

Memory Continuity as a Foundation of User Trust in Emotional AI Avatars: A Literature Review

Tianyu Zhang

Department of Computer Science, University of Auckland, Auckland, 1010, New Zealand

tzha118@aucklanduni.ac.nz

ABSTRACT

Emotional AI avatars are increasingly used for companionship and emotional support, where users may expect continuity across interactions, including being remembered over time. This makes memory continuity relevant to user trust. Prior research shows that people respond to digital systems in social ways, that repeated interaction can support relationship-like expectations, and that trust in emotional support AI involves more than technical performance alone. Recent work also suggests that memory can influence how users judge AI systems, including perceived intelligence, likeability, and trustworthiness. However, existing literature has not clearly explained how memory continuity functions as a condition of user trust in emotional AI avatars. This review addresses this gap by synthesizing research on social responses to AI, long-term human-AI relationships, emotional support AI, and AI memory through a user trust lens. It argues that continuity in remembered context helps shape whether interaction feels stable, coherent, and personally responsive over time. When continuity is maintained, trust may be supported; when it breaks, trust may weaken. Taken together, the reviewed literature suggests a trust-centred way to understand how continuity and forgetting shape emotionally supportive AI interaction.

Keywords: Emotional AI Avatars; Memory Continuity; User Trust; AI Forgetting; Human-AI Relationship.

1. INTRODUCTION

Imagine that users often ask the AI system for emotional support. In several conversations, users share their personal concerns, preferences and important life experiences. Later, the system no longer remembers these details. In this case, the problem is not just technical. It can also change the user's understanding of interaction and whether the system still looks trustworthy. This situation explains why the continuity of memory may be important when using emotional support of artificial intelligence.

In this review, "emotional AI avatar" is used as the main term of AI system for company, emotional support and daily relationship interaction. Direct research on emotional AI avatar is still limited. Therefore, this review also uses social chat bots, embedded agents and the close work of generating artificial intelligence for emotional support. These studies are helpful to explain the same interactive problems. This limitation is important because this review focuses on how the continuity of memory background affects users' confidence in using artificial intelligence for emotional support.

Existing research has already given us some useful ideas. Early research shows that people often respond to computers and other digital systems in social ways, even when they know the system is not human^[1]. Another body of work shows that repeated interactions can create familiarity, continuity, and expectations that resemble a relationship over duration^[2]. Related research on social chatbots also suggests that some users are beginning to treat these systems as companions, and that trust and confidences can grow with continued use^[3]. In addition, recent research suggests that memory can influence how users judge a chatbot, including its intelligence, likeability, and reliability^[4]. Taken together, these studies point a clear path: when AI is socially interpreted, and when the interaction becomes repetitive and personal, the stored context can be part of how users judge the quality and reliability of the system.

However, current literature has not yet clearly explained how memory continuity shapes trust through relational continuity in emotional interaction with AI. Existing studies discuss emotional support, companionship, trust, and AI memory, but they often do so separately. Research on emotional AI often focuses on empathy, support, or companionship. Research on AI memory often focuses on personalisation, context retention, or technical function. Research on trust is often broader and does not explain which kinds of remembered continuity matter most, or how breaks in continuity may weaken trust over time. So, the problem is not only that direct evidence is limited. The literature also lacks clear conceptual integration.

What is still missing is a focused account of how memory continuity should be understood as a trust-relevant part of emotionally supportive AI interaction.

In order to fill this gap, this review uses user trust as its main analysis tool, and studies how memory continuity affects the trust of emotional AI avatars. His main question is: according to the existing literature, how does memory continuity affect users' trust in emotional AI avatars? Because the direct research on emotional AI avatar is still limited, this review also includes the research on social chat robots, embedded agents and generated AI for emotional support. These studies are included because they examine repetitive, personal and emotionally meaningful similar forms of interaction between humans and artificial intelligence. According to this review, in the emotional AI environment, remembering the context helps to determine whether the system looks stable, responsive and reliable.

2. EMOTIONAL AI AVATARS AND LONG-TERM HUMAN-AI RELATIONSHIPS

2.1 Emotional AI avatars as more than tools

The incarnation of Emotional AI is not just to help us complete the task. Users can use them to make friends, get emotional support or talk about everyday things. In these cases, the value of the system is not only because it is useful or effective, but also because it is always available, familiar and encouraging. A recent review paper shows that generative AI for emotional support is increasingly discussed in terms of trust and relationship experience, not just technical performance^[5]. For this reason, the incarnation of emotional AI should not be understood in the same way as ordinary task-based AI.

This difference is important because using emotional AI can create expectations for memory continuity. In the task-based framework, users are most concerned about whether the system gives correct or useful responses. In the emotional support framework, users can also expect the system to remember previous interactions, respond consistently, and maintain a stable sense of familiarity over time. This is not to say that users completely confuse AI with human partners. On the contrary, it shows that when the interaction is repeated and has emotional significance, the function of the system becomes more extensive. Under these conditions, the continuity of memory can be a part for users to judge whether the system is trustworthy and supported ^[2].

2.2 Repeated interaction and social interpretation

Repeated interaction helps to explain why emotional AI avatars are handled differently from ordinary software. Early research on HCI shows that people often respond to computers in a social way, even though they know that the system is not human ^[6]. Later research shows that repetitive interaction can promote familiarity and expectation, similar to a relationship ^[2].

The research on social chat robots provides useful evidence for this. Although these studies mainly focus on conversational chat robots rather than embedded emotionalization, they are still related because both systems involve repetitive social interaction, personal reaction and emotionally meaningful communication. For example, some users have established meaningful relationships with chat bots, and trust and self-disclosure have increased over time ^[3]. Subsequent research also found that some users began to describe chat bots as companions or friends ^[7]. These results cannot be directly applied to all emotional avatars, because avatars may include other functions, such as voice, face, body movements or visual presence. However, the research on chat bots still helps to show how repeated interactions can make users explain AI systems in a more relational way.

At the same time, the current books and studies do not clearly indicate that users are beginning to feel more familiar or expect how much interaction a relationship needs. The research we have seen here shows that repetitive interactions are usually important, but they do not give the exact number or duration that applies to all systems and all users. Therefore, repetitive interaction should be regarded as something that increases gradually, rather than a fixed number.

It is also important to note that the evidence seen in this section mainly comes from a culture. Most of the books and studies discussed here come from research traditions and publishing places, which are mainly western. This is not to say that the result is invalid, but it does mean that we should be careful when drawing conclusions about the companion AI relationship in different cultural backgrounds. At present, the research literature does not provide enough evidence from many cultures to prove how cultural differences change the way people accept artificial intelligence as a relationship partner.

2.3 Remembered continuity as part of relational trust

For this review, the key question is not only whether users can form emotional connections with AI, but why it is important to remember continuity in these relationships. When users often interact with emotional artificial intelligence avatars, they can expect the system to remember past conversations, maintain a consistent interaction style, and respond in a way that matches previous conversations. In this case, memory is part of the user experience to determine whether the system looks consistent, reliable and personal.

The research on generating the trust of emotional support in artificial intelligence supports this point. In these cases, trust involves not only the ability, but also whether the system looks safe, predictable and supported ^[5]. Remember that continuity can support this form of trust by helping the system look stable over time. On the contrary, in the use of emotional AI, remember that the interruption of continuity may be more obvious than normal software use, because users can relate past interactions with their current expectations.

However, the literature does not show that all users have formed a strong emotional connection with AI, nor does it show that all emotional AI systems are used in the same way. Nor does it show how cultural differences or differences in interaction frequency change this process. Therefore, the narrower conclusion in this section is that in repetitive and emotionally meaningful interactions, the continuity of memory may become an important factor in how users judge their trust in emotional AI avatars.

3. MEMORY CONTINUITY AS A FOUNDATION OF USER TRUST

3.1 Defining memory continuity in emotional AI

In this review, continuous memory means that the system seems to be able to stably retain important personal information again and again. This is not to say that this system really has memories like humans. But this means that for us users, the system seems to be able to remember things that were useful before, such as our preferences, past conversations or emotionally important details, and then use them. This practical definition is more in line with the current literature, rather than strongly explaining how AI remembers the inside, because many studies focus on user experience and perception, not just technical architecture ^[5].

This definition is important because the word “memory” is used in different ways in books and articles. In some studies, memory is considered as a function that contributes to personalization or continuity. In other studies, we check it through semantic long-term memory or users’ feelings, that is, chat robots “remember” them ^[4]. Because of these different uses, this review does not regard memory continuity as a single fixed technical mechanism. Instead, she uses the term as a concept at the interactive level: to maintain the continuity of relevant personal background over time ^[5].

An accurate definition is also necessary because the available evidence is still limited. There is not much direct research on emotional AI avatar, and some of the most relevant evidence comes from very close systems, such as social chat bots and generated AI for emotional support. For this reason, the argument here should be interpreted as a targeted explanation of the current evidence, rather than the final report that applies to all AI systems in the same way ^[5].

3.2 Trust beyond technical competence

In the emotional AI system, trust depends not only on technical ability. In task-based systems, users usually judge trust according to accuracy, speed or practicality. But in emotional support interaction, trust also depends on the fact that the system looks safe, predictable and kind. Volpato ^[5] and others said that confidence in generating emotional support for artificial intelligence comes from many fields and cannot be attributed to simple ability judgment. This means that in this case, trust is formed by a wider user interaction experience, not just the technical accuracy of the response. This broader view of trust is also consistent with HCI’s early work on the long-term relationship between people and computers.

Recent evidence further illustrates why memory should not be regarded as a simple solution to trust. Yan et al. ^[4] found that adding semantic long-term memory makes the chat robot smarter and friendlier, but it not only improves the reliability. This result is very important, because it shows that memory can change the perception of the system without automatically generating trust. Therefore, this review regards trust as a specific analytical lens, and memory continuity as a possible condition to shape trust, rather than the direct cause of trust.

3.3 How continuity supports trust over time

Based on this broader view of trust, continuous memory can support trust through three interrelated mechanisms: perceptual stability, recognition and relationship consistency.

3.3.1 Perceived stability

First of all, continuity can help everything look stable. When a system seems to remember what we said before, its later reaction seems less random and more related to the user's history. This is very important when seeking emotional support, because users can not only rely on the system to obtain information, but also get a stable and reliable interactive experience. In this sense, the context we remember can help the system look more predictable over time.

3.3.2 Recognition

Secondly, continuity can help identify. In repeated interactions, users may not only want the right answers. They may also want signs that the system still "knows" in a limited but meaningful sense. Research on intelligent virtual agents shows that familiarity can improve users' perception of these systems, including trust-related judgments^[8]. This does not prove that memory continuity and familiarity are the same thing. However, this supports the view that stable identification of interaction is very important for system evaluation. Because of this, over time, memory continuity can support trust by making AI more familiar and personalized.

3.3.3 Relational coherence

Thirdly, continuity can support the consistency of relationship. Previous research on long-term human-computer relationship shows that repetitive interaction can maintain a sense of continuity and social value^[2]. Recent research on social chat robots also found that with the development of relationships, trust and self-disclosure may increase over time^{[3][7]}. Generally speaking, these studies show that trust in emotional AI is not only based on the unique moment of successful interaction. It may also depend on the fact that the system seems to be consistent from one meeting to another.

Together, these three mechanisms put forward a simple conceptual model: memory and continuity can support users' confidence by making the interaction with emotional AI look stable, recognizable and consistent in relationship over time. This model does not mean that the continuity of memory guarantees trust. On the contrary, it locates the memory background as the relevant condition of trust in a wider interactive system.

A restriction is also important. Persistent memory does not always support trust. If the system has a bad memory, misuses past information, or remembers more than the user expected, then memory may cause discomfort rather than trust. Misinterpreted memory may also make users uncertain about what the system knows and how to use this information. For this reason, continuity should not be understood as automatically positive. It can support trust when it is accurate, relevant and meets users' expectations, but it can weaken trust when it looks incorrect, excessive or unclear.

Therefore, memory continuity is best understood here as a condition that can support user trust, rather than guarantee it. It may foster trust by making the system seem stable, familiar, and personally responsive, but the literature does not show that memory alone is sufficient. Other things, such as interface style, emotional tone, transparency and user expectations, may also be important. What we know at present shows this point more accurately: when we interact with artificial intelligence repeatedly and emotionally, it is very important for it to remember the facts of the context to understand how trust is formed, maintained or weakened over time.

4. AI FORGETTING AND TRUST DISRUPTION

4.1 Forgetting as a relational problem

AI forgetting is usually regarded as a technical problem, such as storage, context management or response quality. However, in the emotional AI system, forgetting it may also become a breakdown of the relationship. When users often talk to the system in a personal and emotional way, they may start waiting for the conversation to be remembered. If the system doesn't remember important personal information in the future, users may not think it's just a system error. They may also regard it as a lack of attention, balance and personal recognition^[2, 5].

This difference is important, because the trust in emotional support AI comes not only from its correct answer. It also depends on whether the communication with AI looks safe, predictable and supportive^[5]. The research on social chat bots shows that some users begin to regard them as companions, and their trust may increase with time^[3,7]. In this case, if AI

forgets, it may break the feeling of continuous dialogue and weaken users' confidence that the system is paying attention to its history. Therefore, when emotion-assisted artificial intelligence forgets, it is not only a technical problem, it may also be the reason why users no longer trust.

4.2 Expectation violation and weakened trust

To understand the influence of artificial intelligence exercise, we can consider expectation. Over time, users will form expectations about how the AI system will respond, what it will remember, and whether it will remain unchanged. Expectation violation theory says that when communication behavior does not meet their expectations, people will react^[6]. Here, AI forgetting becomes a trust problem when it violates the user's expectation of continuity.

In the case of embedded proxy, Burgoon et al.^[6] show that violating users' expectations may affect their views on proxy during interaction. This does not prove that every artificial intelligence exercise will reduce trust in the same way, but it does support a broader view that unsatisfied expectations are important in the communication between people and artificial intelligence. If users expect continuity, and the system responds as if the previous interaction never happened, then trust may decline because the interaction no longer looks stable or reliable.

Recent research on AI memory also points in this direction. Yan et al.^[4] found that memory changed users' judgment on chat bots, but the effects were varied and not entirely positive. Their research results show that memory can increase perceptual intelligence and compassion, and at the same time make more complicated judgments on reliability. This result is very important for this review, because it means that memory shapes the standard of judging system. Once users begin to expect continuity, the lack of continuity may become more obvious in the future.

For this reason, the forgetting of AI can be understood as a potential source of undermining trust. The current evidence does not support the strong statement that forgetting always leads to the collapse of trust. However, they support a narrower conclusion: in repetitive and emotionally meaningful interactions, forgetting may destroy the continuity of expectations and make the system look less reliable, less focused, or less personal, thus weakening trust^[5].

4.3 Different types of AI forgetting

The literature has not provided the typology of AI in the emotional interaction accepted by everyone. Therefore, the following categories should be understood as the analytical structure used in this review, rather than the fixed classification established by previous studies. This structure is still very useful, because the current literature shows that not all forms of forgetting are important to users. Some forms may only affect a little confidence, while others may be more important because they are related to identity, familiarity or emotional importance^[4, 5].

The first type is the memory loss of situational interaction. This means that we have forgotten unimportant and emotionless details in ordinary conversation, such as a small topic mentioned earlier or the usual communication at the last meeting. In many cases, this forgetting can be regarded as a small limitation. Current research does not systematically compare this type of forgetting with other types of forgetting, so its influence should not be exaggerated. However, based on the extensive work on continuity and repetitive interaction, it is reasonable to think that this form of forgetting may have less influence on trust than the more personal form of memory decline^[2].

The second is memory loss related to preference. This means that we forget the stable interests, habits or preferences that often come from users. This type may be more important because it will affect whether users feel recognized over time. Related research shows that familiarity may affect how we view intelligent virtual agents, including trust-related judgments^[8]. Although familiarity is not exactly the same as memory continuity, this connection is related: if the system cannot maintain the user's stable preferences, it may look less familiar and less sensitive to the user's response. For this reason, this forgetting may have a greater impact on trust than losing unimportant background details, although the direct evidence of emotional AI incarnation is still limited.

The third type is relational or commitment memory loss. This refers to forgetting emotionally important events, vulnerable disclosures, or shared commitments from earlier interaction. This is the most influential form of forgetting, because it affects the emotional meaning of the relationship itself. Although the current research has not provided enough direct evidence to show that this effect is always the same, the existing work on long-term human-computer relationship, emotional support for AI and violation of expectations shows that the continuous rupture in the personal sense may be particularly disturbing^[2, 5, 6].

In a word, these three forms of forgetting support a main point: not all forgetfulness is the same. The possible influence of forgetting on trust depends on the importance of emotion and relationship of lost information. Some memory failures can

be regarded as minor inconveniences, while others can really reduce trust by breaking continuity on a more personal or meaningful level. At present, this statement must be understood as a conceptual explanation of the literature, rather than a completely definite empirical conclusion. Even so, organizing discussions in this way helps to clarify why artificial intelligence should not only be regarded as a technical failure, but also as an HCI problem that affects trust.

5. RESEARCH GAPS AND DESIGN IMPLICATIONS

5.1 Gaps in current literature

The current literature shows five main gaps. In short, these gaps indicate that memory continuity, user trust and emotional support interaction with AI are not yet fully connected.

First of all, literature lacks conceptual integration. The research on emotional AI usually focuses on empathy, partnership or anthropomorphic design, while the research on AI memory usually focuses on personalization, context preservation or system performance. Therefore, memory continuity is rarely regarded as the core condition of user trust in emotional support interaction.

Secondly, there is limited direct evidence about emotional AI avatar as an independent category. Many available evidences come from neighboring systems, such as social chat bots, embedded agents, intelligent virtual agents or generative AI tools for emotional support. These studies are useful, but their results cannot be fully extended to the emotional AI avatar, because the system is different in packaging, interactive style and user expectations.

Third, AI memory is usually regarded as a system function rather than a user interaction problem. The existing work shows that users may have an incomplete understanding of the content of the agent's memory, the organization of the memory and how the memory information affects the subsequent behavior^[9, 10]. This means that memory needs to be studied not only as a back-end function, but also as a part of user experience.

Fourthly, forgetting AI as a trust problem is still rarely studied. Recent research shows that memory can affect how users judge chat bots, including their intelligence, compassion and reliability^[4]. However, little is known about how different forms of forgetting affect confidence over time, especially in repetitive and emotional interactions.

Fifth, expectation management and memory transparency are still undeveloped design issues. Users may misunderstand whether the AI system goes from one session to another or doesn't remember information at all during the session. Therefore, future research needs to study how the emotional AI system clearly conveys the limitations of memory and helps users form more accurate expectations^[11].

In a word, these gaps indicate that the field needs to explain the memory continuity of HCI as an interactive problem related to trust more clearly. Therefore, future research should not only study whether the emotional AI system can remember, but also study how users understand, manage and respond to memory and exercise.

5.2 Design implications for future emotional AI avatars

Based on the above literature, several design concepts can be put forward for the future emotional AI incarnation. These ideas should be understood as a reasonable design direction based on current evidence, rather than clear rules. At present, the literature is still under development, and the direct verification of emotional AI avatar is still limited.

First, future systems should support memory transparency. If the continuity of memory affects trust, users need to better understand the content of system memory and how to use this information in the future. Huang et al.^[10] think that the current dialogue agents usually can't provide users with effective methods to check or manage their memories, which may lead to bad mental models and interaction interruption. Similarly, Jones^[9] and others found that users often have incomplete expectations about the role of proxy memory and how memory affects their behavior. In short, these studies show that emotional AI systems should not completely hide their memories. On the contrary, they can benefit from the function, making the memory more visible, readable and understandable to users.

Secondly, the system should help to understand what they remember. If we think the system remembers more or better than the actual situation, then trust will be broken. The study of communication ability is also looking for ways to explain what agents can do in interaction, whether it can improve transparency and how we understand it^[11]. In the emotional AI system, this means that they may need to explain more clearly what they keep: for example, whether the information is kept in only one session, multiple sessions, or not at all. This design will not eliminate technical limitations, but it can reduce the confusion of continuous content.

Third, the future system should give users more memory control. Jones and others said that users want to organize and control their memories in a way suitable for their different tasks and activities. Their research results show that users may prefer to classify their memories and access them at different levels according to the context. Huang et al.^[10] have a similar view in their memory sandbox design, which allows users to check, manipulate and share memory content through dialogue. For emotional AI avatars, this means that user-visible controls may be useful, such as allowing users to view, edit, separate or delete memory information. This control can help build trust by improving user behavior and reducing uncertainty about the content stored in the system.

Fourth, the system must carefully use personalization, not automation. Some evidence shows that personalization can enhance trust when AI seems to know the user's situation better. In the online health consultation experiment, Qin et al.^[12] found that higher personalization and nursing increased cognitive and emotional confidence and led to greater adoption. This result comes from the background of health consultation, so it should not be transferred carelessly. But it puts forward a useful principle: personalization can support trust, when it is experienced as carefulness and care, but personalization alone is not enough. In the emotional incarnation AI, the context of memory should be used in a way that seems relevant, appropriate and not too intrusive.

Fifth, the system can benefit from the repair response after forgetting. The direct empirical work on this point is still limited, so this idea must be regarded as a cautious design inference, rather than an established result. However, the literature studied in this paper shows that it may be important to forget it because it destroys continuity and weakens trust. If so, future systems may need ways to identify their limitations, clarify what has not been preserved, and restore the context when continuity fails. This stems from the wider discovery of trust, transparency and violation of expectations, although the specific repair strategy of emotional AI avatar still needs more direct testing^[5,6].

Generally speaking, these design concepts show us a simple but important principle: if we use the emotional AI avatar for a long time and use it for important personal affairs, then trust can not only come from emotional tone or superficial goodwill. It may also depend on whether users can understand, predict and influence how memory works over time. All these ideas together show that emotional AI memory should be transparent, limited, controllable, personalized and repairable. Therefore, in the embodiment of emotional AI, memory should be designed as a responsible interactive layer, rather than a hidden technical ability. Therefore, future design should not only consider whether the system remembers, but also consider how to communicate, limit, control and repair memory when continuity is destroyed.

6. CONCLUSION

This review studies how memory continuity affects users' trust in emotional AI avatars. In the text, one thing is clear: when the interaction with artificial intelligence becomes repetitive, personal and emotional, users may start to see the system as more than just a tool. In these cases, confidence depends not only on technical performance, but also on the fact that the system looks stable, predictable and concerned by individuals over time^[2,5].

According to the research discussed in this paper, memory continuity can be understood as a condition that contributes to this form of trust. When a system seems to be able to remember the personal context from one interaction to another, it may look more consistent and familiar to users. At the same time, memory does not automatically build trust. Existing evidence shows that memory can improve some judgments, such as perceived intelligence or loveliness, and at the same time make judgments related to trust more complicated^[4]. For this reason, it is better to understand memory continuity as a factor that can help build trust, rather than a factor that guarantees trust.

The commentary also said that artificial intelligence exercise should not be regarded as just a technical issue. In the case that artificial intelligence provides emotional support, it may be important to forget it, because it destroys continuity and weakens the user's feeling that the system is careful and reliable, and remembers past interactions. Although there is still little direct evidence forgotten in the emotional avatar AI, the related work on the long-term relationship between human beings and computers, social chat robots, embedded agents and emotional support AI shows that continuity and expectation are very important for users to make judgments^[3,5,6]. This leads to a more accurate conclusion: in the interaction with emotional importance, forgetting it may be the root of losing trust, not just a systematic error.

This review provides a trust-centric framework, which connects memory continuity, AI preprocessing and user trust to emotional AI avatars. It shows that memory helps to determine whether the interaction with AI is stable, familiar and consistent over time. In doing so, the review highlights several defects in the current literature, including limited conceptual

integration, limited direct evidence of emotional AI incarnation, and insufficient attention to user-side issues such as memory transparency, user control and expectation management.

Overall, the main implication is that designing emotional AI avatars requires more than improving memory as a technical capability. It also requires designing memory as part of the interaction experience. Future systems should consider not only whether they remember, but also how memory is communicated, bounded, controlled, and repaired when continuity breaks. This is important because, in long-term emotionally supportive AI interaction, trust may depend not only on what the system remembers, but also on how users understand and manage that memory.

REFERENCES

- [1] Nass, C., Steuer, J., & Tauber, E. R. (1994, April). Computers are social actors. In *Proceedings of the SIGCHI conference on Human factors in computing systems* (pp. 72-78).
- [2] Bickmore, T. W., & Picard, R. W. (2005). Establishing and maintaining long-term human-computer relationships. *ACM Transactions on Computer-Human Interaction (TOCHI)*, 12(2), 293-327.
- [3] Skjuve, M., Følstad, A., Fostervold, K. I., & Brandtzaeg, P. B. (2021). My chatbot companion-a study of human-chatbot relationships. *International Journal of Human-Computer Studies*, 149, 102601.
- [4] Yan, Z., Fischer, J. E., & Clos, J. (2025, July). Remembering Things Makes Chatbots Sound Smarter, but Less Trustworthy-A Pilot Study. In *Proceedings of the 7th ACM Conference on Conversational User Interfaces* (pp. 1-8).
- [5] Volpato, R., DeBruine, L., & Stumpf, S. (2025). Trusting emotional support from generative artificial intelligence: a conceptual review. *Computers in Human Behavior: Artificial Humans*, 100195.
- [6] Burgoon, J. K., Bonito, J. A., Lowry, P. B., Humpherys, S. L., Moody, G. D., Gaskin, J. E., & Giboney, J. S. (2016). Application of expectancy violations theory to communication with and judgments about embodied agents during a decision-making task. *International Journal of Human-Computer Studies*, 91, 24-36.
- [7] Skjuve, M., Følstad, A., Fostervold, K. I., & Brandtzaeg, P. B. (2022). A longitudinal study of human-chatbot relationships. *International Journal of Human-Computer Studies*, 168, 102903.
- [8] Poletaykina, A., Kruse, L., & Steinicke, F. (2025, September). Knowing is Trusting! The Influence of Familiarity on the Perception of Intelligent Virtual Agents. In *Proceedings of the 25th ACM International Conference on Intelligent Virtual Agents* (pp. 1-11).
- [9] Jones, B., Stemmler, K., Su, E., Kim, Y. H., & Kuzminykh, A. (2025, April). Users' Expectations and Practices with Agent Memory. In *Proceedings of the Extended Abstracts of the CHI Conference on Human Factors in Computing Systems* (pp. 1-8).
- [10] Huang, Z., Gutierrez, S., Kamana, H., & MacNeil, S. (2023, October). Memory sandbox: Transparent and interactive memory management for conversational agents. In *Adjunct Proceedings of the 36th Annual ACM Symposium on User Interface Software and Technology* (pp. 1-3).
- [11] Reimann, M. M., Kunneman, F. A., Oertel, C., & Hindriks, K. V. (2025, July). Transparent Conversational Agents: The Impact of Capability Communication on User Behavior and Mental Model Alignment. In *Proceedings of the 7th ACM Conference on Conversational User Interfaces* (pp. 1-12).
- [12] Qin, H., Zhu, Y., Jiang, Y., Luo, S., & Huang, C. (2024). Examining the impact of personalization and carefulness in AI-generated health advice: Trust, adoption, and insights in online healthcare consultations experiments. *Technology in Society*, 79, 102726.