *From brows to trust: evaluating embodied conversational agents*. Kluwer, 2004, pp. 27–66.
- [5] Y. Zhang, X. Gao, S. Lee, C. Brockett, M. Galley, J. Gao, and B. Dolan, “Consistent dialogue generation with self-supervised feature learning,” *arXiv preprint arXiv:1903.05759*, 2019.
- [6] H. Zhou, M. Huang, T. Zhang, X. Zhu, and B. Liu, “Emotional chatting machine: Emotional conversation generation with internal and external memory,” in *Thirty-Second AAAI Conference on Artificial Intelligence*, 2018.
- [7] Ö. N. Yalçın, “Empathy framework for embodied conversational agents,” *Cognitive Systems Research*, vol. 59, pp. 123–132, 2020.
- [8] P. Morville, “Experience design unplugged,” in *ACM SIGGRAPH 2005 Web Program*, ser. SIGGRAPH ’05. New York, NY, USA: Association for Computing Machinery, 2005, p. 10–es. [Online]. Available: <https://doi.org/10.1145/1187335.1187347>
- [9] J. Nielsen, “Enhancing the explanatory power of usability heuristics,” in *Proceedings of the SIGCHI conference on Human Factors in Computing Systems*, 1994, pp. 152–158.
- [10] —, “Severity ratings for usability problems,” *Papers and Essays*, vol. 54, pp. 1–2, 1995.
- [11] J. W. Castro, R. Ren, S. T. Acuña, and J. d. Lara, “Usability of chatbots: A systematic mapping study,” *Repositorio Académico Institucional, Universidad de Atacama*, 2019.
- [12] W. Albert and T. Tullis, *Measuring the user experience: collecting, analyzing, and presenting usability metrics*. Newnes, 2013.
- [13] T. Bickmore, D. Schulman, and L. Yin, “Maintaining engagement in long-term interventions with relational agents,” *Applied Artificial Intelligence*, vol. 24, no. 6, pp. 648–666, 2010.
- [14] S. Babu, S. Schmugge, T. Barnes, and L. F. Hodges, ““what would you like to talk about?” an evaluation of social conversations with a virtual receptionist,” in *International Workshop on Intelligent Virtual Agents*. Springer, 2006, pp. 169–180.
- [15] D. Benyon and O. Mival, “From human-computer interactions to human-companion relationships,” in *Proceedings of the First International Conference on Intelligent Interactive Technologies and Multimedia*, ser. IITM ’10. New York, NY, USA: Association for Computing Machinery, 2011, p. 1–9. [Online]. Available: <https://doi.org/10.1145/1963564.1963565>
- [16] T. Castle-Green, S. Reeves, J. E. Fischer, and B. Koleva, “Decision trees as sociotechnical objects in chatbot design,” in *Proceedings of the 2nd Conference on Conversational User Interfaces*, 2020, pp. 1–3.
- [17] T. Kanda, R. Sato, N. Saiwaki, and H. Ishiguro, “A two-month field trial in an elementary school for long-term human-robot interaction,” *IEEE Transactions on robotics*, vol. 23, no. 5, pp. 962–971, 2007.