CURRENT RESEARCH IN SOCIAL PSYCHOLOGY

Published Online: March 28, 2025

Predicting Affective Impressions in Human-Computer Interaction

Daniel B. Shank
Grace R. Hammond
Fatima Sajjad
Missouri University of Science and Technology

ABSTRACT

If affect control theory (ACT) can accurately predict affective impressions of technology, it can be expanded to human-computer interaction. We compared ACT's predicted impressions to actual collected impressions in three studies. Predicted impressions were 5-20% less accurate for technology actors than human actors (Study 1), similarly accurate in evaluation and activity for technological actors and objects after communication behavior (Study 2a), and 17-28% more accurate for technological actors than objects after physical behaviors (Study 2b). Overall predictions were accurate for technology two-thirds to three-fourths of the time, suggesting ACT's utility for modeling human-computer interaction, though there is room for improvement.

INTRODUCTION

People interact socially with technological agents, from Twitter Bots harassing users to conversations with Siri. We ask how well ACT's impression formation predictions apply to these social technologies in human-computer interaction.

ACT (Heise 2007, Robinson and Smith-Lovin 2006) is a predictive mathematical social psychological theory that proposes people develop impressions when observing social interactions based on their cultural sentiments. It proposes that both participants and observers adjust their situational impressions to align more closely with cultural sentiments. Both these stable cultural sentiments and the situational transient impressions and measured in terms of evaluation (how good or bad something seems), potency (how weak or strong something seems), and activity (how inactive or active something seems; EPA).

ACT scholarship has previously investigated how algorithms and artificial intelligence can enhance ACT (Hoey et al. 2018, Schröder, Hoey and Rogers 2016), how ACT can be integrated into computer agents to improve interaction with people (Heise 2004, Lin et al. 2014, Morgan et al. 2019), and how people use technology to modify their own affective impressions (Lulham and Shank 2023, Shank and Lulham 2017). In contrast, here we examine how well ACT's impression formation equations predict impressions of the technology itself after a social interaction involving it. As ACT was designed to predict human impressions in human-human interaction, if it could adequately predict technology impressions in human-computer interaction the theory's scope could be expanded.

Impression formation equations predict how the EPA sentiments from an *actor behavior object* event alter the EPA impressions of the *actor, behavior*, and *object* involved in it (Heise 2007). These equations are based on data from people's perceptions of human impression change in human-human interactions. Entering EPA cultural sentiment data for any *actor behavior object* event into the equations produces prediction outcomes of the impressions for that *actor*, *behavior*, and *object*.

Previous research has found that computer agents enacting the same social behavior as humans tend to lead to different, usually less extreme, EPA impressions (Shank 2013, 2014), often based on perception of agency of the computer agent (Shank 2014). Yet, no previous research has formally tested ACT's impression formation equations as applied to technological agents. Therefore, we ask: *how well can ACT predict affective impressions of technology in human-computer interaction?* Study 1 addresses impression predictions for actor and compares the technology predictions to human predictions. Study 2 compares technological actors to objects across two behavioral domains.

STUDY 1: ACTOR IMPRESSIONS FOR TECHNOLOGY AND HUMANS

In Study 1, we address how well ACT can predict the impressions of technological actors in human-computer interaction and compare this to the benchmark of its predictions for human actors in human-human interaction. Because ACT requires cultural sentiments in order to predict impression in interactions, we first collect sentiments of well-known technology which will have the most stable and agreed-upon sentiments.

Technology Sentiments and Selection

We identified 128 well-known technologies, and 151 US adult Prolific participants (see Table A1 for exclusions and demographics) were presented with 50 at random. They rated the technologies on semantic differential scales for evaluation (anchored from "bad, awful" to "good, nice"), potency ("powerless, weak" to "powerful, strong"), and activity ("slow, quiet, inactive" to "fast, noisy, active") with slider-bars from -4 to +4. Each integer was represented with as adverb label above it such "neutral" for 0, "slightly" for ±1, "quite" for ±2, "extremely" for ±3, and "infinitely" for ±4. Between subjects, the order of the scales was randomized, and each scale's orientation was randomized. To collect only well-known technologies, the instruction asked participants to select "Skip/I'm not familiar enough to rate this concept" for any concept they did not know well, and we placed this option before the EPA scales to make it more salient.

We retained the 69 terms that were skipped less than average as the most well-known technologies. Each was rated by at least 35 participants (Table A2). The sentiments of each technology were computed by taking the average for evaluation, potency, and activity respectively, across all participants who rated it.

Human Comparisons

To compare technologies to humans, we matched each of these 69 technologies to the human identity (taken from the Smith-Lovin et al. 2019 dataset) that was the most similar in sentiments using Euclidean distance (Table A2). Several technology terms matched to the same human identity, and in that case, we chose to only keep one of those technologies – the one that we deemed most unique, specific, or well-known. Therefore, we retained 56 technologies and 56 corresponding unique human identities as stimuli (Table A2).

Study Design

Each participant rated the actor in 30 *actor behavior object* sentences, randomly chosen from the 56 technology-human pairs. For each of the 56 pairs, one of eight *actor behavior object* sentences was presented according to a two (actor identity: technology vs human) by four (behaviors) factorial design. We selected four behaviors (from Smith-Lovin et al. 2019) with differing sentiments that could reasonably be done by these technologies: *helps* (Evaluation: 3.43, Potency: 2.56, Activity: 1.62), *fails* (-2.29, -1.00, -.63), *pleases* (2.53, 1.67, .59), and *disappoints* (-2.42, -.30, -.33). The object of the *actor behavior object* sentence was "a person" (.92, .59, .43; Smith-Lovin et al. 2019) for all stimuli. Participants saw the *actor behavior object* sentence, followed by "This *actor* is …", with a skip option and EPA scales. For example, the eight factorial conditions for the *spam bot* technology are "A spam bot *helps/fails/pleases/disappoints* a person. This spam bot is..." and for the matched human "A womanizer *helps/fails/pleases/disappoints* a person. This womanizer is…"

Participants and Ratings

405 US adult Prolific participants rated the impressions with the EPA semantic differential scales with the order and orientation of the scales randomized in the same way as the sentiment data collection (see Table A1 for demographics and exclusions).

After data collection, we realized *Android* may have interpreted it as a humanlike robot or the phone brand, so we removed it and its matched human identity. Each of the remaining 440 sentences (55 technologies * 8 conditions) was rated 24.6 times and skipped 2.2 times on average. Actor impressions were calculated for each sentence by averaging the ratings for evaluation, potency, and activity, respectively, across all participants who rated it for a total of 220 human impressions and 220 technology impressions.

Results

Using ACT's impression formation equations, we predicted the 220 impressions for each of the *technological actors* in each of the *actor behavior object* events. Comparing the ACT predictions with the actual impression data, showed that 66.4% evaluation, 55.9% potency, 77.2% activity impression predictions were accurate, within 1 point.

Using ACT's impression formation equations, we also predicted the 220 impressions of each *human actor* in each *actor behavior object* event. Comparing the ACT predictions with the actual impression data, showed that 86.4% evaluation, 60.9% potency, and 87.3% activity impression predictions were accurate, within 1 point. These were significantly more accurate than the technology impression predictions for evaluation (difference: 20.0%; Chi-Squared(1, 440)=24.38, p<.001) and activity (difference: 10.1%; Chi-Squared(1, 440)=7.54, p=.006), but not for potency (difference: 5.0%; Chi-Squared(1, 440)=1.132, p=.287).

Therefore, while ACT is better at predicting impressions of human actors in human-human interaction than for technological actors in human-computer interaction, it is only by a small amount (5-20%). Compared to this human prediction benchmark, we consider this level of accuracy quite decent considering that none of the ACT equations were designed for human-computer interaction.

STUDY 2: IMPRESSIONS OF TECHNOLOGICAL ACTORS AND OBJECTS

In Study 1, we addressed how well ACT predicts the impressions of technological actors behaving towards people. When people behave toward technology, then the technology becomes the object in the *actor behavior object* event. In Study 2, we compare how well ACT can predict the impression of technology as an actor or an object in human-computer interaction. To enhance the generalizability of our results we use behaviors that systematically differ in sentiments. In order to find such behaviors that can be done both by and to technology, we conducted this study in two domains: with conversational technologies (Study 2a) and with physically moving technologies (Study 2b).

Technology and Behavior Selection

For Study 2a, we selected 8 conversational technologies (*A chatbot, A Twitter bot, A virtual assistant, An Amazon Alexa, An Amazon Echo, A Google Assistant, A Google Home,* and *Siri)* and 15 behaviors (*misunderstands, defers to, taunts, makes fun of, deceives, ignores, harasses, threatens, obeys, complies with, requests something from, understands,* and *listens to)*. For Study 2b, we selected 8 physically interacting technologies (*A delivery drone, A drone, A humanlike robot, A robot, A robotic dog, A robotic toy, A self-driving vehicle,* and *Tesla's self-driving car)* and 12 behaviors (*hides from, avoids, flees from, overlooks, monitors, hits, attacks, watches, observed,* and *examines).* In both studies, we selected two behaviors, when available, from each of the eight octants of sentiment space (i.e., ± Evaluation * ± Potency * ± Activity) that we determined could be reasonable done by or to those technologies.

Design

Study 2a/b was a 2 (technology position: actor or object) by 15/12 (behavior) within-subject design. For each of the 30/24 conditions, participants were randomly assigned to one of the 8 technologies for that study for a total of 240/196 sentences. When the technology was assigned to be the actor, the object was "a person" and when the technology was assigned to be the object the actor was "a person." The sentences were presented as *actor behavior object*, followed by "This *technology* is …" followed by the skip option and three EPA scales. For example, the behavior *misunderstands* was presented twice to all participants in Study 2a, each time with a randomly chosen technology. So once, if *Amazon Echo* was selected for the object, it would read

“A person misunderstands an Amazon Echo. This Amazon Echo is…”. The other time, if *a Twitter bot* was selected for the actor, it would read “A Twitter bot misunderstands a person. This Twitter bot is…”.

Participants and Ratings

US adult Prolific participants, 241 for Study 2a and 236 for Study 2b, rated the impressions with the EPA semantic differential scales with the order and orientation of the scales randomized similar to the previous study (see Table A1 for demographics and exclusions). For Study 3a/b, Each of the 240/196 sentence was rated 27.8/27.6 times and skipped 2.2/1.7 times on average. Technology impressions were calculated for each sentence by averaging the ratings of evaluation, potency, and activity, respectively, across all participants.

Results

We predicted the impressions (240 for Study 2a and 196 for Study 2b) for each of the *technological actors* or *technological object* in each of the *actor behavior object* events using ACT’s impression formation equations. The sentiments used for the technologies in these predictions were the technology sentiments collected in Study 1 (Table A2).

For Study 2a, comparing the ACT predictions with the actual impression data for technological actors indicates that 67.5% evaluation, 67.5% potency, and 66.7% activity impression predictions were accurate, within 1 point. Making the same comparisons for technological objects indicates that 64.2% evaluation, 95.0% potency, and 56.7% activity were accurate within 1 point. Accuracy of technological actors and objects are not statistically different for evaluation (difference: 3.3%; Chi-Squared(1, 240)=0.30, p=.586) nor activity (difference: 10.0%; Chi-Squared(1, 240)=2.54, p=.111), but object potency was statistically more accurate than actor potency (difference: 27.5%; Chi-Squared(1, 240)=29.78, p<.001).

For Study 2b, comparing the ACT predictions with the actual impression data for technological actors indicates that 81.3% evaluation, 88.5% potency, and 72.9% activity impression predictions were accurate, within 1 point. Making the same comparisons for technological objects indicates that 58.3% evaluation, 71.9% potency, and 52.1% activity were accurate within 1 point. Accuracy of technological actors was statistically greater than accuracy for technological objects for evaluation (difference: 28%; Chi-Squared (1, 192) = 11.96, p<.001), potency (difference: 16.6%; Chi-Squared (1, 192) = 8.40, p=.007) and activity (difference: 20.8%; Chi-Squared (1, 192) = 8.89, p = .004).

Therefore, in human-computer interaction that deals with physical behaviors ACT better predicts impressions of technological actors than technological objects by a moderate amount (17-28%). Yet, for conversational behavior ACT was similarly accurate at predicting evaluation and activity impressions of technological actors and objects. Yet 27.5% more potency predictions were accurate for technological objects. Given that the other accuracies in Study 2a range from 56.7-67.5%, the outlier of 95.0% accurate may be due to impressions changing less when being the recipient of a conversational action which often doesn’t weaken or strengthen one’s power.

DISCUSSION

How well do ACT's impression formation equations predict affective impressions of technology in human-computer interaction? Across three studies we showed that ACT can accurately predict impressions of technology, with some qualifications. Across EPA and the actor or object position, ACT's technology impression accuracy ranged from 52.1% to 95.0%. Average accuracy was 67.5% for evaluation, 75.8% for potency, and 65.1% for activity. It was 71.5% for technological actors and 66.4% for technological objects. Therefore, we conclude that ACT's predicted impressions of technology in human-computer interaction are accurate for roughly two-thirds to three-fourths of the time. This is comparable, albeit slightly lower, than ACT prediction accuracy for human impressions in human-human interaction in Study 1. Overall, this shows that ACT predictions of technological impressions are fairly accurate and robust to the specific type of actors and objects involved in the interaction. This justifies using ACT in future predictions and hypotheses in human-computer interaction.

Sensitivity to Accuracy Cutoff

There is no set method for assessing accuracy in the ACT literature and traditional analyses of statistical significance are not appropriate for equation-based impression predictions. For our analyses we arbitrarily selected a 1-point difference to represent prediction accuracy, as 1 point is the smallest meaningful change in impression labels (e.g., from "slightly" to "quite"; Heise 2007, 2010).

Supplemental analyses show that changing this accuracy cutoff simply shifts the results but doesn't change comparison conclusions. For example, shifting to a 2-point difference in accuracy means that ACT accurately predicts impression of over 90% of the technological actors and over 95% of human actors in Study 1. Alternatively, shifting to a 0.5-point difference reduced ACT's accuracy to about 35% for technological actors and about 45% for human actors in Study 1. Thus, the overall accuracy shifts the pattern of results when using a more conservative or liberal metric.

Generalizability

We believe these results are robust and therefore generalizable. In Study 1 we considered 55 different technologies of all different types, but only 4 behaviors. Yet, in Study 2 we narrowed down to only 16 technologies in two different behavioral domains but were able to examine 27 behaviors that were systematically selected from across EPA sentiment space. Therefore, across these studies we examined a broad range of technologies, behaviors, and events. Notability, the newest technologies like ChatGPT and Microsoft Bard were not included, although these new AI-based technologies arguably seem more humanlike and therefore may be better predicted by ACT. Future studies should examine these by building on our foundational results.

Limitations, Challenges, and Future Work

Human-computer interaction has become more social, but still has many challenges since technologies are not people.

First, digital technologies rapidly change from creation to obsolescence, and only somewhere in-between could they qualify as having stable cultural sentiments (Heise 2010). Thus, they may be less stable than human sentiments and more influenced by experiences and technological enculturation (King 2008, Shank 2010). Therefore, we focused on only well-known social technologies, first through our own selection, then through only retaining those that were known by most participants. Still, a limitation is the changing nature of technological systems and their sentiments, compared to sentiments of other terms which change more slowly over time (Heise 2010). Related, despite attempts to find technologies with sentiments across EPA space, the majority of technologies were positive on all three dimensions (Table A2). Future research could sample technologies with distributed sentiments from Table A2 or elsewhere (e.g., Lulham and Shank 2023, Shank 2010).

Second, many behaviors do not make sense to be done by or to many technologies because the technologies are limited in purpose, design, or expectations or behaviors may be interpreted differently for technology and humans. For example, *Siri* cannot *flee from* a person and *a drone* does not typically *thank* someone. Due to this, we were limited to four generic behaviors in Study 1 and were able to expand to only 27 behaviors by having two specific domains for Study 2. Future research may want to employ specific domains, like in Study 2, to ensure that there are not mismatches between technologies and behaviors.

Third, our studies all used *a person* as the human interactant identity, somewhat simplifying the situations that we examined. Using *a person* affected the impression predictions and our ability to generalize to other situations. Bad or powerful humans interacting with technology might prove easier or harder to predict. In addition to future research considering a range of human identities, our research sets the stage for an investigation of affect control theory predicting impressions of *computer-computer* social interaction.

Concluding Implications

Over the past few decades, ACT has greatly expanded in theoretical reach into new domains (MacKinnon and Robinson 2014). Herein, we have added to that by showing how the impression formation process, a major component of ACT, functions for human-computer interaction. Since the impression formation processes is sequentially prior to the control principle, we believe a next step would be to demonstrate the affect control process applied to human-computer interaction. If the control principle functions similarly for human-computer interaction, then the scope conditions could be expanded to include technology. Even so, ever-changing technologies may not have stable cultural sentiments and may not function like many human identities. Future empirical research should consider stable and less stable human and technological identities.

Integrating human-computer interaction into ACT will enhance its capacity to be integrated into social technologies (e.g., Lin et al. 2014). As social interactions with artificial intelligence, algorithms, and smart assistants become more commonplace understanding implications of that ultimately relies on the scientific understanding of those social interaction processes. ACT's modeling of impressions in human-computer interaction shows that while our social interactants might sometimes now be machines, our perceptions of them are still based on affective meaning.

References

Heise, David R. 2004. "Enculturating Agents with Expressive Role Behavior." Pp. 127-42 in *Agent Culture: Human-Agent Interaction in a Multicultural World*, edited by S. Payr and R. Trappl. Mahwah, NJ: Lawrence Erlbaum Associates.

Heise, David R. 2007. *Expressive Order: Confirming Sentiments in Social Actions*. New York: Springer.

Heise, David R. 2010. *Surveying Cultures: Discovering Shared Conceptions and Sentiments*: Wiley.

Hoey, Jesse, Tobias Schröder, Jonathan Morgan, Kimberly B Rogers, Deepak Rishi and Meiyappan Nagappan. 2018. "Artificial Intelligence and Social Simulation: Studying Group Dynamics on a Massive Scale." *Small Group Research* 49(6):647-83.

King, Adam B. 2008. "Finding Online Subcultures in Shared Meanings." *Social Science Computer Review* 26(2):137-51.

Lin, Luyuan, Stephen Czarnuch, Aarti Malhotra, Lifei Yu, Tobias Schröder and Jesse Hoey. 2014. "Affectively Aligned Cognitive Assistance Using Bayesian Affect Control Theory." Pp. 279-87 in *Ambient Assisted Living and Daily Activities*: Springer.

Lulham, Rohan and Daniel B Shank. 2023. "How "Stuff" Matters in Affect Control Theory." *American Behavioral Scientist* 67(2):270-89. doi: https://doi.org/10.1177/00027642211066045.

MacKinnon, Neil J. and Dawn T Robinson. 2014. "Back to the Future: 25 Years of Research in Affect Control Theory." *Advances in Group Processes (Advances in Group Processes, Volume 31) Emerald Group Publishing Limited* 31:139-73.

Morgan, Jonathan H, Jun Zhao, Nikolas Zöller, Andrea Sedlacek, Lena Chen, Hayley Piper, Yliana Beck, Kimberly B Rogers, Jesse Hoey and Tobias Schröder. 2019. "Modeling the Culture of Online Collaborative Groups with Affect Control Theory." Pp. 147-60 in *Conference of the European Social Simulation Association*: Springer.

Robinson, Dawn T. and Lynn Smith-Lovin. 2006. "Affect Control Theory." Pp. 137-64 in *Contemporary Social Psychological Theories*, edited by P. J. Burke. Stanford: Stanford University Press.

Schröder, Tobias, Jesse Hoey and Kimberly B. Rogers. 2016. "Modeling Dynamic Identities and Uncertainty in Social Interactions: Bayesian Affect Control Theory." *American Sociological Review* 81(4):828-55.

Shank, Daniel B. 2010. "An Affect Control Theory of Technology." *Current Research in Social Psychology* 15(10):1-13.

Shank, Daniel B. 2013. "Are Computers Good or Bad for Business? How Mediated Customer–Computer Interaction Alters Emotions, Impressions, and Patronage toward Organizations." *Computers in Human Behavior* 29(3):715-25. doi: 10.1016/j.chb.2012.11.006.

Shank, Daniel B. 2014. "Impressions of Computer and Human Agents after Interaction: Computer Identity Weakens Power but Not Goodness Impressions." *International Journal of Human-Computer Studies* 72(10-11):747-56.

Shank, Daniel B. and Rohan Lulham. 2017. "Products as Affective Modifiers of Identities." *Sociological Perspectives* 60(1):186-205. doi: 10.1177/0731121416629996.

Smith-Lovin, Lynn, Dawn T Robinson, Bryan C Cannon, Brent H. Curdy and Jonathan H. Morgan. 2019. "Mean Affective Ratings of 968 Identities, 853 Behaviors, and 660 Modifiers by Amazon Mechanical Turk Workers in 2015." University of Georgia.

Appendix

Appendix Table A1: Participant Demographics by Study

	Technology Sentiment Data	Study 1	Study 2a	Study 2b
Total Participants	153	430	267	251
Excluded for not completing a majority of the ratings	2	23	24	12
Excluded for not spending a majority of life in the US	0	2	2	3
Total Usable Participants	151	405	241	236
Mean Age [range]	25.1 [18-54]	27.4 [18-68]	35.8 [19-73]	37.5 [18-70]
Completed Some College	85.4%	89.9%	85.5%	90.2%
Gender				
Male	17.2%	14.6%	21.2%	21.2%
Female	81.5%	81.0%	74.7%	77.5%
Other, or prefer not to say	1.3%	4.4%	3.7%	1.3%
Race/Ethnicity				
White	62.5%	67.9%	73.9%	74.2%

Asian	11.2%	8.9%	6.2%	7.2%
Hispanic	11.8%	9.9%	6.2%	4.7%
Black	7.2%	5.2%	6.2%	9.7%
Other/Mixed	7.2%	8.1%	7.4%	4.2%
Region				
West	25.8%	27.0%	18.5%	21.2%
South	31.1%	32.0%	38.7%	37.7%
Midwest	23.2%	20.7%	23.9%	25.4%
Northeast	19.8%	10.5%	18.5%	15.7%

Appendix Table A2: List of 69 Terms Retained from Technology Sentiment Data with Ns (Percent Skipped) and EPA Sentiments from Technology Sentiment Data, Human Matched Identity and Euclidian Distance from the Technology Term for Study 1, and Terms used for Study 2

	Sentiment Data Study				Study 1 Human Match	
Technology Term (Used in Study 1)	**N (Skipped %)**	**Evaluation**	**Potency**	**Activity**	**Identity**	**Euclidian**
A Computer Virus	57 (6.6)	-3.20	1.88	0.60	Murderess	0.750
A Criminal Facial Recognition Algorithm	41 (29.3)	1.78	2.15	0.71	Diplomat	0.259
A Customer Service Bot	64 (12.3)	0.03	-0.26	-0.11	Pagan	0.195
A Delivery Drone[p]	42 (27.6)	1.37	1.53	1.62	building contractor	0.091
A Drone[p]	51 (10.5)	1.50	2.00	2.26	construction foreman	0.101
A Facial Recognition Algorithm	57 (1.7)	1.44	2.02	0.64	Principal	0.256

A Fingerprint Recognition Algorithm	61 (4.7)	2.13	2.49	0.47	Bodyguard	0.282
A Home Security System	60 (10.4)	2.79	2.55	1.30	Mother	0.167
A Humanlike Robot[p]	53 (20.9)	-0.18	1.50	1.63	divorce lawyer	0.365
A Music Playlist Algorithm	49 (2.0)	2.24	1.12	1.02	graduate student	0.184
A Next Word Prediction Algorithm	57 (17.4)	1.63	1.01	0.62	Christian missionary	0.156
A Robot[p]	41 (19.6)	0.69	1.91	1.32	Mayor	0.258
A Robotic Dog[p]	40 (21.6)	0.57	0.38	1.88	street preacher	0.279
A Robotic Toy[p]	50 (15.3)	1.35	0.42	1.75	Miner	0.116
A Roomba	60 (17.8)	1.97	1.17	0.90	social worker	0.229
A Self-Driving Vehicle[p]	52 (14.8)	0.81	1.80	0.96	commissioner	0.196
A Smart Phone	55 (0)	3.15	3.02	2.19	Hero	0.156
A Smart TV	55 (3.5)	2.66	2.01	1.80	Fundraiser	0.369
A Smart Watch	57 (8.1)	2.52	1.92	0.91	pediatrician	0.225
A Spam Bot	43 (27.1)	-2.54	0.42	1.66	womanizer	0.283
A Spam Detector Algorithm	47 (21.7)	1.42	1.00	0.06	church deacon	0.038
A Speech-to-Text Algorithm	45 (8.2)	2.21	1.51	0.94	foster mother	0.038
A Text-to-Speech Algorithm	54 (1.8)	1.93	1.42	1.02	black professional	0.084
A Twitter Bot[c]	46 (17.9)	-1.18	0.26	0.70	Divorcé	0.200
A Vacuum Robot	56 (11.1)	2.03	0.99	0.94	Bride	0.126
A Video Game Enemy Bot	42 (25.0)	0.67	0.91	1.47	Feminist	0.151

A Video Game NPC (Non-Player Character)	43 (25.9)	1.53	-0.33	0.00	houseguest	0.270
A Video Game Teammate Bot	39 (20.4)	0.51	0.35	0.66	Blonde	0.167
A Virtual Assistant[c]	44 (29.0)	1.80	1.60	0.99	Education administrator	0.104
Amazon Alexa[c]	57 (12.3)	1.35	1.35	1.12	Store owner	0.106
Amazon Echo[c]	43 (21.8)	1.66	1.50	1.05	Ship engineer	0.174
Amazon's Product Recommendation Algorithm	41 (16.3)	1.34	1.17	0.96	Collaborator	0.060
An Advertisement Algorithm	51 (15.0)	0.13	1.34	0.93	Critic	0.229
An Algorithm	50 (12.3)	1.20	1.98	1.04	Administrator	0.206
An Android[a]	44 (25.4)	0.85	1.44	1.03	Customs officer	0.112
An Apple Watch	48 (12.7)	2.41	2.01	1.01	Foster father	0.210
Facebook’s Friend Recommendation Algorithm	49 (14.0)	0.40	0.52	0.39	Plaintiff	0.860
Facebook's Group Recommendation Algorithm	40 (23.1)	0.77	0.38	0.47	Black man	0.117
Facebook's News Feed Algorithm	41 (29.3)	-0.22	1.07	1.43	Pawnbroker	0.382
Facebook's Recommended Post Algorithm	47 (17.5)	0.10	1.00	0.96	Bailsman	0.405
Google Assistant[c]	46 (24.6)	1.35	1.33	0.65	Lecturer	0.077
Google Home[c]	49 (19.7)	1.90	1.93	0.97	Airline pilot	0.127
Google's Search Algorithm	47 (2.1)	2.43	2.43	1.60	Parent	0.183

Instagram's Recommended Stories Algorithm	53 (11.7)	1.00	0.83	0.85	Truck driver	0.045
Instagram's Suggested Posts Algorithm	54 (10.0)	0.69	0.86	0.63	In-law	0.115
My Starbucks Barista	38 (24.0)	2.61	1.73	1.69	Elementary School teacher	0.140
Netflix's Recommendation Algorithm	58 (0)	1.76	1.31	0.65	Civil engineer	0.094
Pinterest's Recommend Posts Algorithm	45 (21.1)	1.72	1.08	0.64	Husband	0.096
Roku TV	51 (17.7)	2.36	1.55	1.33	Supporter	0.156
Sims (from The Sims)	44 (22.8)	2.12	-0.02	1.32	Mail carrier	0.357
Snapchat's Recommended Stories Algorithm	48 (12.7)	-0.05	-0.16	0.67	Unwed parent	0.193
Spotify's Playlist Algorithm	60 (9.1)	2.55	1.66	1.10	Optimist	0.048
Tesla's Self Driving Car[p]	49 (16.9)	1.13	2.04	1.25	Supervisor	0.168
Tiktok's "For You" Recommendation Algorithm	49 (9.3)	2.13	2.45	1.96	Police officer	0.320
Twitter's News Feed Algorithm	42 (22.2)	0.97	1.10	0.90	Guy	0.215
YouTube's Recommended Video Algorithm	58 (3.3)	1.56	1.22	1.12	Air Force reservist	0.250
Technology Term (Not Used in Study 1)						

A Chatbot[c]	38 (28.3)	0.09	-0.07	0.74	Unwed parent	0.211
A College Admissions Algorithm	35 (28.6)	0.14	1.10	-0.15	Informer	0.333
A GPS (Global Positioning System)	59 (0)	2.99	2.49	1.53	Mother	0.206
A Job Recruitment Algorithm	46 (27.0)	0.89	0.85	0.36	Polymath	0.109
A Mars Rover	42 (27.6)	2.87	2.83	1.42	Mother	0.405
A Personal Home Assistant	46 (24.6)	1.53	1.79	1.08	Ship engineer	0.199
A Product Recommendation Algorithm	52 (10.3)	0.82	1.13	0.63	Guy	0.125
A Self-Checkout Kiosk	59 (0)	2.64	1.79	1.89	Fundraiser	0.324
An Auto Correct Algorithm	52 (8.8)	1.75	1.39	0.97	Air Force reservist	0.069
Google Maps	60 (0)	2.57	2.11	1.10	Physician	0.208
Google Translate	58 (1.7)	2.66	2.14	1.33	Physician	0.252
Siri[c]	52 (0)	1.87	1.71	1.30	Deputy	0.188
Snapchat's Discover Page Algorithm	54 (15.6)	-0.05	-0.16	0.67	Unwed parent	0.127

[a] Removed from Study 1 due to ambiguity.
[c] Conversational technology used in Study 2a.
[p] Physical movement technology used in Study 2b.

AUTHORS' NOTE

Author Note: This research was sponsored by the Army Research Office under Grant Number W911NF-19-1-0246. The views and conclusions contained in this document are those of the authors and should not be interpreted as representing the official policies, either expressed or implied, of the Army Research Laboratory or the U.S. Government. The U.S. Government is authorized to reproduce and distribute reprints for Government purposes notwithstanding any copyright notation here on. No potential competing interest was reported by the authors. Thanks to Ruijie Liu and Angela Bucher for their contributions to the stimuli development, literature, and early versions of this paper and to Abigail Wilson, Julia Mcintyre, and Hannah Tillery for their contributions to the stimuli development on this project.

AUTHOR BIOGRAPHIES

Dr. Daniel B. Shank is an associate professor of psychological science at Missouri Science & Technology specializing in the areas of social psychology and technology. His research primarily focuses on social psychological interactions with and perceptions of artificial intelligence including morality, emotions, impression, and behavior toward AIs. Email: shankd@mst.edu

Grace R. Hammond is a Human Resources Generalist at CRRC Sifang America, where manages daily operations. She holds a Master of Science (MS) degree in Industrial-Organizational Psychology. Her research interests include human behavior as a group and the various emotional and affective factors impacting decision making and change on a societal or cultural level. Email: ch2dd@mst.edu

Fatima Sajjad is an Area Manager at Amazon, where she oversees operations and ensures efficient workflow. She holds a Master of Science (MS) degree in Industrial-Organizational Psychology. Her research interests include evaluating human behavior under various situations, exploring the intricacies of human decision-making and its impact on job performance and selection processes, and studying human interaction with AI agents. She is passionate about leveraging these insights to enhance organizational effectiveness and employee well-being. Email: fsxxb@mst.edu