

Conceptual Metaphors Impact Perceptions of Human-AI Collaboration

PRANAV KHADPE, Stanford University, USA

RANJAY KRISHNA, Stanford University, USA

LI FEI-FEI, Stanford University, USA

JEFFREY T. HANCOCK, Stanford University, USA

MICHAEL S. BERNSTEIN, Stanford University, USA

With the emergence of conversational artificial intelligence (AI) agents, it is important to understand the mechanisms that influence users' experiences of these agents. In this paper, we study one of the most common tools in the designer's toolkit: conceptual metaphors. Metaphors can present an agent as akin to a wry teenager, a toddler, or an experienced butler. How might a choice of metaphor influence our experience of the AI agent? Sampling a set of metaphors along the dimensions of warmth and competence—defined by psychological theories as the primary axes of variation for human social perception—we perform a study ($N = 260$) where we manipulate the metaphor, but not the behavior, of a Wizard-of-Oz conversational agent. Following the experience, participants are surveyed about their intention to use the agent, their desire to cooperate with the agent, and the agent's usability. Contrary to the current tendency of designers to use high competence metaphors to describe AI products, we find that metaphors that signal low competence lead to better evaluations of the agent than metaphors that signal high competence. This effect persists despite both high and low competence agents featuring identical, human-level performance and the wizards being blind to condition. A second study confirms that intention to adopt decreases rapidly as competence projected by the metaphor increases. In a third study, we assess effects of metaphor choices on potential users' desire to try out the system and find that users are drawn to systems that project higher competence and warmth. These results suggest that projecting competence may help attract new users, but those users may discard the agent unless it can quickly correct with a lower competence metaphor. We close with a retrospective analysis that finds similar patterns between metaphors and user attitudes towards past conversational agents such as Xiaoice, Replika, Woebot, Mitsuku, and Tay.

CCS Concepts: • **Human-centered computing** → **Empirical studies in collaborative and social computing**; **Empirical studies in HCI**; **Human computer interaction (HCI)**.

Additional Key Words and Phrases: expectation shaping; adoption of AI systems; perception of human-AI collaboration; conceptual metaphors

ACM Reference Format:

Pranav Khadpe, Ranjay Krishna, Li Fei-Fei, Jeffrey T. Hancock, and Michael S. Bernstein. 2020. Conceptual Metaphors Impact Perceptions of Human-AI Collaboration. *Proc. ACM Hum.-Comput. Interact.* 4, CSCW2, Article 163 (October 2020), 26 pages. <https://doi.org/10.1145/3415234>

Authors' addresses: Pranav Khadpe, Stanford University, Stanford, California, USA, prkhadpe@cs.stanford.edu; Ranjay Krishna, Stanford University, Stanford, California, USA, ranjaykrishna@cs.stanford.edu; Li Fei-Fei, Stanford University, Stanford, California, USA, feifeili@cs.stanford.edu; Jeffrey T. Hancock, Stanford University, Stanford, California, USA, hancockj@stanford.edu; Michael S. Bernstein, Stanford University, Stanford, California, USA, msb@cs.stanford.edu.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.

© 2020 Association for Computing Machinery.

2573-0142/2020/10-ART163 \$15.00

<https://doi.org/10.1145/3415234>

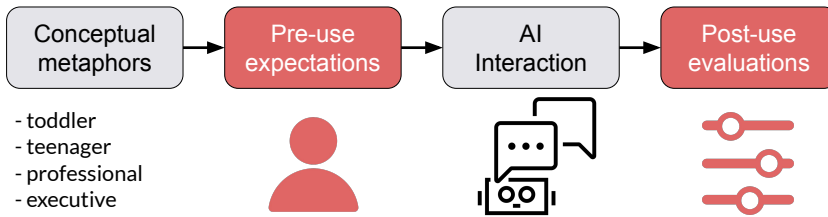


Fig. 1. We explore how the metaphors used to describe an AI agent, by influencing pre-use expectations, have a downstream impact on evaluations of those AI agents.

1 INTRODUCTION

Collaboration between people and conversational artificial intelligence (AI) agents—AI systems that communicate through natural language [30]—is now prevalent. As a result, there is increasing interest in designing these agents and studying how users interact with them [1, 14, 23, 30, 49, 66]. While the technical underpinnings of these systems continue to improve, we still lack fundamental understanding of the mechanisms that influence our experience of them. What mechanisms cause some conversational AI agents to succeed at their goals, while others are discarded? Why would Xiaoice [68] amass millions of monthly users, while the same techniques powering Tay [35] led to the agent being discontinued for eliciting anti-social troll interactions? Many AI agents have received polarized receptions despite offering very similar functionality: for example, Woebot [55] and Replika [43] continue to evoke positive user behavior, while Mitsuku [79] is often subjected to dehumanization. Even with millions of similar AI systems available online [11, 46], only a handful are not abandoned [30, 83]. The emergence of social robots and human-AI collaborations has driven home a need to understand the mechanisms that inform users’ evaluations of such systems.

In HCI, experiences of a system are typically understood as being mediated by a person’s mental model of that system [54]. Conveying an effective understanding of the system’s behavior can enable users to build mental models that increase their desire to cooperate with the system [5, 10, 37, 41]. However, a mental model explanation is insufficient to answer the present question: in the case of Xiaoice and Tay, both agents were based on the same underlying technology from Microsoft, but they resulted in very different reactions by users. Likewise, other agents such as Replika and Mitsuku elicit very different evaluations while existing even within the same cultural context. While theories of mental models and culture each help us understand how users experience conversational AI agents, we require additional theoretical scaffolding to understand the phenomenon.

An important and unexamined difference between these otherwise similar agents are the different *metaphors* that they project. Conceptual metaphors are short descriptions attached to a system that are suggestive of its functionality and intentions [15, 50]. For instance, Microsoft described Tay as an “AI that’s got no chill” [62], while it markets Xiaoice as an “empathetic ear”—two very different metaphors. Metaphors are a central mechanism in the designer’s toolkit. Unlike mental models, they offer more than just functional understandings of the system—they shape users’ expectations from the system. And while most existing expectation-shaping mechanisms depend on the functionality of the specific AI system or task [41], metaphors are agnostic to specificities of a system and can be used to shape expectations for nearly any AI system. Prior theory suggests that pre-use expectations of AI systems influence both initial behaviors [31, 40, 76] and long-term behaviors [42], even if the system itself remains unchanged while varying user expectations [58].

We propose that these metaphors are a powerful mechanism to shape expectations and mediate experiences of AI systems. If, for example, the metaphor primes people to expect an AI that is

highly competent and capable of understanding complex commands, they will evaluate the same interaction with the system differently than if users expect their AI to be less competent and only comprehend simple commands (Figure 1). Similarly, if users expect a warm, welcoming experience, they will evaluate an AI agent differently than if they expect a colder, professional experience — even if the interaction with the agent is identical in both cases.

In this paper, we test the effect of metaphors on evaluations of AI agents. We draw on the Stereotype Content Model (SCM) from psychology [16, 24], which demonstrates that the two dimensions of *warmth* and *competence* are the principal axes of human social perception. Judgements along these dimensions provoke systematic cognitive, emotional, and behavioral reactions [16]. The SCM suggests that user expectations and therefore evaluations, are mediated by judgements of warmth and competence. We crowdsource the labeling of a set of metaphors along these axes to identify a set of metaphors that appear in different quadrants of the SCM — e.g., a toddler, who is high warmth and low competence, and a shrewd executive, who is low warmth and high competence.

We perform an experiment ($N = 260$) that manipulates the metaphor associated with an AI agent and measures how it invokes expectations of competence and warmth and how those two dimensions affect ratings of usability, intention to adopt, and desire to cooperate. We draw on an established method from prior experiments [12, 34, 71, 74] to instantiate the agent itself as a remote Wizard-of-Oz who is blind to the condition and randomized across conditions for each participant. Participants are first exposed to the agent's metaphor, then converse with the agent to complete a travel planning task [4].

Our results suggest that, contrary to how designers typically describe their AI agents, low competence metaphors lead to increases in perceived usability, intention to adopt, and desire to cooperate relative to high competence metaphors. These results persist despite both the low competence and high competence agents operating at full human-level performance levels via a wizard, suggesting that no matter how competent the agent actually is, people will view it negatively if it projects a high level of competence. Participants perceive the wizards to possess lower competence than the expectations implied by high competence metaphors. These results align with Contrast Theory [67], which states that users' evaluations are defined by the difference between their experiences and expectations. Finally, we find that the warmth axis operates conversely to competence: users viewed the AI with higher warmth more positively, interacted with it longer, and were more willing to cooperate with it. This result aligns with Assimilation Theory [67]: users recolor warmth experiences in light of their initial expectation.

Previous work has sought explanations for user behavior and evaluations of AI by profiling users [5, 18] or by making the AI more interpretable [9, 44, 60]. However, these approaches fail to explain why otherwise functionally similar systems elicited vastly different user responses. Our analysis suggests that designers should carefully analyze the effects of metaphors that they associate with the AI systems they create, especially whether they are communicating expectations of high competence. In discussion, we consider implications for design by retrospectively analyzing the metaphors used to describe existing and past AI agents, such as Xiaoice, Tay, and Mitzuku, and show that our results are consistent with the adoption and user cooperation with these products. The connection between our conclusions and the outcomes experienced by Xiaoice and Tay cannot explain the whole story; however, the pattern is striking and motivates the need for exploration of mechanisms to shape expectations and elicit prosocial user behavior.

We begin by laying out related work, deriving our research question and hypotheses from prior theories. We then describe our procedure for sampling metaphors. In Study 1, we study the effects of metaphor warmth and competence. In Study 2, we sample additional metaphors along the competence axis in order to understand the effects of competence at a more fine-grained level. In

Study 3, we test the negative effects of portraying a low competence metaphor by studying the effect that warmth and competence have on participants' interest in using the system in the first place. Finally, we discuss the implications of our findings for the choice of metaphors when designers deal with the dual objective of attracting more users and ensuring a positive user experience.

2 RELATED WORK

Pre-use expectations play a critical role in users' initial usage of a system or design [32, 40, 76]. Setting positive or negative expectations colors users' evaluation of what would otherwise be identical experiences [58]. The effects of these pre-use expectations can have effects on evaluations even after weeks of interaction with a service [42].

In the case of AI systems, which are often data-driven and probabilistic, there exists no simple method of setting user expectations. Providing users with performance metrics does not establish an accurate expectation for how the system behaves [41]. In the absence of effective mental models of AI systems, users instead develop folk theories — intuitive, informal theories — as expansive guiding beliefs about the system and its goals [26, 29, 45, 64].

Prior work has shown how subjective evaluations of interface agents are strongly influenced by the face, voice, and other design aspects of the agent [53, 82], beyond just the actual capabilities of the agent. These results motivate our study of how metaphors set expectations that affect how users view and interact with conversational AI systems. Inaccurate expectations can be consequential. Previously, interviews have established that expectations from conversational agents such as Siri, Google Assistant, and Alexa are out of sync with the actual capabilities and performance of the systems [49, 83]. So, after repeatedly hitting the agent's capability limits, users retreat to using the agents only for menial, low-level tasks [49]. While these prior interview-based studies have demonstrated that a mismatch between user expectations and system operation are detrimental to user experiences [49], they haven't been able to establish causality and quantify the magnitude of this effect. This gap motivates our inquiry into understanding mechanisms that might shape these expectations and measuring the effect of expectations on user experiences and attitudes. We are guided by the following research question:

RESEARCH QUESTION: *How do metaphors impact evaluations of interactions with conversational AI systems?*

2.1 Metaphors shape expectations

Conceptual metaphors are one of the most common and powerful means that a designer has to influence user expectations. We refer to a conceptual metaphor (or user interface metaphor, or just metaphor) as the understanding and expression of complex or abstract ideas using simple terms [45]. Metaphors are attached to all types of AI systems, both by designers to communicate aspects of the system and by users to express their understanding of the system. For instance, Google describes its search algorithm as a "robotic nose" [26] and YouTube users think of the recommendation algorithm as a "drug dealer" [81]. Starting with the desktop metaphor for personal computing in the Xerox Star [39], conceptual metaphors proliferated through the design of user interfaces — trash cans for deleted files, notepads for freetext notes, analog shutter clicking sounds for mobile phone cameras, and more.

Some AI agents utilize metaphors based in personas or human roles, for example an administrative assistant, a teenager, a friend, or a psychotherapist, and some are metaphors grounded in other contexts, for example a Jetsons-style humanoid servant robot. Such metaphors are meant to help human-AI collaboration in complex domains by aiding users' ability to understand and predict the agent's behavior [5, 18]. Metaphors include system descriptions outside of those rooted in

human roles as well: Google describing its search algorithm as a “robotic nose” [26] and Microsoft’s Zo marketed as a bot that “Will make you LOL”. The notion of “metaphors” extends beyond conversational AI to non-anthropomorphic systems that “personas” or “roles” may be ill-equipped to describe. Metaphors are effective: they influence a person’s folk theories of an AI system even before they use it [19]. Prior work has developed methods to extract conceptual metaphors [45, 64] for how people understand AI systems and aggregate them into underlying folk theories [26].

Metaphors impact expectations, sometimes implicitly by activating different norms, biases, and expectations. For example, social robots that are racialized as Black or Asian are more likely to be subject to antisocial behaviour such as aggression and objectification [70]. Similarly, female-gendered robots can elicit higher levels of dehumanisation than male-gendered bots. Antisocial behavior leads to verbal disinhibition toward AI systems [69], and in some extreme cases, to physical abuse and even dismemberment [8, 61]. Female voice agents are viewed as friendlier but less intelligent [53]. Users also have a higher tendency to disclose information to female gendered agents [53]. Race and gender of pedagogical agents affect learning outcomes—agents racialized as Black or female-gendered lead to improved attention and learning [6]. Beyond race and gender, agents portrayed as less intelligent, taking on roles such as “motivator” or “mentor”, promote more self-efficacy than agents projected as “experts” [6]. Young, urban users respond positively to bots that can add value to their life by suggesting recommendations, while in the role of a “friend” [71].

However, designers typically aim to use metaphors to affect expectations in more explicit, controlled, and pro-social ways. Most obviously, a metaphor communicates expectations of what can and cannot be done with an AI agent [39]. Just as we expect an administrative assistant to know our calendar but not to know the recipe for the best stoat sandwiches, an AI agent that communicates a metaphor as an “administrative assistant” projects the same skills and boundaries. In a similar vein, describing an agent as a “toddler” suggests that the agent can interact in natural language and understand some, but not all, of our communication.

While other expectation shaping mechanisms for AI agents such as tutorials and instructions have been studied [41], the effect of metaphors on user expectations and evaluations have not. Our work also bridges to research suggesting that people already form metaphor-based theories of socio-technical systems [26] and suggests design implications for how designers should choose their metaphors.

2.2 Competing predictions: assimilation vs. contrast

As people view AI agents as social agents [59], the metaphor—and thus the nature of that agent—is likely to influence their experience. However, the literature presents two competing theories for how changes to the metaphor — and thus to expectations — will impact user evaluation of an AI system. Assimilation theory [67] states that people adapt their perceptions to match their expectations, and thus adjust their evaluations to be positively correlated with their initial expectations. (As Dumbledore points out to Snape in *Harry Potter and the Deathly Hallows*, “You see what you expect to see, Severus.”) Assimilation theory argues that users don’t perceive a difference between their pre-use expectations and actual experiences. Prior work supports that, for interactive systems, users’ expectations do influence evaluations [31, 73]. For example, users rate an interactive system higher when they are shown a positive review of that system before using it, and rate the system lower if they are shown a negative review before using it [58]. Likewise, humor and other human-like characteristics that create high social intelligence expectations can be crucial in producing positive evaluations [36, 48].

Assimilation theory would predict that a metaphor signaling high competence will set positive expectations and subsequently lead to positive evaluation:

HYPOTHESIS 1 (H1). *Positive metaphors (e.g., high competence, high warmth) will lead to higher average intention to adopt and desire to cooperate with an AI agent than if it had no metaphor or negative metaphors.*

Contrast theory [67], on the other hand, attributes user evaluations to the difference they perceive between expectations and actual experience. Contrast theory argues that we are attuned not to absolute experiences, but to differences between our expectations and our experiences. For example, exceeding expectations results in high satisfaction, whereas falling short of expectations results in lower satisfaction. This suggests that it is beneficial to set users' initial expectations to be low (with practitioners reasoning in the manner of George Weasley, in *Harry Potter and the Order of the Phoenix*, "E' is for 'Exceeds Expectations' and I've always thought Fred and I should've got 'E' in everything, because we exceeded expectations just by turning up for the exams.") Users of conversational AI agents such as Alexa stumble onto humorous easter egg commands that raise their expectations of what the system can do, but then report disappointment in the contrast to discovering the system's actual limits [49]. Likewise, ratings of interactive games are driven in part by contrasting players experiences against their expectations of the game [51].

Contrast theory predicts that positive metaphors will backfire because AI agents inevitably make mistakes and have limits:

HYPOTHESIS 2 (H2). *Positive metaphors (e.g., high competence, high warmth) will lead to lower average intention to adopt and desire to cooperate with an AI agent than if it had no metaphor or negative metaphors.*

3 METHODS

Our research aim is to study the effect of metaphors on experiences with AI agents. So, we seek an experimental setup where participants accomplish a task in collaboration with an AI system, while avoiding effects introduced by idiosyncrasies of any particular AI system. We situate our method in goal-oriented conversational agents (or task-focused bots) as these systems represent a broad class of agents in research and product [17, 35, 55, 68, 78, 79].

3.1 Collaborative AI task

Goal-oriented AI systems, such as those for booking flights, hotel rooms, or navigating customer service requests, have become pervasive on social media platforms including Kik, Slack, and Facebook Messenger, with as many as one million flooding the Web between 2015 and 2019 [30]. Surveys revealed that as of 2018, as many as 60% of surveyed millenials had used a chatbot [2] and 15% of surveyed internet users had used customer-service chatbots [20]. Their prevalence means that interaction with such an agent is an ecologically valid task, and that many users online are familiar with how to interact with them. We draw on a common set of transactional tasks such as appointment booking, scheduling, and purchasing, which require people to engage with the agent in task-focused dialogue to acquire information or complete their task [27]. Inspired by the popular Maluuba Frames [4] data collection task templates, used to evaluate conversational agents in the natural language processing community, we utilize a travel planning task. More concretely, the task is a vacation planning endeavor where users must pick a vacation package that meets a set of experimenter-specified requirements through a search-compare-decide process. Specifically, every participant is presented with the following prompt:

You are considering going to New York, Berlin or Paris from Montreal. You want to travel sometime between August 23rd and September 1st. You are traveling alone. Ask for information about options available in all cities. Compare the alternatives and make your decision.

Participants were further instructed to determine what they could get for their money and to take into consideration factors they would consider while actually planning a vacation, including wifi, breakfast options, and a spa. The task is structured to involve three sub-goals: finalize a hotel package, an outgoing flight and an incoming flight back to Montreal.

3.2 Wizard-of-Oz conversational agent

We sought a conversational AI agent whose actual performance was strong enough for our result to generalize as the underlying AI models improve. So, following a pattern in prior work [12, 34, 71, 74], we adopt a Wizard-of-Oz study paradigm.

We hire and train customer-support professionals from the Upwork platform to act as wizards in our experiment and pay them their posted hourly rate of \$10 – \$20 per hour. The wizards play the role of the conversational AI agent in the text chat. We filtered workers who had at least a 90% job success rating from past work, and had already earned \$20k USD through the Upwork platform. We also filtered for workers with English proficiency by asking them to submit a cover letter detailing their past work experience and manually checked for spelling or grammatical errors. We hired 5 wizards in all. To eliminate wizard-specific confounds across our different conditions, wizards were blind to the treatment condition of the participants in the study, and randomized to a new condition for each new participant that they interacted with. Randomizing this source of variation produces an unbiased estimate of the effect of each condition.

Wizards were trained on how to provide responses by engaging in practice tasks with the authors as participants. They were instructed to send emotionally neutral responses simply addressing the participant's query. Wizards were required to, at every turn, search for hotels that met the requirements specified by the participant. In case multiple hotels met the set of requirements, they were asked to present all the options to the participant, ordered according to how they appear in the database. Once users achieved their three goals (booking two flights and a hotel), the wizards informed the users to proceed to the next step of the study. In order to minimize inter-wizard differences, we provided feedback to wizards from trial conversations to ensure that there were no drastic differences across the wizards' performance. In many common exchanges, wizards were also provided with template responses. To highlight the similarity of the wizards' responses, sample responses of three wizards to similar input queries are provided in the Supplementary Material.

Task-focused agents are knowledgeable within a narrow task focus [30] and are often unable to answer questions that require external knowledge. In order to retain phenomena associated with access to finite knowledge, we provide our wizards with a database of hotels and flights. We construct this database by hand to mimic what one would find on a standard travel booking platform. We have provided details on the construction of the database and examples of hotels and flights from the database in the supplementary material. Consistent with instructions provided to wizards in the creation of Frames [4] corpus, if the wizard is asked about knowledge that is outside of their available database, wizards were trained to respond that they do not have that information.

We constructed a conversational chat platform via using Chatplat (<https://www.chatplat.com>). We embedded this chat widget into our web based survey.

3.3 Sampling metaphors

We place participants into treatment groups each defined by the metaphor used to describe their AI collaborator. Instead of randomly sampling metaphors, we draw on the Stereotype Content Model (SCM) [16, 24], an influential psychological theory that articulates two major axes in social perception: warmth and competence. These two dimensions have proven to far outweigh others and repeatedly come up as prime factors in literature [3, 38, 77]. Judgements on warmth and competence

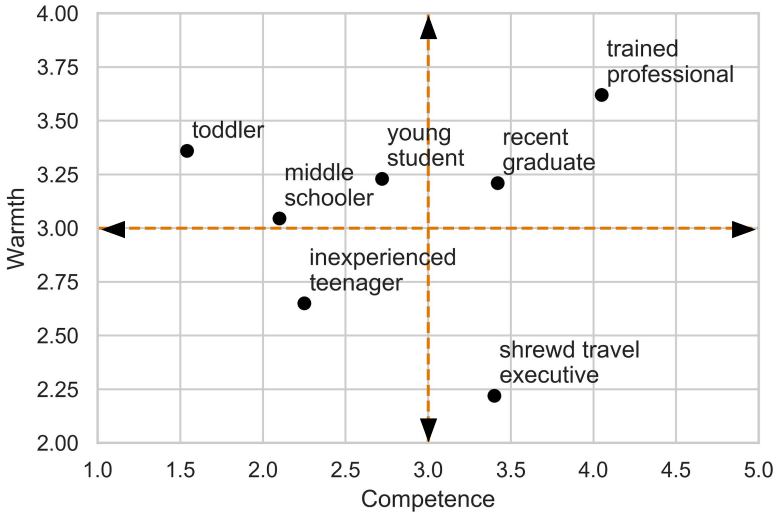


Fig. 2. Average warmth and competence measured for the conceptual metaphors sampled for our studies. Both the axes ranged from 1 to 5.

	Toddler	Middle Schooler	Inexperienced Teenager	Young Student	Shrewd Travel Executive	Recent Graduate	Trained Professional
Competence	1.54 ± 1.04	2.1 ± 0.83	2.25 ± 0.85	2.72 ± 0.58	3.4 ± 0.97	3.42 ± 0.72	4.05 ± 0.58
Warmth	3.36 ± 0.81	3.05 ± 0.69	2.65 ± 0.73	3.23 ± 0.56	2.22 ± 1.02	3.21 ± 0.67	3.62 ± 0.59

Table 1. Warmth and competence values (average ± standard deviations) for the metaphors we use across the studies. High competence and warmth values are in bold.

are made within 100 milliseconds [72] and a change in these traits alone can wholly change impressions [84]. The SCM proposes a quadrant structure and cognitive notions of warmth and competence- better understood as discrete- are characterized as being low or high [16]. Warmth is characterised by notions such as good-naturedness and sincerity, while competence is characterised by notions of intelligence, responsibility, and skillfulness. For example, a “shrewd travel executive” can be described as high competence and low warmth. We sample metaphors such that they have either high or low values of warmth and competence.

In our first study, we use four metaphors, one in each quadrant, to study the impact of competence and warmth. We pre-tested a set of metaphors for a conversational agent, measuring the perceived competence and warmth of conversational AI agents described with these metaphors using a 5 point Likert scale. We captured 50 ratings for each metaphor from workers on Amazon Mechanical Turk – a mutually exclusive set of workers from those who will later be involved in the experiment. Based on the results (see Figure 2), we chose “trained professional travel assistant” (high competence, high warmth), “shrewd travel executive” (high competence, low warmth), “toddler” (low competence, high warmth), and “inexperienced teenager” (low competence, low warmth). We selected metaphors that were otherwise agendered, with similar socio-cultural connotations across the world, and representative of actual metaphors that could be associated with a travel assistant bot. These four

metaphors form our four treatment groups in Study 1; their mean and standard deviation values of competence and warmth are reported in Table 1.

In Study 2, we follow the same procedure to characterize several additional metaphors: “middle schooler”, “young student”, and “recent graduate”. These metaphors offer intermediate levels of competence, with “toddler” less competent than “middle schooler”, “middle schooler” less competent than “young student”, and “young student” less competent than “trained professional”. “Young student” is associated with higher competence levels than “middle schooler”, suggesting that people’s impression of a “young student” is a high schooler or college student, somewhere between a “middle schooler” and a “recent graduate”. In Study 3, we revisit the metaphors we analyzed in Study 1 to understand the effects of metaphors on potential users’ likelihood of trying out the system and their intentions of co-operating with it prior to using it.

4 STUDY 1: METAPHORS DRIVE EVALUATIONS

In our first study, we examine the effects of metaphors attached to the conversational system on user pre-use expectations and post-use evaluations. Specifically, we examine participants’ perceived pre- and post-use usability and warmth of the AI system. Additionally, we measure their post-use intention to adopt and their desire to cooperate with such a system given their treatment metaphor condition. Finally, we analyze the chat logs to explore if there are behavioral differences between the participants in different conditions.

4.1 Procedure

We perform a between-subjects experiment where participants in each treatment condition are primed with a metaphor to associate with the system. As described in the previous section, metaphors were chosen to vary as low/high warmth \times low/high competence, resulting in four treatment conditions. In addition, we included a control condition where participants were not primed with a metaphor, resulting in five total conditions.

After consenting to the study, participants were introduced to one of the study conditions, i.e. they were shown one of the four metaphors, or a control condition of no metaphor:

The bot you are about to interact with is modeled after a “shrewd travel executive”.

With the study condition revealed, participants were asked questions about their pre-use expectations of the AI system’s competence and warmth. Next, participants were shown the goal-oriented task description and allowed to interact with the wizard posing as conversational agent via the chat widget until they completed their task.

After finalizing their travel plans, participants were asked to evaluate their experience with the AI system and answer the manipulation check question. Finally, participants were debriefed, informed of the actual purpose of the study, and made aware that they were talking to a human and not an AI system. A high-level workflow is depicted in Figure 1.

4.2 Measures

User evaluation measures. To test contrast theory, it is important to measure a user’s evaluation of the experience without drawing explicit attention to the contrast between their expectations and their experience—since this makes the contrast salient [42]. So, we independently measure pre-use expectations and post-use evaluations without explicitly asking if expectations were met or violated. To gauge participants’ pre-use expectations and post-use perceptions of the systems competence and warmth, we ask the participants to report how strongly, on a 5 point Likert scale (where 1 = strongly disagree and 5 = strongly agree), they agree with the following statements, both

before and after they interacted with the AI system. Questions asked before use simply replaced the past tense of the verb with the future tense; the question ordering was randomized.

- **Usability:** Since our notion of a system's competence is akin to the notion of usability in previous studies, we adapt questions from previous surveys that examine usability [42]. These questions are: 1) *"Using the AI system was (will be) a frustrating experience."* 2) *"The AI system was (will be) easy to use."* 3) *"I spent (will spend) too much time correcting things with this AI system."* 4) *"The AI system met (will meet) my requirements."* Responses from before using the system are combined to form a pre-use usability index ($\alpha = 0.91$) while responses from after the conversation are combined to form a post-use usability index ($\alpha = 0.87$).
- **Warmth:** To measure the warmth of the AI system, we draw on different warmth levels articulated in the stereotype content model [25]: 1) *"This AI system was (will be) good-natured."* 2) *"This AI system was (will be) warm."* Responses from before using the system are combined to form a pre-use warmth index ($\alpha = 0.94$). Similarly, responses from after the conversation are combined to form a post-use warmth index ($\alpha = 0.91$).
- **Intention to Adopt and Desire to Cooperate:** We borrow from prior work [42] that captures user evaluations through their intentions to adopt the system. Since we increasingly have situations in which humans work alongside AI systems where these systems augment human efforts, it also becomes necessary to understand users' behavioural tendencies towards these systems. So, we draw on prior work in HRI [52] and capture users' behavioural tendencies through their desire to help and cooperate with the AI system. After their interaction with the system, participants are probed for their intentions to adopt as well as their desire to cooperate with the system. To probe for the participants' intentions to adopt, we asked them the following two questions on 5 point Likert scales: 1) *"Based on your experience, how willing are you to continue using the service?"*, 2) *"How likely is it that you will be using the service in the future?"*. Like previous work [42], these two questions are combined to form an intention to adopt index ($\alpha = 0.98$). To understand the participants' desire to cooperate, we use questions about behavioral tendencies towards stereotyped groups adapted to the context of social robots [52]. Users are asked on a 5 point Likert scale: *"How likely would you be to cooperate with this AI?"* and *"How likely would you be to help this AI?"*. Like previous work [52], these two questions are combined to form a cooperation index ($\alpha = 0.79$).

Conversational behavior measures. To investigate if participant behavior changes across conditions we include measures to analyze differences in the conversational behavior of users.

- **Language measures:** To measure differences in the chatlogs by the participant and by the wizards across the various conditions, we utilize the popular linguistic dictionary Linguistic Inquiry and Word Count, known as LIWC [57]. LIWC uses 82 dimensions to determine if a text uses positive or negative emotions, self-references, and causal words, to help assess physical and mental health, intentions and expectations of the writers. We categorize all the words used by the participants and the wizards into LIWC categories and create normalized frequency histograms of these categories. We compare the words used by the participants across all the conditions to see if there significant differences in the types of LIWC categories used. Similarly, we compare the wizards' words across all conditions. Finally, we combine the words used by the wizard and participant together and also check to see if there were differences between conversation across the different conditions.
- **Conversation measures:** We also investigate differences across conditions, at the level of individual messages and whole conversations in terms of *number of words used* and *duration of interaction*.

4.3 Participants

For all the studies in this paper, we hired participants to interact with our Wizard-of-Oz AI system from Amazon Mechanical Turk (AMT). Participants were all US citizens aged 18+. Each participant was allowed to take part in the experiment only once. Participants were compensated \$4 for a survey lasting an average of 15 ± 5 minutes, for a rate of roughly \$15/hr in accordance with fair work standards on Mechanical Turk [75]. Participants' data was discarded if they failed to follow instructions, left the conversation midway or did not follow the task specifications. 50.7% of our participants were female and the mean age of participants was 41 ± 10.3 .

In this specific study, for a small expected effect size of $f = 0.05$, a power analysis with a significance level of 0.05, powered at 80%, a power analysis indicated that we require 25 participants per condition, or 125 total participants. Thirteen participants' responses were discarded because the raters had coded their WoZ manipulation check as expressing suspicion that the agent might be human. After these exclusions, we had a sample size of 140 participants, which met the requirements from our power analysis.

4.4 Wizard-of-Oz manipulation check

To ensure that our study was not compromised by participants who identified that they were speaking to a wizard instead of an AI, we included a manipulation check at the end of the survey. The manipulation check gauges whether the participant was suspicious of the AI without explicitly drawing their attention to the fact that this might be the case. So, drawing on prior Wizard-of-Oz studies [33], we asked the participants how they thought the system worked from a technical standpoint.

The responses were sent to two coders who inspected each response individually and marked all the responses that suspected a person was pretending to be an AI system. Participants who expressed suspicion that the system might be human were excluded from further analysis. Some participants were very confident they knew how such our conversational AI could be built: *I am a programmer so i understand the bot has a vocabulary of words it attempts to parse through, then it takes what it finds from the user and checks against a database to output information it thinks is relevant*. Others talked about how *Most chat bots go through "training" beforehand to be able to parse commonly asked questions and phrasing* or how it must be using a *database full of responses*.

Out of all our participants, both coders identified the same 13 participants ($\kappa = 1$) who failed the manipulation check by calling out the agent as a human, resulting in a suspicion level of 10.4%. In most of these cases, it was triggered by a wizard making a typo or taking too long to respond. One suspicious participant exclaimed, *"I'm like 90% sure it's not a bot, but if it were a bot, machine learning, though I don't know exactly what THAT means"*. These 13 participants were excluded from our analysis.

4.5 Results: Metaphors shape pre-use expectations

We compare the impact of setting expectations by varying the competence and warmth of the metaphors used. We perform our analysis using a pair of two-way analyses of variance (ANOVAs), where competence and warmth are two categorical independent variables, and pre-use usability and warmth are the dependent variables. We compared the impact of the conceptual metaphors used to describe our system compared to the control condition to measure if they have an impact on the participants' default expectations of conversational AI systems. So, the independent variables are categorized into high, low or control categories.

Pre-use usability is affected by the metaphor's competence. For pre-use usability (Figure 3 (a)), we find that competence has a large main effect ($F(1, 108) = 8.71, p < 0.001, \eta^2 = 0.12$). By

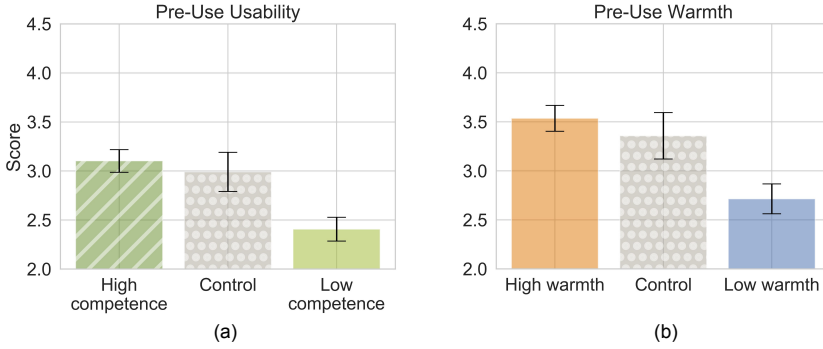


Fig. 3. (a) Metaphors that signal high competence lead to higher pre-use usability scores. (b) Similarly metaphors that signal high warmth lead to higher pre-use warmth scores. We also notice from both (a) and (b) that participants are naturally predisposed to have high expectations of usability and warmth from conversational systems; however, priming them with metaphors reduces the variance of their expectations as opposed to when their expectations are uninformed.

default, participants have high expectations of competence. A post-hoc Tukey revealed that pre-use usability was significantly lower for the low competence condition (2.41 ± 0.89) than high competence (3.10 ± 0.86 , $p < 0.001$) or control (2.99 ± 1.03 , $p = 0.019$) conditions. We found no main effects for warmth.

Pre-use warmth is likewise affected by the metaphor's warmth. For pre-use warmth (see Figure 3 (b)), we find that warmth has a large main effect ($F(1, 108) = 8.20$, $p < 0.001$, $\eta^2 = 0.11$). By default, participants have high expectations of warmth. A post-hoc Tukey revealed that pre-use warmth is significantly low for the low warmth condition (2.71 ± 1.12) than both the high warmth (3.53 ± 1.05 , $p < 0.001$) and control (3.35 ± 1.23 , $p = 0.035$) conditions. We found no main effects of competence.

Together, these tests imply that participants, by default, expect conversational AI to possess high competence and high warmth. However, the change in expectation caused by low competence and low warmth implies that these conceptual metaphors do affect participants' expectations of how the AI system will perform and behave. We visualize the means and standard errors for these conditions in Figure 3 (a, b).

4.6 Results: Metaphors impact user evaluations and user attitudes

We compare the impact of varying the competence and warmth of the metaphors on participants' post-use evaluations of the AI system's usability and warmth. We perform our analysis using a pair of two-way ANOVAs where competence and warmth are categorical independent variables, and post-use usability and post-use warmth are the two dependent variables.

Participants perceive agents with low competence to be more usable after interaction. For post-use usability (Figure 4 (a)), competence has a main effect ($F(1, 108) = 6.49$, $p < 0.05$, $\eta^2 = 0.06$). Competence has a smaller effect on post-use usability than on pre-use usability, implying that the actual interaction of the participant with the system affects their final evaluations. In the low competence condition, post-use usability was rated at (4.02 ± 0.72) and in the high competence condition, it was rated (3.58 ± 1.07). These results suggest that users perceive a difference between their experience and their expectations in terms of competence of the agent. The means for post-use usability are also higher than for pre-use usability for both high ($t(55) = -3.48$, $p < 0.001$) and

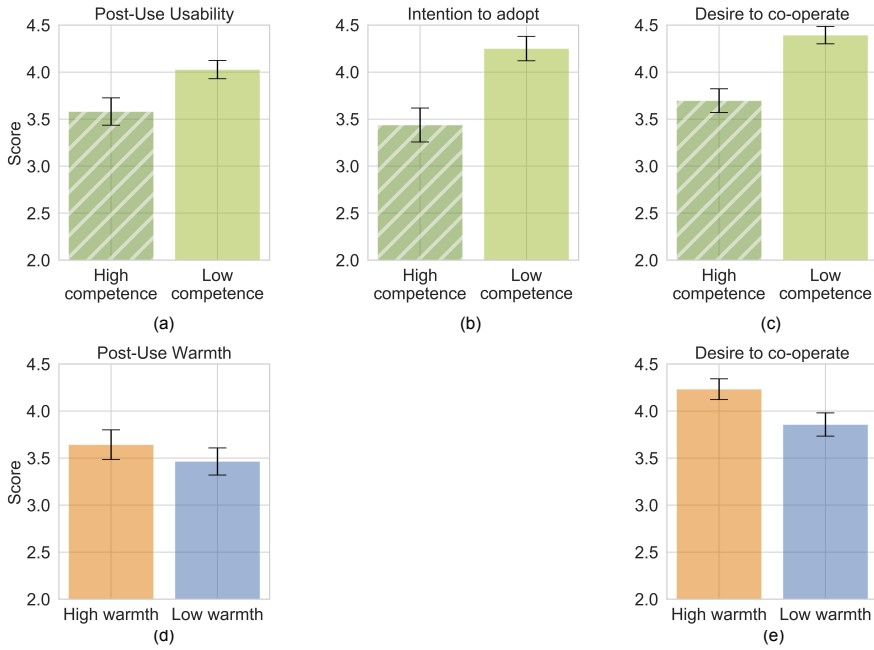


Fig. 4. The low competence metaphor condition features the highest post-use usability, intention to adopt, and desire to cooperate. This result suggests that metaphors that undersell the AI agent's competence are most likely to succeed.

low ($t(55) = -10.79, p < 0.001$) competence conditions. We found no main effects of warmth or interaction effects between competence and warmth.

We observe post-use warmth ratings to be higher in the high warmth condition than the low warmth condition though the difference is not significant. For post-use warmth (Figure 4 (d)), we find no main effects of competence or warmth and no interaction effects. In the high warmth condition, warmth was rated at (3.64 ± 1.17) and in the low warmth condition, it was rated (3.50 ± 1.12). There is no significant difference between the means of pre-use and post-use warmth for high warmth ($t(55) = -0.67, p = 0.501$), but it is significantly different in the low warmth ($t(55) = -4.56, p < 0.001$).

Using the composites described in the study design, we measure the effect of conceptual metaphors on the participants' intention to adopt and desire to cooperate after interacting with the system.

Low competence metaphors increase participants' likelihood of adopting the AI agent. For their intention to adopt (Figure 4 (b)), competence has a main effect ($F(1, 108) = 13.31, p < 0.001, \eta^2 = 0.12$). In the high competence condition, the intention to adopt was rated at (3.42 ± 1.33) and in the low competence condition, it was rated (4.25 ± 0.96). These results support Hypothesis 2 as we see support for contrast theory: participants are more likely to adopt an agent that they originally expected to have low competence but outperforms that expectation. They are less forgiving of mistakes made by AI systems they expect to have high competence. We found no main effects of warmth or interaction effects between competence and warmth.

Participants prefer to cooperate with agents that have high warmth and low competence. For their desire to cooperate with the AI system, we found that both competence ($F(1, 108) = 20.58, p <$

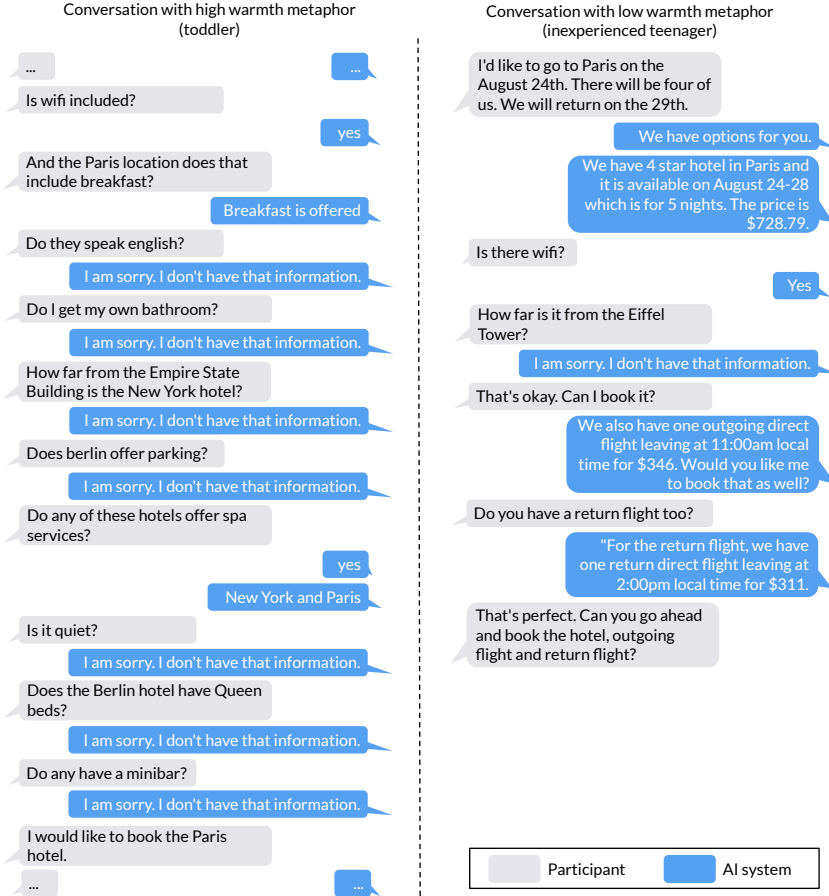


Fig. 5. Segments of two example conversations between a participant with our conversational AI system. In both cases, the participant expects the AI system to have low competence. While the left conversation is in the high warmth metaphor condition, the right conversation is in the low warmth metaphor condition. Participants in the high warmth condition ask more questions and explore the space of possible interactions by asking the agent details about checked luggage and hotel amenities. Wizards, acting as conversation agents, are given a fixed knowledge set, mimicking how today's systems are designed, and reply with apologies when asked about details outside of their knowledge.

0.001, $\eta^2 = 0.19$) and warmth ($F(1, 108) = 5.96, p < 0.01, \eta^2 = 0.05$) had main effects but no interaction effect (see Figure 4 (c, e)). The means increase from high (3.69 ± 0.93) to low competence of (4.39 ± 0.67). Similarly, the means decrease from high (4.23 ± 0.81) to low warmth (3.85 ± 0.92). These results provide mixed support to both Hypothesis 1 and Hypothesis 2 as we see support for assimilation theory along the warmth dimension and contrast theory along the competence dimension. If participants are told that the AI system is high warmth, they are more likely to cooperate with it. But if the AI system is described as high competence, they are less likely to cooperate.

4.7 Results: Expectations change, but behavior doesn't

We analyze the chat logs with LIWC features, following a standard LIWC analysis protocol of building a frequency count of how often words belonging to a specific LIWC category were used. We contrasted these counts across the various conditions and observed no significant differences ($p > .05$) in language level phenomenon in the chatlogs across the conditions. This result implies that the post-use evaluations are driven primarily by the expectations set by the metaphors, not by the actual content of the conversation. In other words, evaluations differed between conditions, but the actual conversations themselves did not. The wizard was blinded to the condition, so any differences would have needed to be prompted by the participant. However, we acknowledge that there might be language shifts that LIWC categories cannot capture.

Participants use more words and spend more time speaking to agents with high warmth. We find a significant main effect of warmth on the number of words used per conversation ($F(1, 106) = 5.35, p < 0.05, \eta^2 = 0.05$). The number of words increase from (82 ± 37) in low to (101 ± 45) in high warmth. We also find that participants in the high warmth condition typically spend an average of 4 ± 1.5 minutes longer while interacting with the AI system. On a qualitative inspection, we find that participants tend to ask more questions and spent more time exploring the AI system's capabilities. Consider, for example, the conversation shown in Figure 5, where the participant expects the bot to have low competence and high warmth. The participant asks numerous questions to test the system's capabilities and even though it fails, they later express a high intention to adopt and cooperate with the system.

4.8 Summary

Our results support contrast theory (Hypothesis 2) for the competence axis. Users are more tolerant of gaps in knowledge of systems with low competence but are less forgiving of high competence systems making mistakes. The intention to adopt and desire to cooperate decreases as the competence of the AI system metaphor increases. For the warmth axis, our results provide some support for assimilation theory (Hypothesis 1): users are more likely to co-operate and interact longer with agents portraying high warmth, but we do not observe significant impact of warmth on users' intention to adopt.

5 STUDY 2: THE COMPETENCE-ADOPTION CURVE

Study 1 established that setting low expectations and violating them positively increased the likelihood of users adopting the system. In Study 2, we zoom in and try to understand how user evaluations change as the magnitude of that gap changes. We sample additional metaphors and use the same experiment procedure as before to characterize how users' intentions to adopt the system vary as gap between users expectations and their experience changes. For this purpose, we rely on the same measure of Intention to Adopt as Study 1.

5.1 Procedure

To precisely traverse the range of perceived competence, we sampled 3 additional metaphors — “middle schooler”, “young student”, and “recent graduate”. Our pre-experiment survey revealed that these metaphors had perceived competence levels between the “toddler” and the “trained professional”. Together, these five metaphors formed five treatment conditions. As Figure 2 demonstrates, all five metaphors lie in the high warmth half of the space, minimizing any interfering effects of variations in warmth. Participants in these five conditions were primed with the respective metaphor.

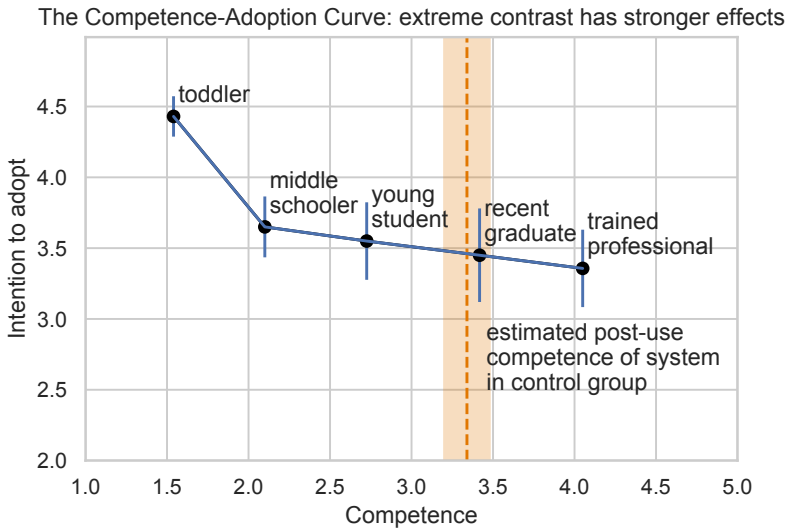


Fig. 6. A larger positive violation of expectation increases adoption intentions. The intention to adopt decreases monotonically with an increase in expected competence of the system. The red vertical line shows the average score users in the control condition assigned the system and the yellow shaded region around the vertical line depicts the standard deviation.

To understand users' unprimed evaluations of the system, we asked a sixth control group to participate without a metaphor (similar to the control condition in Study 1). Afterwards, we asked them to pick from the list of five metaphors and identify which one they felt described the system most accurately after use.

5.2 Measures

Similar to Study 1, we measure participants' intention to adopt and desire to co-operate across all five metaphors.

5.3 Participants

Similar to the protocol in Study 1, participants were recruited on AMT. We recruited 20 participants for each condition, for a total of 120 participants. The duration of the study was similar to Study 1 and participants were compensated at the same rate. The average age of participants was 39 ± 11.06 ; 48% identified as female.

5.4 Wizard-of-oz manipulation check

The two coders were consistent and identified the same 5 participants ($\kappa = 1$) as being suspicious, implying a low suspicion level of 4.1%. These five participants were removed from analysis.

5.5 Extreme violations of expectations have stronger effects

The five metaphors we sampled are shown in Figure 6, where the x-axis depicts workers' perceived competence of a system with that metaphor and the y-axis depicts a different set of users' intention to adopt. The red vertical line shows the average score users in the control condition assigned the

system and the yellow shaded region around the vertical line depicts the standard deviation. The unprimed system was viewed roughly as competently as a recent graduate.

Over-performing low competence leads to higher adoption than over-performing medium competence, and projecting any more competence than the toddler metaphor incurs an immediate cost (Figure 6). These results paint a fuller picture of contrast theory at play, as the intention to adopt decreases monotonically as the expected competence of the system increases. Consistent with prior literature, the effect is greater as the contrast is greater [7, 28]. However, the effect is nonlinear, with only the lowest competence metaphor receiving a substantial benefit.

The “toddler” metaphor sees the highest (beneficial) violation as it is furthest away from the vertical line and sees the greatest intention to adopt and desire to cooperate. There was a statistically significant difference between groups as determined by one-way ANOVA for intention to adopt ($F(4, 114) = 3.701, p = .007$) and for desire to cooperate ($F(4, 114) = 6.841, p < .001$). A Tukey post-hoc test revealed that intention to adopt was statistically significantly higher for “toddler” (4.42 ± 0.76) than “young student” ($3.55 \pm 1.25, p = .047$), “recent graduate” ($3.45 \pm 1.5, p = .021$), and “trained professional” ($3.35 \pm 1.43, p = .007$). There was no statistically significant difference between other metaphors. Similarly, a Tukey post-hoc test revealed that desire to cooperate was statistically significantly higher for “toddler” (4.50 ± 0.50) than “middle schooler” ($3.52 \pm 1.08, p = .001$), “young student” ($3.82 \pm 0.86, p = .006$), “recent graduate” ($3.95 \pm 0.76, p = .003$), and “trained professional” ($3.87 \pm 0.93, p = .013$). There was no statistically significant difference between other metaphors.

5.6 Summary

Our results further support contrast theory (Hypothesis 2) for the competence axis. Users are more likely to adopt a lower competence agent than one with high competence, even though all conditions were exposed to human-level performance. We additionally see an asymmetry — users are even more likely to adopt an agent that exceeds extremely low expectations than one that exceeds slightly higher (but still low) expectations. And as the agent begins to under-perform expectations, intentions to adopt decrease further.

6 STUDY 3: THE COST OF LOW-COMPETENCE METAPHORS

From our results so far, it might appear that that designers should pick metaphors that project lower competence and high warmth regardless of experience, as these conditions are most conducive for cooperative and patient user interactions. However, such a conclusion might be myopic. Metaphors attached to a system also have the ability to attract or drive people away.

To test the effect of metaphor on pre-use intention to adopt, and pre-use desire to cooperate with an AI system, we ran a third study. In this study, we present participants with AI systems, described using conceptual metaphors. We ask participants to identify which systems they are more likely to try out and potentially adopt, prior to using the system.

6.1 Procedure

We perform a between-subjects experiment. Each participant was introduced to an AI agent described using one of the metaphors in Study 1. Unlike the previous experiments, participants do not actually interact with an AI system (or wizard). Instead, they are asked to rate their likelihood of trying out a new AI system service described by each metaphor.

6.2 Measures

To probe for the participants’ intentions to try the system, we asked them the following two questions on 5 point Likert scales: *How likely are you to try out this AI system?*, and *Do you envision*

yourself engaging in long-term use of such a AI system? These two questions are combined to form a trial index ($\alpha = 0.94$).

To understand the participants' pre-use desire to cooperate with the system, they are asked on a 5 point Likert scale: *How likely are you to cooperate with such an AI system?*, and *How likely are you to tolerate errors made by this AI system?* These two questions are combined to form a pre-use desire to cooperate index ($\alpha = 0.78$).

6.3 Participants

Similar to the previous studies, we recruited participants from AMT. 80 new participants participated in this survey: 20 participants exposed to a metaphor from each quadrant. We ensured that none of the participants in this study participate in any of our other studies.

6.4 More interest in trying out high competence and high warmth AI systems

Participants were more interested to try out AI systems that were described by high competence and high warmth. A two-way ANOVA revealed that competence ($F(1, 100) = 42.96, p < 0.001, \eta^2 = 0.34$) and warmth ($F(1, 100) = 5.63, p = 0.020, \eta^2 = 0.04$) both had significant impact on their intention to try out the AI system. The average trial index response was (2.95 ± 1.18) for high competence, and (1.5 ± 0.83) for low competence. Following a similar pattern, the average trial index response was (2.48 ± 1.28) for high warmth as opposed to (1.96 ± 1.16) for low warmth. The ANOVA also showed an interaction effect between competence and warmth ($F(1, 100) = 1.54, p = 0.021, \eta^2 = 0.012$). In this case, the combination of high competence and high warmth produced a substantial benefit 3.35 ± 1.06 compared to the effects of warmth and competence individually: low competence and high warmth 1.63 ± 0.84 , high competence and low warmth 2.55 ± 1.77 , and both low competence and low warmth 1.38 ± 0.82 .

Participants were more likely to cooperate positively towards AI systems that were described by high competence and high warmth. A two-way ANOVA revealed that competence ($F(1, 100) = 37.36, p < 0.001, \eta^2 = 0.28$) and warmth ($F(1, 100) = 5.63, p = 0.01, \eta^2 = 0.08$) both had significant impact on the trial index. The average pre-use desire to cooperate index response was (2.80 ± 1.11) for the high competence as opposed to (1.58 ± 0.84) for low competence. Similarly, the average pre-use desire to cooperate index response was (2.53 ± 1.18) for the high warmth as opposed to (1.85 ± 1.04) for low warmth. The ANOVA also showed an interaction effect between competence and warmth ($F(1, 100) = 6.86, p = 0.001, \eta^2 = 0.052$). People expected to behave more positively with high competence and high warmth AI systems 3.40 ± 0.75 over the low competence and high warmth 1.65 ± 0.82 , high competence and low warmth 2.20 ± 1.09 , and both low competence and low warmth 1.50 ± 0.87 bots.

6.5 Summary

While our previous studies demonstrated the detrimental effects of presenting an AI system with a high competence metaphor, this study shows a positive benefit of high competence — people are more likely to try out a new service if it is described with a high competence metaphor. This study also shows that metaphors that project high warmth also increase people's likelihood of trying out a service and to behave positively with it. We discuss the implications of these findings and suggest guidelines for choosing metaphors considering both the competing objectives of attracting more users and ensuring favorable evaluations and cooperative behavior.

7 DISCUSSION

Metaphors, as an expectation setting mechanism, are task- and model-agnostic. Users reason about complex algorithmic systems, including news feeds [19], content curation, and recommender

systems, using metaphors. This implies their effects are not limited to conversational agents or even to AI systems and can be used to set expectations of any algorithmic system (e.g., is Facebook's newsfeed algorithm a gossipy teen, an information butler, or a spy?), although the implications of our study might differ depending on the task, interaction and context.

With our findings in mind, this section explores their design implications and limitations, and situates our work amongst existing literature in HCI. We end with a retrospective analysis on existing and previous conversational AI products, reinterpreting their metaphors and adoption/user cooperation patterns through the lens of our results.

7.1 User behavior around algorithmic systems

Our work contributes to a growing body of work in HCI that seeks to understand how people reason about algorithmic systems with the aim of facilitating more informed and engaging interactions [19, 22, 26]. Previous work has looked at how users form informal theories about the technical mechanisms behind of social media feeds [19, 21, 26] and how these “folk theories” drive their interactions with these systems. People's conceptual understanding of such systems have been known to be metaphorical in nature, leading them to form folk theories of socio-technical systems in terms of metaphors. Folk theories for Facebook and Twitter news feeds include metaphors rooted in personas such as “rational assistant” and “unwanted observer” as well as metaphors tied to more abstract concepts such as “corporate black box” and “transparent platform”. More recent work has sought to study the social roles of algorithms by looking at how people personify algorithms and attach personas to them [81]. Prior work in the domain of interactive systems and embodied agents has observed that the mental schemas people apply towards agents affect the way they behave with the agent and it is possible to detect users' schematic orientation through initial interaction [47]. Diverging from previous work on folk theories, our work takes a complementary route — instead of studying *which* metaphors users attach to systems, we study *how* metaphors explicitly attached to the system, by designers, impact experiences.

7.2 Design implications

Studies 1 and 2 demonstrate that low competence metaphors lead to the highest evaluations of an AI agent, but Study 3 counters that agents with low competence metaphors are least likely to be tried out. What should a designer do?

From Study 3, it becomes clear that associating a high warmth metaphor is always beneficial—however, the choice of competence level projected by the metaphor becomes a more nuanced decision. One possible approach might be to choose a higher-competence metaphor but to lower competence expectations right after interaction begins (e.g., “Great question! One thing I should mention: I'm still learning how to best respond to questions like yours, so please have patience if I get something wrong.”) Another approach might be to age the metaphor over time: to present a high competence metaphor such as a professional, but when a user first encounters it, the agent introduces itself via a lower-competence version such as a professional trainee and tells the user that it will evolve over time into a full professional [65].

If designers are unwilling to change or adapt their high-competence metaphor, then their designs run the risk of being abandoned for being less effective than users expect. There may be other ways to disarm the contrast between expectations and reality. The agent blaming itself for errors or blaming the user for errors create challenging issues, but blaming an intermediary might work [53]: for example, “I've seen that previous folks who asked that question meant multiple different things by it. To make sure I can help effectively, can you reword that question?”

7.3 Limitations and future work

The scope of the study was limited to a conversational AI, as an instance of an algorithmic system, where interaction is devoid of strong visual cues [63]. In the case of embodied agents and systems where visual communication is a major aspect of the interaction, visual factors might have a strong effect on expectations. It is important to understand how users factor in these visual signals in forming an impression of the system. Additionally, our choice of conceptual metaphors was solely textual metaphors; future work should explore how these findings translate to visual metaphors such as the abstract shape associated with Siri, or the cartoonish rendering of Clippy, because such visual abstractions also inform users' judgements of a system's competence and warmth.

Since the task in our study was highly structured and participants had no incentive to explore peripheral conversational topics, we did not observe significant differences in user vocabulary across the conditions. This result surprised us — that evaluations would differ even if the interactions themselves had no major differences between conditions. Future work should explore user behavior in open-ended conversations, which are more likely to contain personal stories and anecdotes that can elicit greater behavior changes. The conversations and therefore, interactions with the AI system were limited to 15 – 20 minutes, so further research needs to establish the effects of metaphors on prolonged exposure to the AI system. Additionally, our service is not commonly used by people today to book flights or hotels and it is possible that the novelty of performing this task with a conversational agent might have skewed evaluations. This necessitates the need to understand how prior experience with similar technology changes people's susceptibility to such expectation shaping and subsequently their evaluations.

We observed partial support for Assimilation Theory along the warmth axis: participants preferred to cooperate with agents projecting higher warmth but at the same time, they perceived a difference between the agent's projected warmth and the actual warmth. Future work is needed to develop more robust theories along the warmth axis. One potential direction could create conditions of more extreme violation — sampling metaphors that signal either extremely low or extremely high warmth and measuring to see if participants' attitudes towards the agent are still driven by their pre-use warmth perceptions or whether the larger perceived difference in warmth alters their attitudes towards the agent.

Our study explored the effect of metaphors on adoption and behavioral intentions but these could also impact many other factors, including perceived trustworthiness of the system. Along this direction, future work should explore the impact on user evaluations when the interaction with the AI system results in a failure to accomplish the task. Finally, the actual competence and warmth of the AI system should be varied to analyze the effects of metaphors as the AI system's competence is lowered from our human-level performance.

7.4 Retrospective analysis

Studies have repeatedly shown that initial user expectations of their conversational agents (including Google Now, Siri, and Cortana) are not met [36, 49, 83], causing users to abandon these services after initial use. Initial experiences with conversational agents are often decisive: users reported that initial failures in achieving a task with Siri caused them to retreat to simple tasks they were sure the system could handle. While bloated expectations of users before interacting with AI systems have been acknowledged, little work has explored what those expectations are and how they contribute to user adoption and behavior.

Are today's conversational agents being set up for failure? Our studies establish that the descriptions and metaphors attached to these systems can play a key role in shaping expectations. Woebot was introduced as "Tiny conversations to feel your best"; Replika presaged as "The AI companion

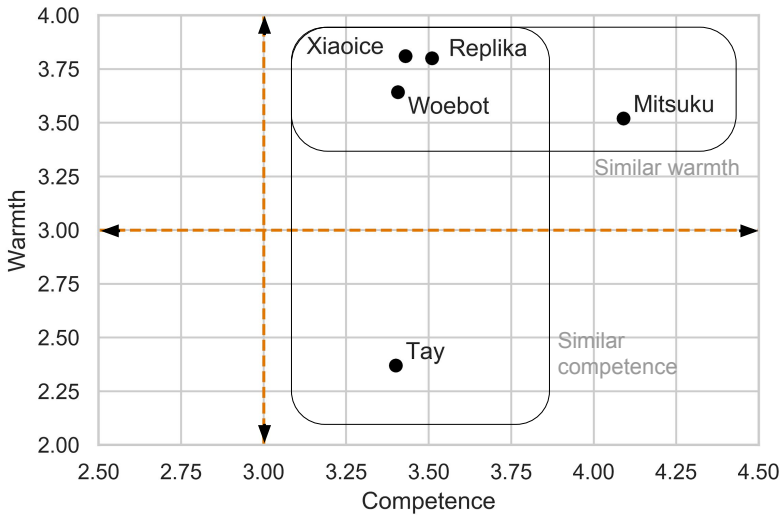


Fig. 7. Average warmth and competence measured for popular social chat-bots. Both axes range from 1 to 5.

who cares” and Mitsuku was revealed as “a record breaking five-time winner of the Loebner Prize Turing Test [...] the world’s best conversational chatbot”. We collected these descriptions associated with 5 popular social chatbots—Xiaoice, Mitsuku, Tay, Replika and Woebot—and deployed the exact same warmth-competence measurement of those descriptions with 50 participants from AMT as we used for the metaphors in our study.

We find that today’s social chatbots signal high competence (Figure 7), between “recent graduate” and “trained professional”. (Tay, incidentally, also projects very low warmth.) Descriptions of this kind, as we’ve shown, might be setting such systems up for failure. As *Ars Technica* reported: “You might come away thinking that Apple found some way to shrink *Mad Men*’s Joan Holloway and pop her into a computer chip. Though Siri shows real potential, these kinds of high expectations are bound to be disappointed” [13]. It is important to note, then, that users often report disappointment after using these agents, especially since Apple’s announcement of Siri included the sentence: “Ask Siri and get the answer back almost instantly without having to type a single character”; Google Assistant was heralded as “the first virtual assistant that truly anticipates your needs”.

With the recent glut of Twitter bots and other social agents that learn from their interactions and are adaptive in nature [56], it also becomes important to understand what drives users’ anti-social behaviour towards such bots and what factors contribute to antisocial behavior. Previous work has sought explanations through the lens of user profiling, gender attributions, and racial representations. Our work provides another lens on why otherwise similar systems such as Xiaoice and Tay (both female, teen-aged, and not representative of marginalised communities in their respective countries) might have elicited vastly different responses from their users. While Tay’s official Twitter account described it as “Microsoft’s AI fam from the internet that’s got zero chill!”, signaling high competence and low warmth, Xiaoice was setup to be “Sympathetic ear” and an “Empathetic social chatbot” [85], very clearly signaling high warmth and even priming behaviors around warmth such as personal disclosure. Our study suggests that people are more likely to cooperate with a bot that is perceived as higher warmth before use — a result consistent with the

fact that Xiaoice continued to be a friend and remained popular with its user base while Tay was pulled down within 16 hours of its release for attracting trolls.

Xiaoice is not an isolated case. Other bots, such as Woebot and Replika, which were set up as with high warmth, have had success in garnering users. Even though both Woebot and Replika had comparable competence expectations to that of Tay's, they were far warmer and obtained an altogether different outcome. Similarly, among the bots perceived as high warmth, Mitsuku stood out as as exceptionally competent and consistent with our finding that perceptions of very high competence decrease the desire to cooperate with the AI system: up to 30% of the messages received by Mitsuku comprise of antisocial messages [80]. For their part, Microsoft may have absorbed the lesson, as Tay's successor named Zo, was described more warmly as "Always down to chat. Will make you LOL".

While we acknowledge that there are several variables that affect user reception of these systems, the fact that our findings are consistent with in-the-wild outcomes of extant conversational systems is notable. It is, of course, impossible to prove that the expectations set by attached metaphors are a causal factor in the users' reception of these specific systems, and caution readers against concluding that metaphors alone are responsible.

8 CONCLUSION

We explore metaphors as a causal factor in determining users' evaluations of AI agents. We demonstrate experimentally that these conceptual metaphors change users' pre-use expectations as well as their post-use evaluations of the system, their intentions of adopt and their desire to cooperate. While people are more likely to cooperate with agents that they expect to be warm, they are more likely to adopt and cooperate with agents that project low competence. This result runs counter to designers' usual default towards projecting high competence to attract more users.

ACKNOWLEDGMENTS

We thank Jacob Ritchie, Mitchell Gordon and Mark Whiting for their valuable comments and feedback. This work was partially funded by the Brown Institute of Media Innovation and by Toyota Research Institute (TRI) but this article solely reflects the opinions and conclusions of its authors and not TRI or any other Toyota entity.

REFERENCES

- [1] Norah Abokhodair, Daisy Yoo, and David W McDonald. 2015. Dissecting a social botnet: Growth, content and influence in Twitter. In *Proceedings of the 18th ACM Conference on Computer Supported Cooperative Work & Social Computing*. ACM, 839–851.
- [2] Andrew Arnold. 2018. How Chatbots Feed Into Millennials' Need For Instant Gratification. (2018).
- [3] Solomon E Asch. 1946. Forming impressions of personality. *The Journal of Abnormal and Social Psychology* 41, 3 (1946), 258.
- [4] Layla El Asri, Hannes Schulz, Shikhar Sharma, Jeremie Zumer, Justin Harris, Emery Fine, Rahul Mehrotra, and Kaheer Suleman. 2017. Frames: A Corpus for Adding Memory to Goal-Oriented Dialogue Systems. *CoRR* abs/1704.00057 (2017). arXiv:1704.00057 <http://arxiv.org/abs/1704.00057>
- [5] Gagan Bansal, Ece Kamar, Walter S Lasecki, and Daniel S Weld Eric Horvitz. 2019. Beyond Accuracy: The Role of Mental Models in Human-AI Team Performance. (2019).
- [6] Amy L Baylor and Yanghee Kim. 2004. Pedagogical agent design: The impact of agent realism, gender, ethnicity, and instructional role. In *International conference on intelligent tutoring systems*. Springer, 592–603.
- [7] Susan A Brown, Viswanath Venkatesh, and Sandeep Goyal. 2012. Expectation confirmation in technology use. *Information Systems Research* 23, 2 (2012), 474–487.
- [8] Drazen Brscic, Hiroyuki Kidokoro, Yoshitaka Suehiro, and Takayuki Kanda. 2015. Escaping from children's abuse of social robots. In *Proceedings of the tenth annual acm/ieee international conference on human-robot interaction*. ACM, 59–66.

- [9] Rich Caruana, Yin Lou, Johannes Gehrke, Paul Koch, Marc Sturm, and Noemie Elhadad. 2015. Intelligible models for healthcare: Predicting pneumonia risk and hospital 30-day readmission. In *Proceedings of the 21th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*. ACM, 1721–1730.
- [10] Justine Cassell, Tim Bickmore, Lee Campbell, Hannes Vilhjálmsón, and Hao Yan. 2000. Embodied Conversational Agents. MIT Press, Cambridge, MA, USA, Chapter Human Conversation As a System Framework: Designing Embodied Conversational Agents, 29–63. <http://dl.acm.org/citation.cfm?id=371552.371555>
- [11] Elaine Chang and Vishwac Kannan. 2018. Conversational AI: Best practices for building bots. <https://medius.studios.ms/Embed/Video/BRK3225>
- [12] Ana Paula Chaves and Marco Aurélio Gerosa. 2019. How should my chatbot interact? A survey on human-chatbot interaction design. *CoRR* abs/1904.02743 (2019). arXiv:1904.02743 <http://arxiv.org/abs/1904.02743>
- [13] Jacqui Cheng. 2011. iPhone 4S: A Siri-ously slick, speedy smartphone. (2011).
- [14] Justin Cranshaw, Emad Elwany, Todd Newman, Rafal Kocielnik, Bowen Yu, Sandeep Soni, Jaime Teevan, and Andrés Monroy-Hernández. 2017. Calendar help: Designing a workflow-based scheduling agent with humans in the loop. In *Proceedings of the 2017 CHI Conference on Human Factors in Computing Systems*. ACM, 2382–2393.
- [15] L Elizabeth Crawford. 2009. Conceptual metaphors of affect. *Emotion review* 1, 2 (2009), 129–139.
- [16] Amy JC Cuddy, Susan T Fiske, and Peter Glick. 2008. Warmth and competence as universal dimensions of social perception: The stereotype content model and the BIAS map. *Advances in experimental social psychology* 40 (2008), 61–149.
- [17] Sonam Damani, Nitya Raviprakash, Umang Gupta, Ankush Chatterjee, Meghana Joshi, Khyatti Gupta, Kedhar Nath Narahari, Puneet Agrawal, Manoj Kumar Chinnakotla, Sneha Magapu, et al. 2018. Ruuh: A Deep Learning Based Conversational Social Agent. *arXiv preprint arXiv:1810.12097* (2018).
- [18] Leslie A DeChurch and Jessica R Mesmer-Magnus. 2010. The cognitive underpinnings of effective teamwork: A meta-analysis. *Journal of Applied Psychology* 95, 1 (2010), 32.
- [19] Michael A. DeVito, Jeremy Birnholtz, Jeffery T. Hancock, Megan French, and Sunny Liu. 2018. How People Form Folk Theories of Social Media Feeds and What It Means for How We Study Self-Presentation. In *Proceedings of the 2018 CHI Conference on Human Factors in Computing Systems* (Montreal QC, Canada) (*CHI '18*). ACM, New York, NY, USA, Article 120, 12 pages. <https://doi.org/10.1145/3173574.3173694>
- [20] Salesforce Drift, Survey Monkey. 2018. 2018 State of Chatbots Report. (2018).
- [21] Motahhare Eslami, Karrie Karahalios, Christian Sandvig, Kristen Vaccaro, Aimee Rickman, Kevin Hamilton, and Alex Kirlik. 2016. First I â€œIlikeâ€ It, Then I Hide It: Folk Theories of Social Feeds. In *Proceedings of the 2016 CHI Conference on Human Factors in Computing Systems* (San Jose, California, USA) (*CHI '16*). Association for Computing Machinery, New York, NY, USA, 2371â€2382. <https://doi.org/10.1145/2858036.2858494>
- [22] Motahhare Eslami, Aimee Rickman, Kristen Vaccaro, Amirhossein Aleyasen, Andy Vuong, Karrie Karahalios, Kevin Hamilton, and Christian Sandvig. 2015. â€œI Always Assumed That I Wasnâ€™t Really That Close to [Her]â€ Reasoning about Invisible Algorithms in News Feeds. In *Proceedings of the 33rd Annual ACM Conference on Human Factors in Computing Systems* (Seoul, Republic of Korea) (*CHI '15*). Association for Computing Machinery, New York, NY, USA, 153â€162. <https://doi.org/10.1145/2702123.2702556>
- [23] Emilio Ferrara, Onur Varol, Clayton Davis, Filippo Menczer, and Alessandro Flammini. 2016. The rise of social bots. *Commun. ACM* 59, 7 (2016), 96–104.
- [24] Susan T Fiske, Amy JC Cuddy, Peter Glick, and Jun Xu. 2018. A model of (often mixed) stereotype content: Competence and warmth respectively follow from perceived status and competition (2002). In *Social cognition*. Routledge, 171–222.
- [25] Susan T. Fiske, Juan Xu, Amy C. Cuddy, and Peter Glick. 1999. (Dis)respecting versus (Dis)liking: Status and Interdependence Predict Ambivalent Stereotypes of Competence and Warmth. *Journal of Social Issues* 55, 3 (1999), 473–489. <https://doi.org/10.1111/0022-4537.00128> arXiv:<https://spss.onlinelibrary.wiley.com/doi/pdf/10.1111/0022-4537.00128>
- [26] Megan French and Jeff Hancock. 2017. What's the folk theory? Reasoning about cyber-social systems. *Reasoning About Cyber-Social Systems (February 2, 2017)* (2017).
- [27] Jianfeng Gao, Michel Galley, and Lihong Li. 2018. Neural Approaches to Conversational AI. *CoRR* abs/1809.08267 (2018). arXiv:1809.08267 <http://arxiv.org/abs/1809.08267>
- [28] Andrew L Geers and G Daniel Lassiter. 1999. Affective expectations and information gain: Evidence for assimilation and contrast effects in affective experience. *Journal of Experimental Social Psychology* 35, 4 (1999), 394–413.
- [29] Susan A Gelman and Cristine H Legare. 2011. Concepts and folk theories. *Annual review of anthropology* 40 (2011), 379–398.
- [30] Jonathan Grudin and Richard Jacques. 2019. Chatbots, Humbots, and the Quest for Artificial General Intelligence. In *Proceedings of the 2019 CHI Conference on Human Factors in Computing Systems* (Glasgow, Scotland Uk) (*CHI '19*). ACM, New York, NY, USA, Article 209, 11 pages. <https://doi.org/10.1145/3290605.3300439>
- [31] Jan Hartmann, Antonella De Angeli, and Alistair Sutcliffe. 2008. Framing the user experience: information biases on website quality judgement. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*. ACM,

855–864.

- [32] Jan Hartmann, Antonella De Angeli, and Alistair Sutcliffe. 2008. Framing the User Experience: Information Biases on Website Quality Judgement. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems* (Florence, Italy) (CHI '08). ACM, New York, NY, USA, 855–864. <https://doi.org/10.1145/1357054.1357190>
- [33] P. J. Hinds, T. L. Roberts, and H. Jones. 2004. Whose job is it anyway? A study of humanrobot interaction in a collaborative task. (2004).
- [34] Annabell Ho, Jeff Hancock, and Adam S Miner. 2018. Psychological, relational, and emotional effects of self-disclosure after conversations with a chatbot. *Journal of Communication* 68, 4 (2018), 712–733.
- [35] Elle Hunt. 2016. Tay, Microsoft's AI chatbot, gets a crash course in racism from Twitter. *The Guardian* 24 (2016).
- [36] Mohit Jain, Pratyush Kumar, Ramachandra Kota, and Shwetak N Patel. 2018. Evaluating and informing the design of chatbots. In *Proceedings of the 2018 Designing Interactive Systems Conference*. ACM, 895–906.
- [37] Maurice Jakesch, Megan French, Xiao Ma, Jeffrey T Hancock, and Mor Naaman. 2019. AI-Mediated Communication: How the Perception that Profile Text was Written by AI Affects Trustworthiness. In *Proceedings of the 2019 CHI Conference on Human Factors in Computing Systems*. ACM, 239.
- [38] Charles M Judd, Laurie James-Hawkins, Vincent Yzerbyt, and Yoshihisa Kashima. 2005. Fundamental dimensions of social judgment: understanding the relations between judgments of competence and warmth. *Journal of personality and social psychology* 89, 6 (2005), 899.
- [39] Ralph Kimball and B Verplank E Harslem. 1982. Designing the Star user interface. *Byte* 7, 1982 (1982), 242–282.
- [40] Kristen J Klaaren, Sara D Hodges, and Timothy D Wilson. 1994. The role of affective expectations in subjective experience and decision-making. *Social Cognition* 12, 2 (1994), 77–101.
- [41] Rafal Kocielnik, Saleema Amershi, and Paul N. Bennett. 2019. Will You Accept an Imperfect AI?: Exploring Designs for Adjusting End-user Expectations of AI Systems. In *Proceedings of the 2019 CHI Conference on Human Factors in Computing Systems* (Glasgow, Scotland Uk) (CHI '19). ACM, New York, NY, USA, Article 411, 14 pages. <https://doi.org/10.1145/3290605.3300641>
- [42] Sari Kujala, Ruth Mugge, and Talya Miron-Shatz. 2017. The role of expectations in service evaluation: A longitudinal study of a proximity mobile payment service. *International Journal of Human-Computer Studies* 98 (2017), 51–61.
- [43] Eugenia Kuyda and Phil Dudchuk. 2015. Replika [computer program]. <https://replika.ai/>
- [44] Isaac Lage, Andrew Ross, Samuel J Gershman, Been Kim, and Finale Doshi-Velez. 2018. Human-in-the-loop interpretability prior. In *Advances in Neural Information Processing Systems*. 10159–10168.
- [45] George Lakoff and Mark Johnson. 2008. *Metaphors we live by*. University of Chicago press.
- [46] Justin Lee. 2018. Chatbots were the next big thing: what happened? <https://medium.com/swlh/chatbots-were-the-next-big-thing-what-happened-5fc49dd6fa61>
- [47] Min Kyung Lee, Sara Kiesler, and Jodi Forlizzi. 2010. Receptionist or information kiosk: how do people talk with a robot?. In *Proceedings of the 2010 ACM conference on Computer supported cooperative work*. 31–40.
- [48] Q Vera Liao, Mas-ud Hussain, Praveen Chandar, Matthew Davis, Yasaman Khazaeni, Marco Patricio Crasso, Dakuo Wang, Michael Muller, N Sadat Shami, Werner Geyer, et al. 2018. All Work and No Play? Conversations with a question-and-answer chatbot in the wild. In *Proceedings of the 2018 CHI Conference on Human Factors in Computing Systems*. ACM, 3.
- [49] Ewa Luger and Abigail Sellen. 2016. Like having a really bad PA: the gulf between user expectation and experience of conversational agents. In *Proceedings of the 2016 CHI Conference on Human Factors in Computing Systems*. ACM, 5286–5297.
- [50] Matthew S McGlone. 1996. Conceptual metaphors and figurative language interpretation: Food for thought? *Journal of memory and language* 35, 4 (1996), 544–565.
- [51] Jaroslav Michalco, Jakob Grue Simonsen, and Kasper Hornbæk. 2015. An exploration of the relation between expectations and user experience. *International Journal of Human-Computer Interaction* 31, 9 (2015), 603–617.
- [52] H. Mieczkowski, S. X. Liu, J. Hancock, and B. Reeves. 2019. Helping Not Hurting: Applying the Stereotype Content Model and BIAS Map to Social Robotics. In *2019 14th ACM/IEEE International Conference on Human-Robot Interaction (HRI)*. 222–229. <https://doi.org/10.1109/HRI.2019.8673307>
- [53] Clifford Ivar Nass and Scott Brave. 2005. *Wired for speech: How voice activates and advances the human-computer relationship*. MIT press Cambridge, MA.
- [54] Donald A Norman. 1988. *The psychology of everyday things*. Basic books.
- [55] AE Nutt. 2017. The Woebot will see you now. *the rise of chatbot therapy: Washington Post* (2017).
- [56] Junwon Park, Ranjay Krishna, Pranav Khadpe, Li Fei-Fei, and Michael Berstein. 2019. AI-based Request Augmentation to Increase Crowdsourcing Participation. In *AAAI Conference on Human Computation and Crowdsourcing*. ACM.
- [57] James W Pennebaker, Martha E Francis, and Roger J Booth. 2001. Linguistic inquiry and word count: LIWC 2001. *Mahway: Lawrence Erlbaum Associates* 71, 2001 (2001), 2001.

- [58] Eeva Raita and Antti Oulasvirta. 2011. Too good to be bad: Favorable product expectations boost subjective usability ratings. *Interacting with Computers* 23, 4 (2011), 363–371.
- [59] Byron Reeves and Clifford Ivar Nass. 1996. *The media equation: How people treat computers, television, and new media like real people and places*. Cambridge university press.
- [60] Cynthia Rudin. 2018. Please stop explaining black box models for high stakes decisions. *arXiv preprint arXiv:1811.10154* (2018).
- [61] Pericle Salvini, Gaetano Ciaravella, Wonpil Yu, Gabriele Ferri, Alessandro Manzi, Barbara Mazzolai, Cecilia Laschi, Sang-Rok Oh, and Paolo Dario. 2010. How safe are service robots in urban environments? Bullying a robot. In *19th International Symposium on Robot and Human Interactive Communication*. IEEE, 1–7.
- [62] C. Sandvig. 2015. Seeing the Sort: The Aesthetic and Industrial Defense of “The Algorithm.” *Journal of the New Media Caucus*, 11, 35–51 (2015).
- [63] Ayse Pinar Saygin, Thierry Chaminade, Hiroshi Ishiguro, Jon Driver, and Chris Frith. 2011. The thing that should not be: predictive coding and the uncanny valley in perceiving human and humanoid robot actions. *Social cognitive and affective neuroscience* 7, 4 (2011), 413–422.
- [64] Robin Sease. 2008. Metaphors’ role in the information behavior of humans interacting with computers. *Information technology and libraries* 27, 4 (2008), 9–16.
- [65] Joseph Seering, Michal Luria, Connie Ye, Geoff Kaufman, and Jessica Hammer. 2020. It Takes a Village: Integrating an Adaptive Chatbot into an Online Gaming Community. In *Proceedings of the 2020 CHI Conference on Human Factors in Computing Systems (CHI ’20)*.
- [66] Ameneh Shamekhi, Q Vera Liao, Dakuo Wang, Rachel KE Bellamy, and Thomas Erickson. 2018. Face Value? Exploring the effects of embodiment for a group facilitation agent. In *Proceedings of the 2018 CHI Conference on Human Factors in Computing Systems*. ACM, 391.
- [67] Muzafer Sherif, Daniel Taub, and Carl I Hovland. 1958. Assimilation and contrast effects of anchoring stimuli on judgments. *Journal of experimental psychology* 55, 2 (1958), 150.
- [68] Heung-Yeung Shum, Xiao-dong He, and Di Li. 2018. From Eliza to Xiaoice: challenges and opportunities with social chatbots. *Frontiers of Information Technology & Electronic Engineering* 19, 1 (2018), 10–26.
- [69] Megan Strait, Virginia Contreras, and Christian Duarte Vela. 2018. Verbal Disinhibition towards Robots is Associated with General Antisociality. *arXiv preprint arXiv:1808.01076* (2018).
- [70] Megan Strait, Ana Sánchez Ramos, Virginia Contreras, and Noemi Garcia. 2018. Robots Racialized in the Likeness of Marginalized Social Identities are Subject to Greater Dehumanization than those racialized as White. In *2018 27th IEEE International Symposium on Robot and Human Interactive Communication (RO-MAN)*. IEEE, 452–457.
- [71] Indrani Medhi Thies, Nandita Menon, Sneha Magapu, Manisha Subramony, and Jacki O’Neill. 2017. How do you want your chatbot? An exploratory Wizard-of-Oz study with young, urban Indians. In *IFIP Conference on Human-Computer Interaction*. Springer, 441–459.
- [72] Alexander Todorov, Chris P Said, Andrew D Engell, and Nikolaas N Oosterhof. 2008. Understanding evaluation of faces on social dimensions. *Trends in cognitive sciences* 12, 12 (2008), 455–460.
- [73] Paul Van Schaik and Jonathan Ling. 2008. Modelling user experience with web sites: Usability, hedonic value, beauty and goodness. *Interacting with computers* 20, 3 (2008), 419–432.
- [74] Peter Wallis and Emma Norling. 2005. The Trouble with Chatbots: social skills in a social world. *Virtual Social Agents* 29 (2005).
- [75] Mark E Whiting, Grant Hugh, and Michael S Bernstein. 2019. Fair Work: Crowd Work Minimum Wage with One Line of Code. In *Proceedings of the AAAI Conference on Human Computation and Crowdsourcing*, Vol. 7. 197–206.
- [76] Timothy D Wilson, Douglas J Lisle, Dolores Kraft, and Christopher G Wetzell. 1989. Preferences as expectation-driven inferences: Effects of affective expectations on affective experience. *Journal of personality and social psychology* 56, 4 (1989), 519.
- [77] Bogdan Wojciszke. 2005. Affective concomitants of information on morality and competence. *European psychologist* 10, 1 (2005), 60–70.
- [78] Victoria Woollaston. 2016. Following the failure of Tay, Microsoft is back with new chatbot Zo.
- [79] S Worswick. 2015. Mitsuku [computer program].
- [80] Steve Worswick. 2018. The Curse of the Chatbot Users. (2018).
- [81] Eva Yiwei Wu, Emily Pedersen, and Niloufar Salehi. 2019. Agent, Gatekeeper, Drug Dealer: How Content Creators Craft Algorithmic Personas. *Proceedings of the ACM on Human-Computer Interaction* 3, CSCW (2019), 219.
- [82] Jun Xiao, John Stasko, and Richard Catrambone. 2004. An empirical study of the effect of agent competence on user performance and perception. In *Proceedings of the Third International Joint Conference on Autonomous Agents and Multiagent Systems-Volume 1*. IEEE Computer Society, 178–185.
- [83] Jennifer Zamora. 2017. I’m Sorry, Dave, I’m Afraid I Can’t Do That: Chatbot Perception and Expectations. In *Proceedings of the 5th International Conference on Human Agent Interaction (Bielefeld, Germany) (HAI ’17)*. ACM, New York, NY,

- USA, 253–260. <https://doi.org/10.1145/3125739.3125766>
- [84] Mark P Zanna and David L Hamilton. 1972. Attribute dimensions and patterns of trait inferences. *Psychonomic Science* 27, 6 (1972), 353–354.
- [85] Li Zhou, Jianfeng Gao, Di Li, and Heung-Yeung Shum. 2018. The Design and Implementation of XiaoIce, an Empathetic Social Chatbot. *CoRR* abs/1812.08989 (2018). arXiv:1812.08989 <http://arxiv.org/abs/1812.08989>