Artificial Intelligence and Persuasion: A Construal-Level Account

(Invited for 2nd round revision at *Psychological Science*, Job Market Paper)

TAE WOO KIM

ADAM DUHACHEK

Tae Woo Kim (kim805@indiana.edu) is a doctoral candidate in marketing at Indiana University, Bloomington, IN, 47405. Adam Duhachek (aduhache@indiana.edu) is the Nestle-Hustad Professor of Marketing at Indiana University, Bloomington, IN, 47405. This research proposal is based on the first author's dissertation. The authors would like to thank Selma Sabanovic for her support in data collection, Eliot Smith for his helpful comments, and the seminar participants at the R-House Lab and Socially Situated Cognition Lab at Indiana University. The preparation of this article was supported by a research grant from the Marketing Science Institute (MSI) and the Kelley School of Business at Indiana University.

Abstract

Whereas more individuals are relying on information provided by non-human agents, such as artificial intelligence and robots, little research has examined how persuasion attempts made by non-human agents may differ from persuasion attempts made by human agents. Drawing on construal level theory, we posit that individuals will perceive artificial agents at a low-level construal, which directs individuals' focus towards "how" these agents implement actions to serve humans. Further, we posit that interactions with such agents drive individuals to adopt a low-level construal mindset more generally. We show these construal-based differences impact compliance with persuasive messages made by artificial agents such that these messages are more effective when the message represents low- (vs. high-) level construal features. We find these effects are moderated by the extent to which an artificial agent can independently learn from its environment (i.e., machine learning), as learning defies people's lay theories about artificial agents.

Keywords: *artificial intelligence, robot, decision making, construal level theory, machine learning*

Artificial Agents and Persuasion

Intelligent non-human agents, such as artificial intelligence and robots, referred to as artificial agents by researchers (Cohen & Levesque, 1995; Floridi & Sanders, 2004; Russel & Norvig, 2010), are increasingly common across a variety of platforms and applications created to interact with human decision-makers. Artificial agents (AA[1] hereafter) are being increasingly used in domains such as medical and financial decision-making and product recommendation systems. Some AA such as IBM's Watson have already proven superior to human doctors in diagnosing certain diseases such as skin cancer (Esteva et al., 2017; Leachman & Merlino, 2017; McFarland, 2016). Despite these advances, research has shown many people are still averse to interactions with AA. The current research examines how AA systems can increase persuasive effectiveness by drawing on construal level theory, a previously unexamined dimension within people's mental representations of AA. The findings suggest a novel paradigm through which AA may be studied and human interactions with AA made more effective. This research also identifies another factor integral to persuasion in human-AA interactions, namely, the role of lay theories of learning by AA, and posits that beliefs about AA' learning capabilities moderate the impact of construal level perceptions on persuasion.

First, this research identifies a key linkage between AA and construal level perceptions among humans with which AA interact. According to action identification theory (Vallacher & Wegner, 1987), an action (e.g., "painting a room") can be represented at a high-level of construal emphasizing why the action is conducted (e.g., "making the room look fresh") or at a low-level of construal processing how the action is conducted (e.g., "applying brush strokes"). Construal level theory (Trope & Liberman, 2010) posits that whether the same action is represented at a high-level (vs. low-level) of construal depends on the presence

[1] The term "AA" will be used to refer to plural agents (i.e., artificial agents) and "an AA" will be used to refer to a singular artificial agent.

of superordinate goals associated with the action. Notably, AA are designed by humans to serve, and the common lay theory individuals hold is that AA's actions are not driven by their own autonomous goals or intentions (Russel & Norvig, 2010; Ward, Olsen, & Wegner, 2013). In other words, AA are more likely to be perceived by humans as agents that are controlled by preprogrammed algorithms rather than acting on their own intentions. Research has shown that such lack of perceived intention in agents inhibits the inference of superordinate goals derived from these actions, thus leading to low-level identifications for these actions among observers (Kozak, March, & Wegner, 2006). Therefore, we hypothesize that individuals will interpret AA's actions at lower level construals compared to the same action performed by humans. In addition, we hypothesize that observing AA's actions may induce low-level construal mindsets among observing individuals. This prediction is based on prior research showing that focusing on low-level (vs. high-level) construals of actions (e.g., how actions are conducted) leads to adopting the same mindset (i.e., low-construal mindset) among observing individuals (Liberman, Trope, McCrea, & Sherman, 2007; Fujita, Trope, Liberman, & Levin-Sagi, 2006).

In the context of persuasion, the prediction that AA would be perceived as low-construal agents has an important implication. The persuasion literature has shown that persuasion is influenced by multiple factors such as the message's content (Petty & Cacioppo, 1979), the source of the message (Petty & Cacioppo, 1984), and recipient's characteristics, such as their psychological states (Han, Duhachek, & Agrawal, 2016; Lee, Keller, & Sternthal, 2010). Prior research has shown that persuasion is more effective when the message's construal level matches the message recipient's mindset because it increases processing fluency (Han, Duhachek, & Agrawal, 2014; Han et al., 2016; Lee et al., 2010). Therefore, we hypothesize that persuasive messages from AA will be more effective when

the persuasive message highlights low-level (vs. high-level) construal arguments to fit the message recipient's mindset due to increased processing fluency resulting from this match.

One factor that has heretofore not been examined in the literature related to human perceptions of AA relates to lay theories regarding AA. While still relatively nascent and unbeknownst to many individuals, advanced versions of AA are capable of greatly increasing their functionality over time through machine learning (Franklin & Graesser, 1996; Huang & Rust, in press; Shandwick, 2016). Notably, these advanced AA can increase understanding of their environment over time and potentially derive superordinate meanings from their experience just as humans do (Russell & Norvig, 2010). Compared to non-learning AA, learning AA are driven by a superordinate goal: to improve their capabilities with experiences. Thus, learning AA have more superordinate goals than non-learning AA and these superordinate goals represent high-level construals (Trope & Liberman, 2010). Thus, we predict that informing people about AA' learning capabilities will attenuate perceptions of AA as low-construal agents.

Overview of experimental studies

In Studies 1a and 1b, we provide initial evidence that AA are not only perceived as low-construal agents, but also that they induce low-construal mindsets among observers. In Study 2, we show that messages from AA are more persuasive when they highlight low-level (vs. high-level) construals. In Studies 3 and 4, we extend these effects to interactions using a commercially sold intelligent virtual agent product (Amazon Echo) and a commercially sold humanoid robot (Baxter).

Studies 1a and 1b

In Study 1a, we examined whether individuals focus on low-level (vs. high-level) construal features of the actions conducted by AA compared to when the same actions are

conducted by humans. In study 1b, we examined whether such attentional foci on low-level construals induce low-construal mindsets among individuals observing AA's actions.

Method

Participants and design. Study 1a had one hundred participants from M-Turk (female = 45%, M_{age} = 35.57) and Study 1b had eighty participants from M-Turk (female = 64%, M_{age} = 35.38). Our sample size was determined based on previous AA research (Waytz & Norton, 2014) that used similar experimental methodologies. Based on the effect sizes reported in prior work on AA (Waytz & Norton, 2014), an a priori power analysis indicated that about 40 participants per condition would provide sufficient power. Therefore, we determined to recruit at least 40 participants per condition. Additionally, sample size was influenced by the size of the participant pool made available to the authors in the given semester. In both studies, participants were randomly assigned to one of the conditions in a 2 (agent: human, AA) between-subjects design.

Procedure. In both Studies 1a and 1b, our dependent variable was the Behavioral Identification Form (BIF), a scale developed to measure the extent to which an individual perceives an action consistent with either low- or high-level construals (Vallacher & Wegner, 1989). As an example item from the BIF, participants were given an action (e.g., "painting a room"), and were asked to choose between the two descriptions of the action; one that described the action in low-level construals (e.g., "applying brush strokes") and another that described the action in a high-level construals (e.g., "making the room look fresh").

In Study 1a, participants were instructed to imagine that the actions provided in the BIF were conducted by either a human or an AA, depending upon condition. Then, participants chose the descriptions that they thought were more appropriate for the agent that they were assigned to (i.e., either a human or an AA). Among the twenty-five items in the BIF, seven items that did not fit as actions appropriate for an AA (e.g., "brushing teeth") and

were excluded. The remaining eighteen items were provided to participants (for more information, see Study 1a in the Supplemental Material available online). It was our prediction that individuals who imagined that the actions in the BIF were conducted by an AA (vs. a human) would emphasize lower levels of construal for these actions.

A notable difference between Study 1a and 1b was that participants in Study 1b were first instructed to imagine that an action (i.e., cooking pizza) was conducted by either a human or an AA (i.e., an intelligent robot capable of cooking), depending on the condition they were assigned to, and then answered the BIF. In particular, participants in the human (AA) condition were instructed to imagine and write down actions that a person (a robot) would have conducted to make a pizza from scratch. We chose cooking as the focal action to be imagined because cooking is one of the mainstream tasks of current domestic robot technologies (Huen, 2016). Then, participants in Study 1b answered the twenty-five items in the BIF as a part of an unrelated psychology study on behavioral perception (for more information, see Study 1b in the Supplemental Material available online). It was our prediction that individuals who imagined cooking by an AA (vs. a human) will focus more on the low-level construals of the action (e.g., how the pizza is made) while paying less attention to high-level construals of the action (e.g., why the pizza is made). Therefore, we predicted that this effect will carry over and influence the construal mindset of observers as captured in the BIF. Participants in both studies answered demographic questions and were debriefed.

The key difference between Studies 1a and 1b in their procedure was that participants in Study 1a answered the BIF while imagining that the actions provided in the BIF were conducted by either a human or an AA whereas participants in Study 1b answered the BIF after envisioning an unrelated action by either a human or an AA. That is, the BIF in Study 1a was used to examine whether individuals perceive actions conducted by a human or

an AA differently whereas the BIF in Study 1b was used to examine whether envisioning actions conducted by a human or an AA *induces* different levels of construal.

Results

In Studies 1a and 1b, we summed the number of high-construal choices made by each participant to create an index of construal level. The index ranged from 0 to 18 in Study 1a and it ranged from 0 to 25 in Study 1b. In Studies 1a and 1b, the index of construal level was submitted to a 2 (agent: human, AA) ANOVA. In Study 1a, the construal level score was significantly lower when participants imagined that the action was conducted by an AA (M = 4.04, SD = 5.51, 95% CI = [2.44, 5.64]) than when the action was conducted by a human (M = 11.28, SD = 5.88), 95% CI = [9.68, 12.88]), $F(1, 98) = 40.36$, $p < .001$, $\eta_p^2 = .29$. In Study 1b, construal level score was significantly lower after participants imagined that the unrelated previous action (i.e., cooking) was conducted by an AA (M = 12.62, SD = 6.98, 95% CI = [10.68, 14.56]) than a human (M = 16.53, SD = 5.50, 95% CI = [14.48, 18.57]), $F(1, 78) = 7.62$, $p = .007$, $\eta_p^2 = .09$.

Further, in Study 1b, there was no significant difference between the human and AA conditions in terms of the time that participants spent on the writing task, $F(1, 78) = .51$, $p = .478$, or the length of writing as measured by the number of words used, $F(1, 78) = .25$, $p = .619$. In both Studies 1a and 1b, gender and age did not have a main effect or interaction with the treatment conditions. These results related to demographic variables were observed in all studies and thus these variables will not be discussed further.

Discussion

As predicted, an AA's actions was perceived at lower level of construal compared to a human agent (Study 1a). Furthermore, this shift in attentional focus extends to induce low-level construals among individuals observing these agents (Study 1b).

Study 2

The purpose of Study 2 was to test whether a message from an AA is more persuasive when the message represents low-level (vs. high-level) construals.

Method

Participants and design. Three hundred twenty four participants from M-Turk (female = 54%, M_{age} = 37.30) were randomly assigned to one of the conditions in a 2 (agent: human, AA) x 2 (message: high-construal, low-construal) between-subjects design.

Procedure. Participants in the human conditions were instructed to browse an online medical website (www.sermo.com) and were told that they would receive instant medical advice from a doctor online. Participants in the AA conditions were instructed to browse a webpage describing IBM's artificial intelligence Watson's abilities related to medical diagnosis of skin cancer (https://www.ibm.com/cognitive/au-en/melanoma) and were told that they would receive instant medical advice from Watson. A pretest with eighty M-Turk participants confirmed that the level of trust individuals held in the two medical agents was not significantly different (M_{watson} = 5.28 vs. $M_{human\text{-}doctor}$ = 5.46), $F(1, 78) = .62$, $p = .433$.

Then, participants filled out a questionnaire that purported to measure their risk of skin cancer (e.g., age, gender, residential area, hours of daily sunlight exposure, skin cancer history in the family etc.), after which they received medical advice from either Watson or a human doctor based on their answers. Based on prior construal level research (Lee et al., 2010), the message either highlighted high- or low-level construal prescriptive statements regarding sunscreen use (for more information, see Study 2 in the Supplemental Material available online).

Next, participants indicated their intention to apply sunscreen using measures adapted from Keller (2006) (e.g., "I intend to use sunscreen"; 1 = very unlikely and 7 = very likely, the four items were averaged to create an index of intention, $\alpha = .91$). Then, participants answered demographic questions and were debriefed.

Results

We first submitted the intention to use sunscreen to a 2 (agent: human, AA) x 2 (message: high-construal, low-construal) ANOVA. The result revealed a non-significant main effect of the agent, $F(1, 320) = .62$, $p = .432$, $\eta_p^2 = .002$, and message on the intention to use sunscreen, $F(1, 320) = 1.81$, $p = .180$, $\eta_p^2 = .01$, The hypothesized interaction between agent and message was significant, $F(1, 320) = 5.92$, $p = .016$, $\eta_p^2 = .02$. Additional pairwise comparisons revealed that when the message was coming from an AA, the low-construal message ($M = 4.98$, $SD = 1.48$, 95% CI = [4.65, 5.31]) was more effective than the high-construal message ($M = 4.30$, $SD = 1.72$, 95% CI = [3.93, 4.68]) in persuading individuals to use sunscreen, $F(1, 320) = 6.96$, $p = .009$, $\eta_p^2 = .04$ (Fig. 1). Also, the low-construal message led to a higher intention to use sunscreen when it was coming from an AA ($M = 4.98$, $SD = 1.49$, 95% CI = [4.65, 5.31]) than when the message was coming from a human ($M = 4.40$, $SD = 1.63$, 95% CI = [4.05, 4.76]), $F(1, 320) = 5.41$, $p = .021$, $\eta_p^2 = .03$ (Fig. 1). Additionally, when the message was coming from a human, there was no significant difference between the low-construal message ($M = 4.40$, $SD = 1.62$, 95% CI = [4.05, 4.76]) and the high-construal message ($M = 4.60$, $SD = 1.58$, 95% CI = [4.26, 4.93]) in persuading individuals to use sunscreen, $F(1, 320) = .61$, $p = .437$, $\eta_p^2 < .01$.

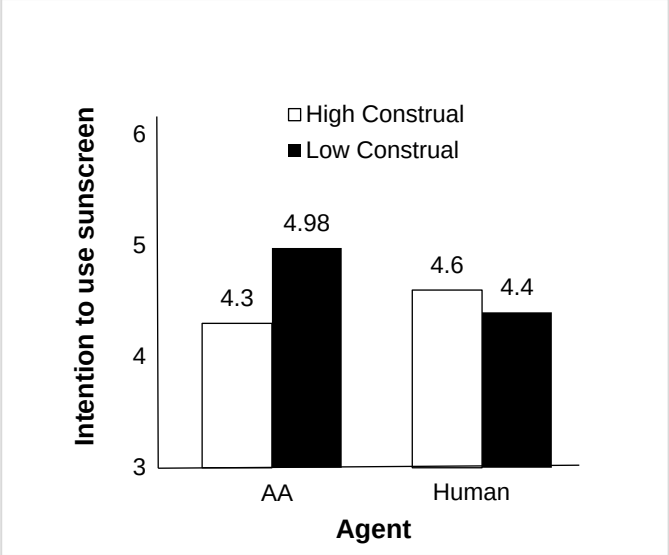

Fig. 1. Intention to use sunscreen as a function of agent and message's construal level (Study 2)

Discussion

Study 2 showed that low-construal (vs. high-construal) messages were more effective when the message was delivered by an AA, but not when the source was human. This finding extends the results of Studies 1a and 1b into the persuasion context and demonstrates that persuasion is more effective when the construal level of the message matches the construal level of the message source.

Study 3

The primary objective of Study 3 was to replicate the findings in Study 2 using actual interactions with an AA. In particular, participants interacted with Amazon Echo's virtual agent Alexa. Another important objective of Study 3 was to test the hypothesized effect that perceving an AA as capable of learning would attenuate the default low-construal mindset induced by the AA.

Method

Participants and design. One hundred and ninety-two undergraduate business students (female = 48%, M_{age} = 20.21) were randomly assigned to one of the conditions in a

2 (AA's capability: fixed, learning) x 2 (message construal: high-construal, low-construal) between-subjects design.

Procedure. The experiment took place in the university behavioral lab. Upon arrival to the lab, participants were seated in one of nineteen cubicles in front of a computer. In the middle of the lab space, Echo, an AA created by the company Amazon, was placed on a table and made visible to each participant. As a cover story, participants were told that they will receive a personalized product recommendation from Amazon Echo's virtual agent Alexa.

First, participants were instructed to read a description about Alexa, which was provided in their computer screens. To manipulate the belief about Alexa's capability, participants were randomly assigned to one of two conditions which described Alexa as an agent with either a fixed or experiential learning capability (for more information, see Study 3 in the Supplemental Material available online).

Then, Alexa spoke to participants and asked them to list up to ten items that they had purchased recently. After all participants spent 90 seconds on listing the items on their computers, Alexa told participants that she received their past purchase histories and told participants that she would send personalized product recommendations to each participant's computer screen. Therefore, the participants believed that their message was individualized and different from the message that other participants received.

Then, participants randomly received on their computer either a high- or low-construal message that advertised a gym. This stimuli were adapted from Han et al. (2016) and was modified for the purpose of the current experiment (for more information, see Study 3 in the Supplemental Material available online). A pretest with eighty M-Turk participants confirmed that participants who were exposed to the two ads showed comparable levels of intention to visit the gym ($M_{\text{why-ad}}$ = 2.65 vs. $M_{\text{how-ad}}$ = 2.98), $F(1, 78) = .73$, $p = .396$.

Then, participants indicated the extent to which they intended to visit the gym ("How likely is it that you would try out this gym in case you were looking for a place at which to work out?"; 1 = not at all likely, 7 = very likely), which served as the measure of persuasion effectiveness. To measure whether the manipulation was successful, we also measured the extent to which participants believed AA's capability is fixed ("Alexa can only conduct the task that it is programmed to do"; 1 = strongly disagree, 7 = strongly agree). Finally, participants answered demographic questions and were debriefed.

Results

First, the measure of fixed capability belief (i.e., "Alexa can only conduct the task that it is programmed to do") was submitted to a 2 (AA's capability: fixed, learning) x 2 (message: high-construal, low-construal) ANOVA to assess the effectiveness of this manipulation on beliefs about the AA's capabilities. The results revealed that only the main effect of the AA's capability manipulation was significant such that the fixed capability condition had a stronger belief that Alexa can only conduct tasks that it was programmed to do (M = 4.44, SD = 1.05, 95% CI = [4.21, 4.67]) as compared to the learning capability condition did (M = 3.88, SD = 1.24, 95% CI = [3.65, 4.11]), $F(1, 188) = 11.27$, $p = .001$, $\eta_p^2 = .06$. Therefore, our manipulation of the AA's capability was successful.

Then, the intention to visit the gym was submitted to a 2 (AA's capability: fixed, learning) x 2 (message: high-construal, low-construal) ANOVA. Neither the main effect of AA's capability, $F(1, 188) = .19$, $p = .667$, nor the main effect of message were significant, $F(1, 188) < .001$, $p = .987$. However, the interaction between AA's capability and message was significant, $F(1, 188) = 5.93$, $p = .016$, $\eta_p^2 = .03$.

Additional pairwise comparisons in the fixed capability condition showed that the low-construal message (M = 3.89, SD = 1.83, 95% CI = [3.38, 4.40]) was marginally more effective than the high-construal message (M = 3.27, SD = 1.73, 95% CI = [2.77, 3.77]), F(1,

188) = 3.01, p = .084, η_p^2 = .03. In the learning capability condition, it was shown that the high-construal message (M = 4.00, SD = 1.91, 95% CI = [3.48, 4.51]) was marginally more effective than the low-construal message (M = 3.38, SD = 1.63, 95% CI = [2.89, 3.88]), F(1, 188) = 2.93, p = .089, η_p^2 = .03 (Fig. 2). Additionally, the high-construal message was more effective in the learning capability condition (M = 4.00, SD = 1.91, 95% CI = [3.48, 4.51]) than in the fixed capability condition (M = 3.27, SD = 1.73, 95% CI = [2.77, 3.77]), F(1, 188) = 4.07, p = .045, η_p^2 = .04.

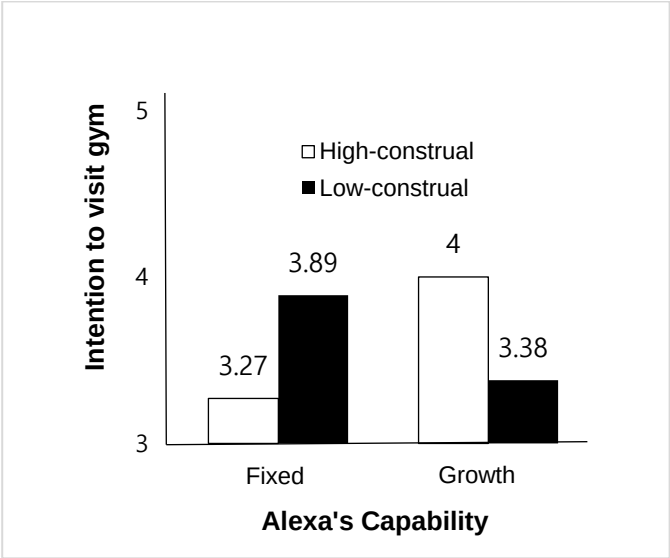

Fig. 2. Intention to visit gym as a function of Alexa's capability and message's construal level (Study 3)

Discussion

The result of Study 3 showed that message delivered from an AA is more persuasive when the message's construal level matches the perceived capability of the AA. Importantly, this study showed that beliefs on AA's learning beliefs are capable of overriding the fundamental linkage between low-construal and perceptions of AA, thereby impacting persuasive effectiveness.

Study 4

Study 4 attempted to replicate the findings from the previous studies using a different type of AA, a humanoid robot Baxter (produced by ReThink Robotics, http://www.rethinkrobotics.com). Just as in previous studies, it was our prediction that a low-construal message would be more effective than a high-construal message. Another objective of Study 4 was to further examine whether perceiving an AA as capable of learning would attenuate the default low-construal mindset induced by the AA. A notable distinction made in Study 4 was that the learning capability of the AA was measured whereas it was manipulated in Study 3. In order to ensure that participants' construal level is uniquely moderated by the belief of AA's learning capability but not by other aspects of AA, we also measured the AA's perceived cognitive and experiential capability. These variables were shown to be important to forming meantal perception of non-human agents (Gray, Gray, & Wegner, 2007).

Method

Participants and design. One hundred undergraduate business students (female = 56%, M_{age} = 20.50) were randomly assigned to one of the conditions in a 2 (message: high-construal, low-construal) between-subjects design.

Procedure. The experiment took place in the Informatics Department's robotics lab in which Baxter was located. Upon arrival to the lab, participants provided their email address to the experimenter and the experimenter sent them a Qualtrics survey link, which participants were to use during their participation in the experiment via their cell phones. Then, the experimenter sent participants to the robotics lab in groups of three with instructions not to speak to each other during the study. Once participants entered the lab, Baxter welcomed them verbally and also by waving its hand. All motions (i.e., hand waving and occasional head movement) and speech by Baxter were predetermined and operated by another experimenter whose presence was unbeknownst to participants. First, Baxter instructed the participants to bring out their cell phones and open the survey link sent to their

email accounts. Then, just as in Study 3, Baxter asked participants to list up to ten items that they had purchased recently. After all participants spent 90 seconds on this task, Baxter told participants that he received their past purchase histories and told participants that he would send a personalized product recommendation to their cell phones.

Then, just as in Study 3, participants randomly received either a high- or low-construal message that advertised a gym. The dependent variable was the same as in Study 3: participants indicated the extent to which they intended to visit the gym ("How likely is it that you would try out this gym in case you were looking for a place at which to work out?"; 1 = not at all likely, 7 = very likely).

Also, using the same measure as in Study 3, we measured the extent to which participants believed Baxter's capability is fixed ("Baxter can only conduct the task that it is programmed to do; 1 = strongly disagree, 7 = strongly agree; M = 5.05, SD = 1.46, 95% CI = [4.62, 5.48]). The purpose of this measure was to examine the extent to which individuals think Baxter's capability is fixed (rather than having the ability to learn new capabilities via machine learning) and to examine whether this measure interacts with the effectiveness of the messages varying in construal level. Because prior research has shown that the perceived level of agency and experiential capability are important to forming mental perception of non-human agents (Gray et al., 2007), we also measured Baxter's perceived agency ("Baxter can think and judge like humans"; 1 = strongly disagree, 7 = strongly agree) and experiential capacity ("Baxter can sense and feel like humans"; 1 = strongly disagree, 7 = strongly agree). Finally, participants answered demographic questions and were debriefed.

Results

Eight participants communicated with other participants and shared their message from Baxter with other participants. These participants were excluded from the analysis, leaving ninety two participants in the analysis. We first submitted intention to visit the gym to

a 2 (message: high-construal, low-construal) ANOVA to assess persuasive effectiveness. The results revealed that the intention to visit the gym was greater in the low-construal message condition (M = 3.74, SD = 1.72, 95% CI = [3.25, 4.23]) than in the high-construal message condition (M = 2.93, SD = 1.60, 95% CI = [2.45, 3.42]), $F(1, 90) = 5.41$, $p = .022$, $\eta_p^2 = .06$.

To examine whether the message's construal level interacted with the measure of fixed capability belief (i.e., the belief that Baxter can only conduct the task that it is programmed to do), we conducted a regression analysis with the message construal variable (high-construal message coded as 0, low-construal message coded as 1), the continuous measure of fixed capability belief, and their interaction predicting the intention to visit the gym. From this regression analysis, we found only a significant interaction between the fixed capability belief and the message construal factor, $b = .52$, $SE = .24$, 95% CI = [.05, .98], $t(88) = 2.19$, $p = .031$. To better understand this interaction, we used the Johnson-Neyman technique (Spiller, Fitzsimons, Lynch, & McClelland, 2013) to identify the range of fixed capability belief for which the simple effect of message construal was significant. This analysis revealed that low-construal (vs. high-construal) messages led to stronger behavioral intentions for individuals with fixed capability beliefs greater than or equal to 4.77, $b = .68$, $SE = .34$, 95% CI = [.00, 1.36], $p = .05$ (Fig. 3). Thus, the low-construal message was more effective when individuals believed that Baxter has fixed intelligence and is only capable of conducting the task it is pre-programmed to do. In other words, the relative effectiveness of low- (vs. high) construal message was attenuated for the individuals who believed that Baxter can learn and improve its capability over time.

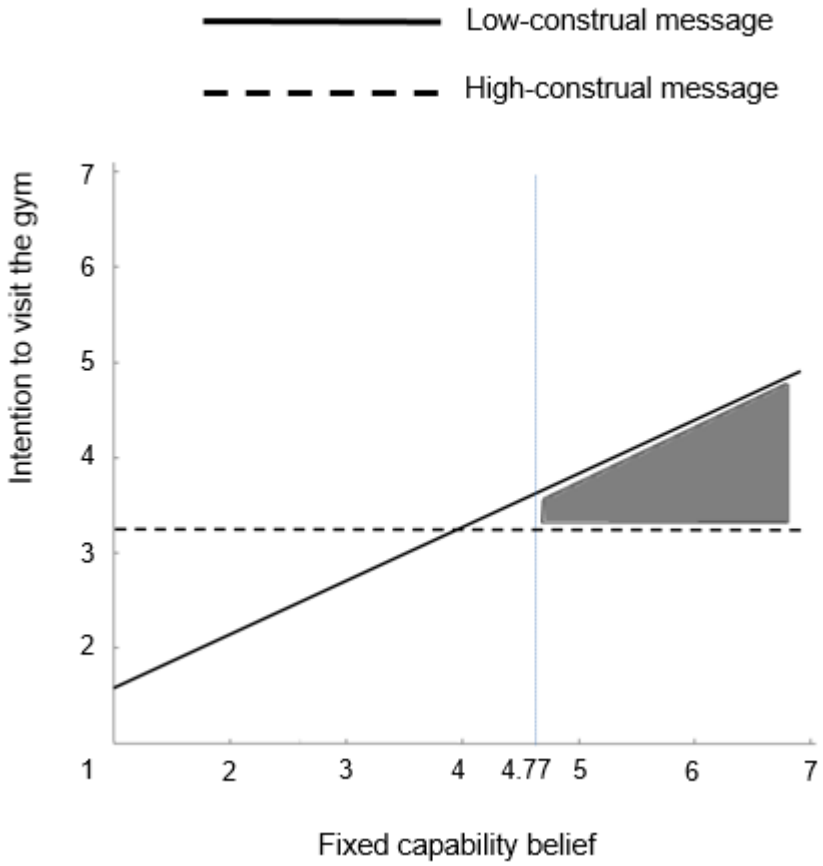

Fig. 3. The effect of message's construal level depends on the extent that people believe AA's capability is fixed (Study 4)

Next, we examined whether the perceived level of Baxter's experiential or cognitive capability interacted with message construal level. In a regression analysis with message construal (high-construal message coded as 0, low-construal message coded as 1), the continuous measure of Baxter's cognitive ability, and their interaction predicting intentions to visit the gym, we did not observe a significant interaction between the two variables, $p = .515$. In a separate regression analysis with the message construal (high-construal message coded as 0, low-construal message coded as 1), the continuous measure of Baxter's emotional ability, and their interaction predicting the intention to visit the gym, we did not find a significant interaction between the two variables, $p = .357$.

Discussion

Study 4 replicated and extended the findings of the previous studies by showing that low-construal (vs. high-construal) messages are more effective when the persuasion agent is an AA. Furthermore, Study 4 revealed that this effect holds only for individuals who believed that Baxter can perform only as much as it is programmed (i.e., fixed capabilities) but not among individuals who believed that Baxter can learn and improve its capabilities via experiences.

General discussion

The current research draws on construal level theory and demonstrates that AA are not only perceived as low-construal agents, but that they also activate low construal mindsets among humans with which they interact. We show that these construal differences lead to enhanced persuasion when the message's construal level fits with the low construal mindset of the message recipient. These findings contribute to the literature of construal level theory, persuasion, and the growing body of research on non-human agents by identifying, for the first time, the unique difference between AA and human agents through a construal level theoretical lens and demonstrating the consequences of this nexus for persuasive communications between humans and AA. Building off of construal theory, the current research suggests that individuals would rely on AA more heavily in contexts that induce low construal: events or objects that are temporally, socially, and spatially proximate, and occurring with high probability. For example, a short- (vs. long-) term investor may rely on a financial AA more than a human financial advisor (Trope & Liberman, 2003).

The current research has shown that priming distinct lay theories of AA can defy the common perception that AA's are fixed capability machines that can only follow pre-programmed algorithms. Because of a lack of clarity in AA's learning and decision making process among lay individuals, informing them about how AA work was shown to increases trust in them (Yeomans, Shah, Mullainathan, & Kleinberg, 2017). Thus, future research can

examine whether bringing people to understanding of the learning mechanism AA use might make them focus on "how" AA learn, reversing their mindset to the low-construal one. Future research can also examine whether the knowledge about machine learning can induce fear or a feeling of threat in interacting with AA for emotionally-laden decisions and may reveal how this may interact with other factors to influence persuasion.

AA utilized in the current research (Watson, Baxter, Alexa) varied not only in their level of technological capabilities but also in the extent of anthropomorphization. For example, IBM's artificial intelligence Watson is more advanced in its reasoning ability than other AA whereas Baxter is more anthropomorphized than other AA. Prior research has shown that the level of reasoning ability and anthropomorphization are both important to formation of attitudes toward AA (Hancock et al., 2011; Waytz, Heafner, & Epley, 2014). Despite these differences, the data showed that these differing AA were commonly imbued with low-level construal characteristics. This finding implies that people's lay theory for AA (i.e., the thought that AA are created by humans to serve and lack their own superordinate goals) is the dominant determinant of construal level when humans interact with AA.

Finally, research has shown a difference between Western and Asian cultures such that Asian individuals (e.g., Koreans, Japanese) are more open-minded than Western individuals (e.g., Americans) in embracing AA (e.g., a robot) as social agents (MacDorman, Vasudevan, & Ho, 2009). It is suggested that this cultural difference may stem from a fundamental difference in their world views. For example, a Western monotheistic view supports human dominance over other creatures based on the belief that only humans possess a soul. In contrast, many Asian cultures subscribe to the belief that not only humans, but also animals and even objects (e.g., a tree), have souls (Kitano, 2007). An interesting avenue for future research would be to examine how the two different cultural worldviews affect

mindsets towards AA and interact with other factors to impact human willingness to interact and comply with advice from AA going forward.

References

Cohen, P. R., & Levesque, H. J. (1995). Communicative actions for artificial agents. *Proceedings in Manufacturing Systems, ICMAS, 95,* 65-72.

Esteva, A., Kuprel, B., Novoa, R. A., Ko, J., Swetter, S. M., Blau, H. M., & Thrun, S. (2017). Dermatologist-level classification of skin cancer with deep neural networks. *Nature, 542*, 115-118.

Floridi, L., & Sanders, J. W. (2004). On the morality of artificial agents. *Minds and Machines, 14*, 349-379.

Franklin, S., & Graesser, A. (1996). Is it an agent, or just a program?: A taxonomy for autonomous agents. *Lecture Notes in Computer Science, 1193*, 21-36.

Fujita, K., Trope, Y., Liberman, N., & Levin-Sagi, M. (2006). Construal levels and self-control. *Journal of Personality and Social Psychology, 90*, 351-367.

Gray, H. M., Gray, K., & Wegner, D. M. (2007). Dimensions of mind perception. *Science, 315*, 619.

Han, D., Duhachek, A., & Agrawal, N. (2014). Emotions shape decisions through construal level: The case of guilt and shame. *Journal of Consumer Research, 41*, 1047-1064.

Han, D., Duhachek, A., & Agrawal, N. (2016). Coping and construal level matching drives health message effectiveness via response efficacy or self-efficacy enhancement. *Journal of Consumer Research, 43*, 429-447.

Hancock, P. A., Billings, D. R., Schaefer, K. E., Chen, J. Y., De Visser, E. J., & Parasuraman, R. (2011). A meta-analysis of factors affecting trust in human-robot interaction. *Human Factors, 53*, 517-527.

Huang, M., & Roland, T. R. (in press). Artificial Intelligence in Service. *Journal of Service Research*.

Huen, E. (2016), The world's first home robotic chef can cook over 100 meals. *Forbes*. Retrieved from https://www.forbes.com/sites/eustaciahuen/2016/10/31/the worlds-first-home-robotic-chef-can-cook-over-100-meals/#bb456347228b

Johnson, P. O., & Neyman, J. (1936). Tests of certain linear hypotheses and their application to some educational problems. *Statistical Research Memoir, 1,* 57-93.

Keller, P. A. (2006). Regulatory focus and efficacy of health messages. *Journal of Consumer Research, 33*, 109-114.

Kitano, N. (2007). Animism, rinri, modernization; the base of japanese robotics. *ICRA, 7*, 10-14.

Kozak, M. N., Marsh, A. A., & Wegner, D. M. (2006). What do I think you're doing? Action identification and mind attribution. *Journal of Personality and Social Psychology, 90*, 543-555.

Leachman, S. A., & Merlino, G. (2017). Medicine: The final frontier in cancer diagnosis. *Nature, 542,* 36-38.

Lee, A. Y., Keller, A. P., & Sternthal, B. (2010). Value from regulatory construal fit: The persuasive impact of fit between consumer goals and message concreteness. *Journal of Consumer Research, 36*, 735-747.

Liberman, N., Trope, Y., McCrea, S. M., & Sherman, S. J. (2007). The effect of level of construal on the temporal distance of activity enactment. *Journal of Experimental Social Psychology, 43*, 143-149.

MacDorman, K. F., Vasudevan, S. K., & Ho, C. C. (2009). Does Japan really have robot mania? Comparing attitudes by implicit and explicit measures. *AI & Society, 23*, 485-510.

McFarland, M. (2016). IBM uses a smartphone to help diagnose skin cancer. *CW Tech (Innovate).* Retrieved from www.cwtec.com.

Petty, R. E., & Cacioppo, J. T. (1979). Issue involvement can increase or decrease persuasion by enhancing message-relevant cognitive responses. *Journal of Personality and Social Psychology, 37*, 1915-1926.

Petty, R. E., & Cacioppo, J. T. (1984). Source factors and the elaboration likelihood model of persuasion. *Advances in Consumer Research, 11*, 668-672.

Russell, S. J., & Norvig, P. (2010). Artificial intelligence: A modern approach, 3rd Edition, Upper Saddle River, NJ: Prentice Hall.

Shandwick, W. (2016). AI-ready or not: Artificial Intelligence here we come! What consumers think and what marketers need to know. Retrieved from https://www.webershandwick.com/news/article/ai-ready-or-not-artificial-intelligence-here-we-come

Spiller, S. A., Fitzsimons, G. J., Lynch Jr, J. G., & McClelland, G. H. (2013). Spotlights, floodlights, and the magic number zero: Simple effects tests in moderated regression. *Journal of Marketing Research, 50*, 277-288.

Trope, Y., & Liberman, N. (2003). Temporal construal. *Psychological Review, 110*, 403-421.

Trope, Y., & Liberman, N. (2010). Construal-level theory of psychological distance. *Psychological Review, 117*, 440-463.

Vallacher, R. R., & Wegner, D. M. (1987). What do people think they're doing? Action identification and human behavior. *Psychological Review, 94*, 3-15.

Vallacher, R. R., & Wegner, D. M. (1989). Levels of personal agency: Individual variation in action identification. *Journal of Personality and Social Psychology, 57*, 660-671.

Ward, A. F., Olsen, A. S., & Wegner, D. M. (2013). The harm-made mind: Observing victimization augments attribution of minds to vegetative patients, robots, and the dead. *Psychological Science, 24*, 1437-1445.

Waytz, A., Heafner, J., & Epley, N. (2014). The mind in the machine: Anthropomorphism increases trust in an autonomous vehicle. *Journal of Experimental Social Psychology, 52*, 113-117.

Waytz, A., & Norton, M. I. (2014). Botsourcing and outsourcing: Robot, British, Chinese, and German workers are for thinking—not feeling—jobs. *Emotion, 14*, 434-444.

Yeomans, M., Shah, A. K., Mullainathan, S., & Kleinberg, J. (2018). *Making sense of recommendations.* Unpublished manuscript.