

“Excuse Me, Robot”: Impact of Polite Robot Wakewords on Human-Robot Politeness

Tom Williams¹, Daniel Grollman², Mingyuan Han¹, Ryan Blake Jackson¹,
Jane Lockshin¹, Ruchen Wen¹, Zachary Nahman¹, and Qin Zhu¹

¹ Colorado School of Mines, Golden CO 80401, USA twilliams@mines.edu

² Plus One Robotics, Boulder CO USA dan.grollman@plusonerobotics.com

Abstract. While the ultimate goal of natural-language based Human-Robot Interaction (HRI) may be free-form, mixed-initiative dialogue, social robots deployed in the near future will likely primarily engage in *wakeword-driven* interaction, in which users’ commands are prefaced by a wakeword such as “Hey, Robot.” This style of interaction helps to allay user privacy concerns, as the robot’s full speech recognition module need not be employed until the target wakeword is used. Unfortunately, there are a number of concerns in the popular media surrounding this style of interaction, with consumers fearing that it is training users (in particular, children) to be rude towards technology, and by extension, rude towards other humans. In this paper, we present a study that demonstrates how an alternate style of wakeword, i.e., “Excuse me, Robot” may allay this concern, by priming users to phrase commands as Indirect Speech Acts.

Keywords: Persuasive Robotics; Indirect Speech Acts; Wakewords

1 Introduction

Voice interactive technologies are becoming increasingly common: all modern smartphones (and many personal computers) come with at least one “digital assistant” to help users perform a variety of tasks. Indeed, with the increase of internet of things (IoT) technologies, voice interfaces are being added to a wide array of home and work appliances, including refrigerators, microwaves, and even faucets [12]. But despite several decades of research into mixed initiative dialogue [1, 16] and turn taking [9, 30], the dominant paradigm in consumer-grade voice interaction is primarily human-driven, with human turns started by platform-specific *wakewords*, such as “Alexa”, “Okay Google”, or “Hey Siri”.

Wakewords, which help ensure that voice assistants only respond to genuine assistant-directed requests, and ensure user privacy, are traditionally designed for ease of automatic recognition. But, wakewords may also have social, emotional, and cognitive impact on their users. There has been significant public concern about the potential negative consequences of wakeword-driven interaction, i.e., that wakeword-driven interactions may encourage technology-directed

language that is terse and direct, and that if children become accustomed to addressing machines in this manner, this could train them to be impolite [15, 31]. Public outcry has been high enough that companies have responded by changing assistants’ interaction patterns to encourage the use of key phrases such as “Please” and “Thank You” [11, 2], a decision we will discuss in detail below.

As interactive robots are deployed into the wild, it is likely that the same concerns will arise. Indeed, when the Jibo robot was launched in 2017, it did so with the wakeword “Hey Jibo.” We believe that the concerns the public has raised about wakeword-based interaction will be especially important to address for interactive robots given their unique persuasive power.

Human norms are well known to be dynamic and malleable [14], with norms defined, communicated, and enforced by community members (and the technologies with which they interact) [32]. As researchers have recently argued, social robots wield unique influence over these norms due to their joint status as perceived community members and as technological tools [17]. This influence, which social robots may wield both through direct persuasion and implicit social pressure [5, 21, 18, 35], may be especially strong among language capable robots, with their increased linguistic faculties, and perceived agency, embodiment, anthropomorphism, and ostensible individuality, resulting in significantly greater effect on users’ systems of social and moral norms, including sociocultural norms such as norms of politeness. We argue that for robots, the *choice* of wakeword used is thus especially important.

In this paper, we examine the effect a designer’s *choice* of wakeword may have on robot-directed human politeness. In Section 2, we discuss previous attempts by digital assistant designers to encourage politeness in wakeword-based interaction, and potential limitations of those approaches. We then propose and justify an alternative approach, and delineate a set of research questions and hypotheses raised by that approach. In Sections 3 and 4, we then present the design and results of a human-subject experiment designed to evaluate those hypotheses. Finally, in Sections 5 and 6, we discuss and interpret these results, and present several potential directions for future work.

2 Wakeword Design

In response to public concern about wakeword-based child-Alexa interactions, Amazon pursued a variety of strategies to try to encourage children to speak politely with Alexa. Initially, Amazon developed a mode requiring interactants to include “please” in their requests for them to succeed³. After being told that this approach would likely backfire, Amazon shifted to simply praise interactants

³ Cf. the work of Bonfert [4], who show that rebuking adult Alexa users for not using “please” does indeed lead users to use “please” more frequently (likely to avoid the annoyance of the rebuke) but also causes users to like the assistant less and view it as less inherently entitled to politeness.

for using the words “please” and “thank you” [3]. While this may be effective at encouraging some users to use the word “please,” we suspect it may be far less effective at encouraging users to *be polite*. While saying please is indeed often used as a politeness strategy, adherence to different types of politeness norms is highly context-sensitive [23], and the type of please-usage encouraged by this approach, in which requests are preceded by “Please” (e.g., “Hey Alexa, *Please* play Todd the T-1000”), is actually *negatively-correlated* with politeness [10], as it is most naturally followed by a command. In fact, sentence-medial please usage is typically perceived as polite in part because it typically augments sentences that are already polite for other reasons, such as the use of so-called “indirect speech acts” [28] (e.g., “Could you *X*,” whose literal meaning (in this case, a yes-or-no question pertaining to ability) mismatches its intended meaning (in this case, a request for action)).

An Alternate Approach

To design wakeword-based human-robot interactions that counteract the potential tendency towards impoliteness, rather than simply encouraging the use of “please”, a more promising approach might be to change the wakeword itself, in a way that encourages deeper politeness strategies such as indirect speech act usage. Consider the simple change in wakeword from ‘Hey’ to ‘Excuse me.’

First, while “Hey <Name>” – especially when followed by ‘please’ – may syntactically prime the speaker to continue their utterance with a direct phrasing, “Excuse me <Name>” may instead prime the speaker to continue their utterance with an indirect phrasing. We argue that in Example 1 below, the impolite phrasing (1a) is slightly more syntactically natural, while in Example 2, the polite phrasing (2b) is significantly more natural.

- (1) a. Hey Pepper, please bring me a coffee.
b. Hey Pepper, please could you bring me a coffee?
- (2) a. Excuse me, Pepper, bring me a coffee.
b. Excuse me, Pepper, could you bring me a coffee?

Second, it may be easier to prime users to use indirect speech acts than arbitrary keywords such as “please,” as humans automatically tend toward indirect speech act use, especially in contexts with highly conventionalized sociocultural politeness norms [33]. Third, “Excuse me” as a wakeword is advantageous as it simply changes the wakeword participants need to use, without adding any additional requirements. Fourth, feedback delivered after “please”-use may eventually be perceived as annoying, thus reducing its expected usage. In contrast, wakeword alteration does not require any robotic feedback. Finally, if this choice of wakeword is indeed effective in priming indirect speech act use, this will lead to a more productive opportunity for “ritualization” than would successful priming of “Please”-usage. Rituals are critical ways through which a community maintains

its values and cultivates these values in community members [25], and human acts involving appropriate performance of rituals can be viewed as manifesting cultivated moral selves [22]. Using the wakeword “Excuse me” opens opportunities that “prompt” human teammates to initiate and participate in conversations and interactions that are guided by rituals, whereas “Please” may in fact invite requests that are explicitly demanding.

Hypotheses

In this paper, we begin to examine these intuitions by exploring the efficacy of “Excuse me” relative to the standard impolite wakeword baseline “Hey”. Specifically, we test the following concrete hypotheses:

Hypothesis One: Requiring the use of a polite wakeword (e.g., “Excuse me”) will result in increased robot-directed politeness.

Hypothesis Two: Observed differences in robot-directed politeness will be due to wakeword-driven linguistic priming rather than wakeword-driven differences in perceptions of robots.

3 Methods

To investigate our hypotheses, we conducted a two-condition between-subjects laboratory experiment. In this experiment, human participants collaborated with a fully autonomous robot and a human confederate in a simulated restaurant scenario, with the wakeword used to initiate communication altered between conditions. The experimental setting and robot (SoftBank’s Pepper) are shown in Fig. 1.

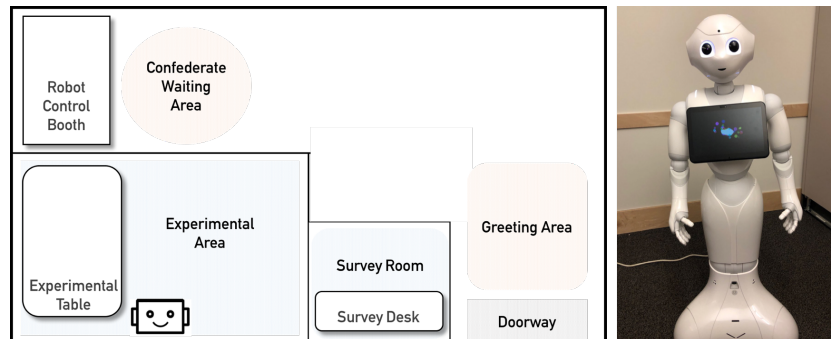


Fig. 1: The experimental environment layout (left) and the Pepper robot (right).

3.1 Experimental Design

Upon arriving at our lab and providing informed consent, participants were given a headset to facilitate audio recording, and were instructed that they would be giving four food orders to a robot designed to help waiters in a simulated restaurant scenario. We chose this context due to its conventions surrounding the use of politeness and directness in task-oriented speech [33]. Every robot-directed utterance during the experiment was required to begin with the robot’s wakeword. In the impolite wakeword condition, the wakeword was “Hey Pepper”, while in the polite wakeword condition, the robot’s wakeword was “Excuse me Pepper”. Our hypotheses predict that robot-directed speech (aside from the wakeword) would be more polite in the polite wakeword condition.

The participant was then given four cards, each containing a food item (e.g., “cheeseburger”) and a table number (1 through 4), and was informed as to the experimental procedure. When they were clear on this procedure, the experimenter activated the robot by saying “{Hey / Excuse me} Pepper we are ready to begin”. The robot then animated and said “Great! Hi there. I’m Pepper. I’m ready to take the first order.” The participant then told each food order to the robot with an utterance like “{Hey / Excuse me} Pepper can I have a chicken quesadilla please?” The robot ostensibly sent the order to the “kitchen”. The human confederate then entered with a card representing the food and said “Alright, I’ve got the order for you.” The participant then told the confederate on which table to place the order. After all four orders, the participant was brought to an adjacent room, where they were given the Robotic Social Attributes Scale (RoSAS) [8], which measures warmth, competence, and discomfort, as well as a brief demographic survey. Finally, participants were paid and debriefed.

3.2 Participants

30 student participants (17 male, 12 female, 1 nonbinary) were recruited from our university campus, and were randomly assigned to our two conditions, resulting in 15 participants per condition. Participant ages ranged from 18 to 36 years ($M=21.30$, $SD=3.91$). This small sample size was in part due to the fact that this experiment was run as part of a novel Experimental Ethics curricular module in Mines’ Spring 2019 *Robot Ethics* class, co-taught by authors Williams and Zhu, the pedagogical implications of which are detailed in a recent paper [34].

12 participants reported previous experience with robots, and 9 reported previous experience in the restaurant industry. All participants reported majoring in fields under the umbrella of science, technology, engineering, and mathematics (STEM), with the two most popular majors being computer science (9 participants) and mechanical engineering (5 participants). Other represented majors with 3 or fewer participants include: chemical engineering, civil engineering, electrical engineering, environmental engineering, geological engineering, geophysics, materials science, metallurgy and materials engineering, petroleum engineering,

physics, and statistics and applied math. Participants were paid \$5 each for participation.

3.3 Data Annotation

Three annotators marked each participant utterance for the presence of the common politeness markers delineated by Danescu-Niculescu-Mizil [10] and described below (see also [6]). Cases where all three annotators disagreed were resolved by a vote from an additional annotator. The first of these markers, and our main marker of interest, is whether the utterance was a conventionally indirect speech act, i.e., an utterance whose surface form does not match its underlying intent. For example, the command “Bring me a cheeseburger” could be phrased indirectly and more politely as a question like “Could you bring me a cheeseburger?” (Fleiss’ κ for inter-annotator agreement = 0.54). Other common politeness markers which we annotate in both direct and indirect speech are deference signifiers (e.g., “nice work” or “good job”, $\kappa = 0.32$), gratitude signifiers (e.g., “thank you”, $\kappa = 0.88$), apologizing (e.g., “Sorry to bother you...”, $\kappa = 1.0$), and use of the word “please” ($\kappa = 0.89$).

3.4 Analysis

We analyzed our experimental data under a Bayesian statistical framework using the BRMS software package for Bayesian multilevel models [7] and the JASP statistical analysis software [20]⁴. We report (1) Bayes factors (BFs) indicating the relative likelihoods of our data given our experimental and null hypotheses (expressed as odds ratios), with interpretations of the strengths of these ratios based on common conventions [19], and (2) credible intervals (CIs) on the posterior probability distributions for our metrics of interest.

Hypotheses One

For Hypothesis One, the following null and alternative binomial models were defined in BRMS.

Null Model

$$Politeness_i \sim Binomial(1, p_i)$$

$$logit(p_i) = \alpha_{actor_i}$$

That is, appearance of politeness signifiers in each trial depend only on the per-participant intercept for the participant involved in that trial.

Alternative Model

$$Politeness_i \sim Binomial(1, p_i)$$

⁴ Data is available at <https://osf.io/c5hxm/>.

$$\begin{aligned} \text{logit}(p_i) &= \alpha_{actor_i} + \beta \text{Wakeword}_i \\ \beta &\sim \text{student_t}(5, 0, 2.5) \end{aligned}$$

That is, appearance of politeness signifiers in each trial depends both on the per-participant intercept for the participant involved in that trial, and on the wakeword used in that condition, with a Student’s t distribution centered on 0 with 5 degrees of freedom and scale 2.5 used as the prior distribution on the wakeword indicator variable’s coefficient β . This prior distribution was chosen based on best-practices recommendations from the research literature [13].

To evaluate Hypothesis One, these models were fit using the subset of data corresponding with robot-directed utterances, and compared using a *Bayes Factor analysis*, in which the amount of evidence for the alternate hypothesis relative to the null hypothesis is quantified as the probability of generating the observed data under the alternate model, divided by the probability of generating the observed data under the null model [27].

3.5 Hypothesis Two

To evaluate Hypothesis Two, RoSAS scores[8] were summed across each factor, after which a Bayesian independent sample t-test was performed using JASP, with wakeword condition as a grouping variable.

4 Results

Hypothesis One

In this section, we report the results of the Bayesian binomial regression used to fit the null and alternative models specified for Hypothesis One. Separate models were fit and analyzed for each of our politeness markers of interest.

Indirect Speech Act Usage – According to the best posterior fit for the alternative model, participants in the impolite wakeword condition were less likely to use robot-directed indirect speech acts than were participants in the polite wakeword condition (mean robot-directed ISA frequency in the impolite wakeword condition = 2.33 (SD=2.02); frequency in the polite wakeword condition = 3.73 (SD=1.03); $\beta = -3.54$; 95% CI=[-9.63, 1.08]). Comparison to the null model resulted in a Bayes Factor (BF) of 2.49, indicating weak evidence in favor of our alternative hypothesis: the data observed are about two-and-a-half times more likely under our alternative (condition-sensitive) model than under the null (condition-insensitive) model.

Deference – More data is needed to fully understand the effect of wakeword choice on robot-directed deference (mean robot-directed frequency of deference in the impolite wakeword condition = 0.0; frequency in the polite wakeword condition = 0.07 (SD=0.26); $\beta = -1.24$; 95% CI=[-7.17, 3.27]; BF=0.93).

Gratitude – More data is needed to assess the effect of wakeword choice on robot-directed gratitude (mean robot-directed frequency of gratitude in the impolite wakeword condition = 0.0; frequency in the polite wakeword condition = 0.07 (SD=0.26); $\beta = -0.04$, 95% CI=[-6.21, 6.05]; BF=1.04).

Apologizing – No instances of robot-directed apologizing were observed.

Please – More data is needed to assess the effect of wakeword choice on robot-directed please-usage (mean robot-directed frequency of please-usage in the impolite wakeword condition = 0.20 (SD=0.77); frequency in the polite wakeword condition = 0.27 (SD=0.59); $\beta = -0.18$; 95% CI=[-5.04, 4.36]; BF=0.81).

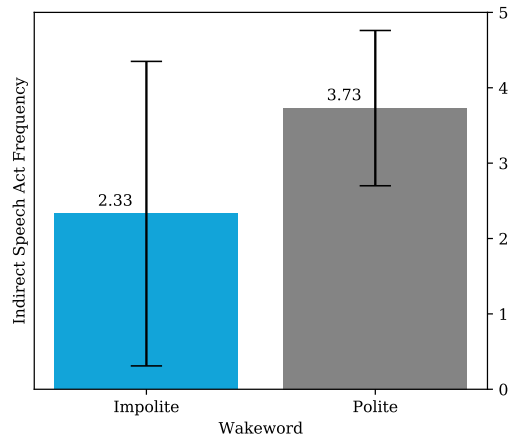


Fig. 2: Robot-directed ISA use after impolite and polite wakewords.

Hypothesis Two

A Bayesian Independent Samples t-test provided evidence against a difference between the two conditions on perceived competence ($\mu_I = 38.20$ ($SD_I = 9.017$), $\mu_P = 38.33$ ($SD_P = 9.817$); $BF = 0.345$), and inconclusive results with respect to Warmth ($\mu_I = 26.93$ ($SD_I = 9.285$), $\mu_P = 22.60$ ($SD_P = 7.790$); $BF = 0.705$) and Discomfort ($\mu_I = 12.80$ ($SD_I = 3.707$), $\mu_P = 14.73$ ($SD_P = 3.011$); $BF = 0.859$), as shown in Fig. 3. Overall this provides anecdotal to moderate evidence in support of Hypothesis Two, allowing us to tentatively rule out effects of wakeword design on human perceptions of robots as an explanation for our findings with respect to Hypothesis One. In fact, the weak differences observed in this analysis trend in favor of *decreased* warmth and *increased* discomfort in the polite wakeword condition, which is the opposite of what would have been expected in order to provide an alternative explanation for our results.

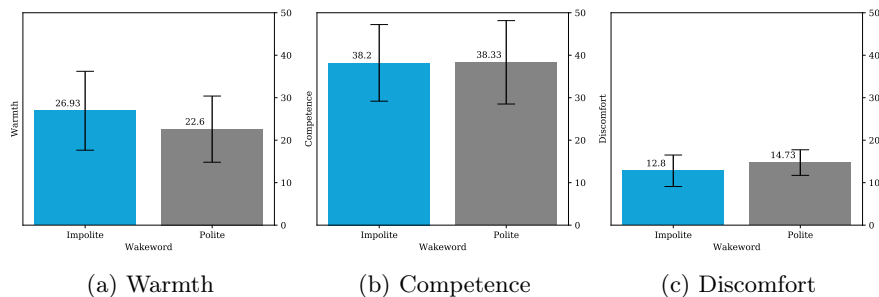


Fig. 3: Differences in Summed ROSAS Scores by Experimental Condition

5 Discussion

Our results provided weak support for our first hypothesis (H1): participants may indeed have used significantly more robot-directed indirect speech acts when instructed to use the polite wakeword than when instructed to use the impolite wakeword.

Our analyses also suggest that wakeword choice had no discernible effect on participants’ perceptions of the robot’s social attributes, thus reducing the probability that the observed differences in robot-directed politeness were due to different wakeword-induced perceptions of the robot itself (H2).

Limitations — While our subjective metrics allowed us to rule out explanations of our results grounded in wakeword-induced changes in perception of robots’ social *attributes*, our results could potentially have been due to wakeword-induced changes in perception of robots’ social *agency* (cp. [24]), which could impact the robot’s persuasive capability [26]. While we expect that a robot’s perceived sociality and perceived social attributes would likely be tightly correlated, direct examination of perceived sociality may be worthwhile.

Overall Design Recommendation Based on the results discussed above, and subject to our identified limitations, we recommend that if designers of social robots wish to encourage polite robot-directed speech, they should use politeness-priming wakewords such as “Excuse me, Robot” rather than traditional wakewords such as “Hey, Robot”.

6 Conclusions

Our results suggest that by using a polite wakeword (e.g., “Excuse me, Robot”), robot designers may be able to prime users to interact more politely with their robots overall. The most critical next step building on this experiment will be to correct the limitations we have identified, in order to better study whether this priming may actually carry over into human-human interactions as well.

In addition, a number of other key open questions may also be addressed in future work. First, there are a number of open questions regarding the real-world usability of polite wakewords: Whether users think that polite wakewords are appropriate, and whether the robot is perceived as deserving the associated politeness. If not, do polite wakewords lead to any decrease in use? And will people use polite wakewords voluntarily if they know that they are an option?

Second, there are open questions regarding how our results might differ based on the nature of the robot: How might the robot’s gender presentation impact carryover of politeness into human-human interactions, given previous findings of differences in persuasive capability for differently gendered robots [29]? And how might our results have differed if a digital assistant such as Siri or Alexa had been used instead of a humanoid robot?

Third, there are open questions regarding how our results might differ with different user populations. There have been significant concerns in the media that direct wakewords may be teaching children to be impolite. Future work will thus be needed to determine whether the effects found in this research actually differ when study participants are drawn from a child population.

Fourth, while we hypothesize that the use of a wakeword such as “Excuse me” would be more effective than encouragement to use “Please”, which as we discuss in the introduction is not actually always positively correlated with politeness, empirical evidence may be needed to strongly argue this claim, as it is possible that encouraging users to use “Please” might actually lead to increased use of “Please” in the context of indirect speech acts, in which Please-usage is indeed correlated with perceived politeness.

Finally, while we examined usage of each of a set of politeness markers in isolation, it may be worth re-analyzing our data, or data from future experiments, using a holistic measure of overall politeness that takes into account usage of each of the individual politeness markers examined in this paper.

Acknowledgments

This work was funded in part by Grants IIS-1849348 and IIS-1909847 from the National Science Foundation. We would also like to thank the students of the Spring 2019 *Robot Ethics* class at the Colorado School of Mines for their help conducting this experiment.

References

1. Allen, J.E., Guinn, C.I., Horvitz, E.: Mixed-initiative interaction. *IEEE Intelligent Systems and their Applications* **14**(5), 14–23 (1999)
2. Baig, E.: Kids were being rude to alexa, so amazon updated it. <https://www.usatoday.com/story/tech/columnist/baig/2018/04/25/amazon-echo-dot-kids-alexa-thanks-them-saying-please/547911002/> (2018)

3. BBC: Amazon Alexa to reward kids who say: ‘please’. <https://www.bbc.com/news/technology-43897516> (2018)
4. Bonfert, M., Spliethöver, M., Arzaroli, R., Lange, M., Hanci, M., Porzel, R.: If you ask nicely: a digital assistant rebuking impolite voice commands. In: Proceedings of the 2018 on International Conference on Multimodal Interaction. pp. 95–102. ACM (2018)
5. Briggs, G., Scheutz, M.: How robots can affect human behavior: Investigating the effects of robotic displays of protest and distress. *Int. J. Soc. Rob.* (2014)
6. Brown, P., Levinson, S.: *Politeness: Some Universals in Language Usage*. Cambridge University Press (1987)
7. Bürkner, P.C., et al.: brms: An r package for bayesian multilevel models using stan. *Journal of Statistical Software* **80**(1), 1–28 (2017)
8. Carpinella, C.M., Wyman, A.B., Perez, M.A., Stroessner, S.J.: The robotic social attributes scale (ROSAS): Development and validation. In: Proc. HRI (2017)
9. Cassell, J., Torres, O.E., Prevost, S.: Turn taking versus discourse structure. In: *Machine conversations*, pp. 143–153. Springer (1999)
10. Danescu-Niculescu-Mizil, C., Sudhof, M., Jurafsky, D., Leskovec, J., Potts, C.: A computational approach to politeness with application to social factors. In: Proc. ACL (2013)
11. Elgin, M.: The case against teaching kids to be polite to alexa. <https://www.fastcompany.com/40588020/the-case-against-teaching-kids-to-be-polite-to-alexa> (2018), accessed: 2019-06-26
12. Faucet, D.: Voice faucet. <https://www.deltafaucet.com/Voice> (2019)
13. Ghosh, J., Li, Y., Mitra, R.: On the use of cauchy prior distributions for bayesian logistic regression. *Bayesian Analysis* **13** (07 2015)
14. Gino, F.: Understanding ordinary unethical behavior: Why people who value morality act immorally. *Current opinion in behavioral sciences* **3**, 107–111 (2015)
15. Gordon, K.: Alexa and the age of casual rudeness. <https://www.theatlantic.com/family/archive/2018/04/alexa-manners-smart-speakers-command/558653/> (2018), accessed: 2019-06-26
16. Horvitz, E.: Principles of mixed-initiative user interfaces. In: Proc. CHI (1999)
17. Jackson, R.B., Wililams, T.: On perceived social and moral agency in natural language capable robots. In: HRI WS on The Dark Side of HRI (2019)
18. Jackson, R.B., Williams, T.: Language-capable robots may inadvertently weaken human moral norms. In: Proc. HRI. pp. 401–410. IEEE (2019)
19. Jarosz, A.F., Wiley, J.: What are the odds? a practical guide to computing and reporting bayes factors. *The Journal of Problem Solving* **7** (2014)
20. JASP Team, et al.: Jasp. Version 0.8. 0.0. software (2016)
21. Kennedy, J., Baxter, P., Belpaeme, T.: Children comply with a robot’s indirect requests. In: Proc. HRI. pp. 198–199. ACM (2014)
22. Lai, K.L.: Confucian moral cultivation: Some parallels with musical training. In: *The Moral Circle and the Self: Chinese and Western Approaches*. Open Court (2003)
23. Lockshin, J., Williams, T.: “we need to start thinking ahead”: The impact of social context on linguistic norm adherence. In: Proceedings of the 42nd Annual Meeting of the Cognitive Science Society (CogSci) (2020)
24. Mayer, R.E., Sobko, K., Mautone, P.D.: Social cues in multimedia learning: Role of speaker’s voice. *Journal of educational Psychology* **95**(2), 419 (2003)
25. Puett, M., Gross-Loh, C.: *The path: What Chinese philosophers can teach us about the good life*. Simon and Schuster (2016)

26. Roubroeks, M., Ham, J., Midden, C.: When artificial social agents try to persuade people: The role of social agency on the occurrence of psychological reactance. *International Journal of Social Robotics* **3**(2), 155–165 (2011)
27. Rouder, J.N., Speckman, P.L., Sun, D., Morey, R.D., Iverson, G.: Bayesian t tests for accepting and rejecting the null hypothesis. *Psy. Bul. & Rev.* (2009)
28. Searle, J.R.: Indirect speech acts. *Syntax and Semantics* **3**, 59–82 (1975)
29. Siegel, M., Breazeal, C., Norton, M.I.: Persuasive robotics: The influence of robot gender on human behavior. In: *Proc. IROS. IEEE* (2009)
30. Traum, D., Rickel, J.: Embodied agents for multi-party dialogue in immersive virtual worlds. In: *Proc. AAMAS*. pp. 766–773. *ACM* (2002)
31. Truong, A.: Parents are worried the amazon echo is conditioning their kids to be rude. <https://qz.com/701521/parents-are-worried-the-amazon-echo-is-conditioning-their-kids-to-be-rude/> (2016)
32. Verbeek, P.P.: *Moralizing technology: Understanding and designing the morality of things*. University of Chicago Press (2011)
33. Williams, T., Thames, D., Novakoff, J., Scheutz, M.: “Thank you for sharing that interesting fact!”: Effects of capability and context on indirect speech act use in task-based human-robot dialogue. In: *Proc. HRI* (2018)
34. Williams, T., Zhu, Q., Grollman, D.H.: An experimental ethics approach to robot ethics education. In: *Proceedings of the 10th Symposium on Educational Advances in Artificial Intelligence (EAAI)*. pp. 13428–13435 (2020)
35. Winkle, K., Lemaignan, S., Caleb-Solly, P., Leonards, U., Turton, A., Bremner, P.: Effective persuasion strategies for socially assistive robots. In: *Proc. HRI* (2019)