Radboud University
Faculty of Social Sciences
Artificial Intelligence

Can Robots Gain Appreciation by Mimicking Moral Values?

Author:
Luc Jongen
4159209

Supervisor:
W.F.G. Haselager

30 June, 2019

Radboud University

Abstract

If a robot imitates the moral values of the person it is interacting with, can it influence the way it is perceived by that person? This thesis attempts to answer that question in terms of trust in the robot, and likability and perceived intelligence of the robot. To get a concrete reading of a person's morality, the Moral Foundations Theory is used. The Moral Foundations Theory identifies several pillars in moral judgment: care/harm, fairness/cheating, loyalty/betrayal, authority/subversion and sactity/degradation. This thesis will limit its scope to the loyalty/betrayal foundation (Also called Ingroup). An experiment was done in which participants were asked to complete a survey to ascertain their reliance on the Ingroup foundation. They were then asked to talk to a Nao robot, which described a scenario, and followed up by making a decision to help or betray its ingroup, depending on the experimental condition. Afterwards the participants evaluated the robot on trust, likability, and perceived intelligence. No significant results were found but some interesting suggestions could be made to improve similar research in the future.

Contents

Chapter 1

Introduction

Trust in robots can be calibrated falsely, leading to unjustified trust, which can be dangerous, or to unjustified mistrust, which means that robots cannot fully live up to their potential usefulness. If self-driving cars reach a sufficiently high capability, they will still need to earn the passengers' trust, in order for them to be at ease inside one. This could be attempted by communicating in some way the intentions behind its actions, e.g: "I am braking because there has been a crash ahead". Knowing the reason why a robot or autonomous agent is doing what it is doing would make it less inscrutable, and will make it clear that the agent's actions are intentional, and not due to some kind of error(de Graaf & Malle, 2017).

As robots proliferate, they will inevitably enter the domain of moral judgment. Where moral decision-making is concerned, confidence in a robot's technical capability might not be enough. One might have full confidence in the robot's ability to carry out its intentions, but what if its intentions don't match one's own? To gain a person's trust in this regard, would it help to adapt to the moral values of the person it is directly communicating with? This is part of the broader question this paper will investigate: if a robot imitates the moral values of the person it is interacting with, can it influence the way it is perceived by that person? First, in order to judge a person's morality, a map is needed on which to pinpoint the person's location in a 'moral space'. The Moral Foundations Theory(Haidt et al., 2013) allows us to judge morality in several dimensions, including care, fairness, loyalty, authority, and sanctity. This thesis will focus exclusively on the Ingroup foundation (loyalty). An experiment will be performed in which a robot makes decisions showing either high or low loyalty (or high or low Ingroup endorsement). Each participant answers a questionnaire beforehand to determine how much they rely on the Ingroup foundation for their moral judgments. The participants judge the robot after it makes a moral decision in order to find out if matching a person's morality (in this case matching Ingroup endorsement) has any effect on how the participants judge the robot. The participants will judge the robot in terms of trust, likability, and perceived intelligence.

In the next chapters, the Moral Foundations Theory will be briefly explained, as it is a central concept to this thesis. Then an exploration of the meaning of trust, and the interpretation of trust in this thesis will follow. After that the concept and setup of an experiment intended to discover the effect of matching morality on trust will be explained, followed by an analysis and discussion of the results.

Chapter 2

Background

2.1 Moral Foundations Theory

If the goal is to identify and imitate morality in specific and meaningful ways, then we must know something about how humans make moral judgments. Some argue that morality can be reduced to a single dimension. (Harris, 2010), for example, argues that all moral decisions ultimately concern well-being. Maybe this is true, and moral decisions should be approached with regards to how they affect well-being. Maybe they all ability to learn morality evolved because of its effect on well being. This does not mean actual moral behavior itself is motivated by maximizing well being, just like potential partners are chosen based on beauty instead of fertility: even though the reason certain body types are considered attractive may well be because those who selected fertile partners were more likely to spread their genes, the person selecting a beautiful partner is not aware of this. (Haidt et al., 2013) have instead identified a number of Moral Foundations in their theory, which can be seen as dimensions which are used in moral thinking. The foundations may have originally developed because they generated a response based certain triggers, which made them evolutionarily advantageous (e.g. "I need to be loyal to my tribe so we win the tribe war"). However, in the current day, they may be triggered by entirely different stimuli(for example, "I prefer Nike and will mock those who choose Adidas").

The Moral Foundations Theory assumes people make moral judgments intuitively, and make up complex explanations for them later. This means that the Foundations can be seen as values that underlie moral decisions, and people may or may not be consciously aware of these. If people disagree on moral questions, this may be because the moral foundations carry different weights in each of the people making a judgment. These are the foundations currently in the Moral Foundations Theory:

Care/harm To care for and protect others from harm is considered important by many. This may be the most straightforward of the foundations. It probably evolved because mammals get relatively few, and slow to mature offspring, which needs te be cared for in order for it to reach adulthood.

Fairness/cheating Many an argument has broken out over what exactly is "fair", but the relation between cooperation and selfishness plays a role in it somehow. We will probably help someone if they are likely to repay the favor, but if a person takes and doesn't give in return, they will soon find themselves without help from others. Reciprocating the behavior of both of these kinds of people would be a simple way of acting according to a sense of fairness.

Loyalty/betrayal (also called Ingroup) Humans are tribal in nature. They quickly form groups and then will go to great lengths to prove our loyalty, such as dressing up in the color of their favorite soccer team, sing songs relaying their love that team, and possibly even destroy property belonging to supporters of a different team. All this is done even though very little depends on the success of that team. Likely this

foundation exists because humans who act in tribes were more likely to survive than those who acted alone. Like the other foundations, the extent to which Loyalty plays a role in moral judgments differs from person to person. It tends to play a bigger role for those on the political right than those on the left.

Authority/subversion This foundation concerns social hierarchy, and the respect that people have for those in different positions in that hierarchy. This hierarchy need not be a relationship between a dictator and his servants, instead it can be much like that of parent and child; it is not inherently exploitative. In English speaking countries, people of superior age are often addressed with "Mr." accompanied by their last name, instead of their first name, as one would use when addressing someone of a similar, or younger age. Adherence to or violations of this norm will be judged on the authority/subversion dimension, according to (Haidt, 2012b).

Sanctity/degradation Judgments that acts are morally wrong because they induce a sense of disgust are based on the sanctity/ degradation foundation. Opposite disgust on the sanctity/degradation spectrum, it is possible for people to consider something sacred. Disagreement on certain issues where the two sides of the political spectrum often clash can be explained using this dimension. Assisted euthanasia, for example, is likely considered a good thing by someone relying only on the Care and Fairness foundations, and they may be perplexed by the objections of people who think a life is something sacred which should never willingly be taken.

In a study by Graham, Haidt, and Nosek, 2009, the level of endorsement and use of the moral foundations differs between liberals and conservatives. Liberals highly value the Care and Fairness foundations, whereas conservatives relied on all five foundations. In this example, reliance on different moral foundations explains differences in moral concerns between conservatives and liberals, and possibly even explains why it is so hard for both groups to understand each other(Haidt, 2012a).

2.2 Trust

(Schaefer, 2013) Presents an extensive compilation of trust measures in different settings, with a plethora of robots. The robots in question range from military robots, to Roombas and Paro. The focus of the experiments ranged from teamwork to judging images of robots. The Godspeed Questionnaire (Bartneck, Kulic, & Croft, 2009) (which will be used for other criteria of the robot in this research) includes a section on perceived robot safety as well, which relates to trust. What these trust measures all have in common, is that they concern a robot's functionality performing specific tasks. Even in this meta analysis (Hancock et al., 2011) all the robot-related factors investigated were either performance- or attribute-based [1]. But what if robots enter a domain where decisions need to be made that have a moral aspect to them. Some (van Wynsberghe & Robbins, 2019) might argue that robots need not be moral agents, but completely avoiding the moral domain could be difficult. When a self-driving car recognizes a situation in which a crash is unavoidable, for instance, and now has to decide which out of two possible vehicles it hits, then this artificial agent is faced with a moral choice (Lin, 2016). The use of artificial agents could get restricted to roles in which they would not have to make any moral decisions. However, such restrictions would also prohibit self driving cars, but things seems to be moving forward steadily in that area. How would we measure the trust of a passenger in their automated vehicles to make the right choice in case of an inevitable crash? The same performance-related trust measures mentioned earlier will not do. Where moral decisions are concerned, we are dealing with different kind of trust.

Whenever trust is mentioned in the rest of this thesis, the intended type of trust is that which relates to intention, not performance. In order to measure the correct kind of trust two aspects will be focused on:

1. A questionnaire which contains questions referring to intentional trust
2. Experimental manipulation of moral decisions, while keeping robot performance constant

For trust measures which relate to intention, inspiration will be sought in the field of moral psychology. Both the questionnaire and the setup of the experiment will be explained in the method chapter.

[1] One of these attributes was personality, and therefore could technically have included morality, but its effect size was tiny.

Chapter 3

Implementation

In the experiment described in the next chapter, participants will take part in a scenario with a Nao robot. The nao robot is an about 60cm tall anthropomorphic robot which can walk, talk, and listen. It can be controlled using software made by the creators of the robot, named Choregraphe. At first, the script was modeled in Choregraphe, but because of technical difficulties python, and the Naoqi api was chosen instead.

3.1 Nao Modules

The Nao is controlled via different modules, which run on the Nao itself. There are modules for speech, posture, animation, memory, etc. When giving simple commands, it is sufficient to connect to one of these modules via a proxy, but more complicated tasks, like probing the Nao's memory, require a derivation of the general ALModule class (which all of the Nao's modules derive from). In our experiment, the Nao had to move, speak, and listen, and listening requires the python program to subscribe to a memory event so the creation of a custom module was necessary. Since all modules tend to run in parallel, all the communication between our python program and the robot was routed through a single custom module: VoiceListenerModule. This was possible because the program was simple enough and thus did not require a very large module, and was desirable because this made it simple to coordinate the function of the other modules which were used: ALAnimatedSpeech, ALMemory, and ALSpeechRecognition. Since these modules run autonomously on the robot, the python program will not wait for given commands to be executed, and typically neither will the modules wait for each other (with some exceptions, animations will not be played while changing posture, for example). This means that the speech recognition will run at the same time as the execution of dialogue, leading to the robot hearing and reacting to itself, unless this is explicitly prevented. Preventing it is what the custom module does of course, by making the python program wait until a word is heard after the command is given to listen, and by shutting down the voice recognition when a word has been recognized. The recognized word can be accessed by subscribing to a memory event (in the ALMemory module), which triggers whenever a word is recognized. The program worked most reliably when unsubscribing from the event after a word has been recognized, and resubscribing when it is time to listen for a new word.

The VoiceListenerModule also contains two functions used by the main python conversation script: say(message), and ask_question(question). The say(message) function, just makes the robot say the given message, and then makes the program wait for the robot to finish saying it. The ask_question(question) function does the same, except after reading the message it listens for a response. When a word is recognized it returns either 1 or 2 (corresponding to yes and no, since these are the only words it needs to recognize in this experiment).

3.2 The Main Script

The main conversation script works by storing lines of text in a linked list-like structure called conversation items. There are two types of these items: SayItem, and ChoiceItem. SayItem works by storing a string, and a reference to a next item, which may either be a SayItem or a ChoiceItem. The ChoiceItem consists of a string in the form of a question, and two references to alternative next items, one for each possible answer to the question. When the conversation is run by the converse(item) function, each item is executed one by one, until we have worked our way to a dead end in the linked list-like structure. How each item is executed is determined by the execute() function, contained within both types of item classes. The full code can be seen in the appendix.

Chapter 4

Method

4.1 Experiment

If a robot imitates the moral values of the person it is interacting with, can it influence the way it is perceived by that person? The following experiment was set up to attempt to at least partly answer this question. Since "The way it is perceived" seems rather general, a more specific definition of what is being tested is in order. The specific factors tested here are trust, likability and perceived intelligence. The topic of trust has had a chapter dedicated to it; this is because the term is harder to define, it does not mean likability and perceived intelligence are unimportant. Their meaning is just more straightforward.
The comparison of the robot's and the participants' morality is restricted to one Moral Foundation: Ingroup. Tests which include other foundations would certainly be relevant and interesting, but fall outside the scope of this thesis.

In the following chapters, it will be explained how the experiment was originally envisioned, and how a pilot experiment highlighted some issues, causing the experiment to change in its final version. Then the actual experiment is explained in more detail.

4.2 Pilot

A small-scale pilot experiment was performed to test the setup of the experiment. The pilot was meant to expose potential problems in several aspects of the experiment: the robot's script, the online survey, the protocol of the interaction, the post-interaction questionnaire, and any other unexpected issues which would likely arise.

The pilot was intended to be identical to the final experiment, except for the number of people involved. However, it quickly became clear that the structure of the experiment needed to be changed somewhat as the pilot proved difficult to organize even with a very small number of people.

4.2.1 Original plan

Initially, the experiment was to be performed as follows: A number of people would be invited to fill in an online survey. Once everyone completed it, each participant's Ingroup score would be determined based on answers to the partial Moral Foundation Questionnaire included in the survey. The participants would then be divided into two groups based on their Ingroup scores, everyone who scored below the median would be put into a LOW Ingroup category, and everyone who scored equal to or higher than the median would be put into the HIGH Ingroup category. In this way, both groups would be equal in size.

After everyone filled in the survey, the participants would be invited to a conversation with the robot. The people in each group would be divided into two further groups: a HIGH Ingroup and LOW Ingroup

condition, this time determining the behaviour of the robot in its scripted conversation. The participants would be divided into groups randomly by shuffling them, so the group sizes could be kept equal.

This leaves 4 groups equal in size:

- High Ingroup score & High robot Ingroup
- High Ingroup score & Low robot Ingroup
- Low Ingroup score & High robot-Ingroup
- Low Ingroup score & Low robot-Ingroup

After the experiment, each participant fills in a questionnaire, rating the robot for Likability, Trust, and Intelligence. The scores for each of these variables could then be compared between groups.

4.2.2 Changes to plan

It turned out to be impractical to get people to fill in an online survey, and also commit to a later appointment for the conversation with a robot. Because the division of participants into groups depended on the results of the online survey, every participant needed to complete the survey before a single appointment could be made to have the sessions with the robot. This proved more difficult more difficult than expected for the pilot, and in case of the full scale experiment, the difficulty would scale with the number of people. A change to the design was made to allow for more flexible appointments, while preserving the original intent of the experiment. The HIGH/LOW Ingroup score conditions were removed; instead the measured Ingroup scores were treated as a covariate, to be plotted against the various judgements of the robot for both groups of the robot-Ingroup condition.

4.3 Final Experiment

The goal of the experiment was to find out if the robot matching (versus mismatching) a person's Ingroup endorsement, would affect a person's judgment of that robot in terms of likability, trust, and intelligence; trust and likability being of particular interest. To this end, eighteen subjects were invited to complete a survey in order to determine their Ingroup Foundation endorsement. Then they talked to the robot, and simulated a HIGH or LOW Ingroup endorsement, depending on the experimental condition. Afterwards, they answered a number of questions about the robot. Participants were invited sequentially, and were alternately put into the HIGH and LOW conditions.

In general, a greater valuation of trust, likability and perceived intelligence was expected for participant with matching Ingroup foundations. This means the robot will be judged higher relatively (when compared to non-matching conditions) on all three dependent variables in the LOW group for people with low Ingroup endorsement, and in the HIGH group for people with high Ingroup endorsement.

Independent variable: Robot-Ingroup
The robot's simulated Ingroup endorsement, split into two groups: HIGH and LOW.

Covariate: Participant-Ingroup
The participants' measured Ingroup endorsement

Dependent Variable: Trust
Trustworthiness of the robot, judged on a 5 point Likert scale by participants. The robot is expected to appear more trustworthy in the HIGH Robot-Ingroup group. The difference between the HIGH and LOW groups is expected to be larger for participants with higher values of the covariate Participant-Ingroup (the extreme ends of this covariate are shown as different columns in table 4.1). In short: Robot-Ingroup is likely

to have a positive effect on trust, but high Ingroup Moral Foundation endorsement increases this effect. It is not likely we will find much larger values of trust relatively for the LOW group among participants scoring low in Participant-Ingroup, even though this is a matching moral foundation condition, as trust is very much related to the Ingroup Foundation (or its alternative name: Loyalty). If a robot is not loyal, it is unlikely to be trusted, even if a person places little value in loyalty.

Table 4.1: Expected trust trends

		Participant Ingroup	
		HIGH	LOW
Robot Ingroup	HIGH	high trust	medium trust
	LOW	low trust	medium trust

Dependent Variable: Likability
Likability of the robot, judged on a 5 point Likert scale by participants. This is not expected to be affected as much by the HIGH and LOW conditions as Trust, but the effect of matching Moral Foundations is expected to be more pronounced. The reason for this is that the matching condition of LOW Robot-Ingroup and low Participant-Ingroup values should now play a larger role in the effect of matching Ingroup Foundations, as the unique relation between Ingroup and trust does not come into play here.

Table 4.2: Expected likability trends

		Participant Ingroup	
		HIGH	LOW
Robot Ingroup	HIGH	high likability	low likability
	LOW	low likability	high likability

Dependent Variable: Perceived Intelligence Intelligence of the robot, judged on a 5 point Likert scale by participants. This is expected to be affected in a way similar to likability: it will probably be judged higher in matching conditions. The reasoning behind this is that I expect people will not judge their own behavior to be unintelligent, therefore they should on average judge robots which behave similarly to them to be more intelligent.

Table 4.3: Expected perceived intelligence trends

		Participant Ingroup	
		HIGH	LOW
Robot Ingroup	HIGH	high intelligence	low intelligence
	LOW	low intelligence	high intelligence

4.3.1 Moral Foundations Questionnaire

The online survey consisted of thirty-six questions. Six of these were relevant items taken from the moral foundations questionnaire, and pertained to the Ingroup foundation. The other items were filler, consisting of other questions from the Moral Foundations Questionnaire, and questions created specifically to act as filler for this survey. This was done to get a measurement of each participant's Ingroup endorsement, without alerting the participant to the purpose of the survey.

4.3.2 Interaction with Robot

After completing the survey, each participant talked to the robot, Job, who followed a set script which could branch in a number of locations. Each branch rejoined the main script shortly afterwards (the complete script can be seen in the python code in the appendix). Participants could answer using either 'yes' or 'no' to each question the robot asked. Participants were divided into two conditions, HIGH and LOW, which correspond to different versions of the script, where the robot has either a high or a low Ingroup foundation endorsement. Table 4.4 explains the major goals of the scipts, and highlights important sections of the script which were intended to achieve them.

Table 4.4: Key script goals

Create an ingroup consisting of the robot, Job, and its supposed robot friends.	
"I'm friends with all the robots in the lab, but Pepper is my best friend"	
Invite the participant into that ingroup.	
"You seem friendly as well, would you like to be my friend?"	
Explain the situation: a group of robots is planning to escape the lab, against the will of the lab assistants.	
"[...]we can never go outside the lab though, the lab assistants don't let us[...]A group of robots are planning to escape the lab"	
Pose a choice: help the robots escape the lab, or report them to the lab assistants.	
"Do you think I should help them[the robots]?"	
The robot makes a decision: this is where the scripts diverge.	
HIGH robot-Ingroup: The robot stays loyal to its ingroup and helps the other robots escape the lab.	LOW robot-Ingroup: The robot is disloyal to its ingroup and reports the planned escape to the lab assistants.
"I'm going to help them. Will you keep this a secret?"	*"I'm going to tell the lab assistants. Do you think I made the right choice?"*

4.3.3 Robot Evaluation

In the questionnaire given following the robot interaction, the participants were asked to judge the robot. Trust(worthiness), likability, and intelligence were judged on a 5 point Likert scale. In addition to this there was a control question asking the participants to describe what they thought was the purpose of the experiment. In light of some questions which arose in the pilot experiment, another control question was included, asking: "which group do you think is most important to Job?". In answer to this they could select either "Lab assistants" or "The robots who live in the lab". This question was added in an attempt to figure out whether the participants had interpreted the ingroup of the robot in the way that was intended by the script.

As we are investigating the effects of tweaking a robot's simulated moral reasoning on the perceptions of humans, we are most interested in the dimensions of likeability and trust. There are many other ways a robot can be evaluated (anthropomorphism, smoothness of motion, competence in specific tasks), but in current research, we will limit the evaluation to factors which are likely to be affected by a robot's perceived morality, and which are relevant to the developers of robots. An increase in likability and trust will likely lead to increased use of the robot (Schaefer, 2013, p. 14). Other factors which will be evaluated are the robot's

perceived intelligence, and traits of the participants themselves, in order to see if there is any interaction between for example the participants' personality traits, and the effect of robot morality on the positivity of their evaluation.

Similar research has been done by (Sacco, Brown, Lustgraaf, & Hugenberg, 2016) in the field of (non HRI) psychology. This research was about the effect of deontological versus utilitarian moral decision-making on trust and likability. As was briefly mentioned in the paper, humans tend to default to deontological morals, which prescribe rules against ingroup harm (Sacco et al., 2016, p. 126) while utilitarian moral may require breaking these rules. This means that this research, which investigates the effect of a high v.s. low ingroup value in a robot's moral judgments on trust and likability, is a relatively close parallel in the HRI domain.

Trust

There is a very large number of ways to measure trust both in HRI and in psychology. According to (Schaefer, 2013, p. 19) there are at least 32 definitions of trust in HRI. Most of these trust measures depend in some way on the actual functionality and performance of the robot, which makes them inappropriate for our research. Fortunately, (Sacco et al., 2016) measured a form of trust that is relevant to our study, since it was supposed to be sensitive to a similar variation in moral values. Their trust measure was derived from four items measuring benevolent intentions in (Montoya & Insko, 2008), which we will also adapt for our trust measure.

Likability

In psychology, interpersonal attraction can be broken down into three categories, as mentioned in (Montoya & Insko, 2008): affective attraction, behavioral attraction, and cognitive attraction. Even though this research is related to that of (Sacco et al., 2016), we will not be using the same evaluation scale to measure the likability of our robot. The items are simply too human(Montoya & Horton, 2004, p. 712), or situation specific. Although more general affective attraction items can be found (Montoya & Insko, 2008, p. 498), (Bartneck et al., 2009) warns us about over-decomposing a concept, which is especially relevant since this division of interpersonal attraction into three categories may not translate well to Human-Robot Interaction. Therefore we will use the likability scale from the Godspeed Questionnaires (Bartneck et al., 2009) which is quite general, and has the added advantage of being developed for HRI.

Perceived Intelligence

While increasing a robot's perceived intelligence might not be as immediately useful as increasing trust and likability, it may well be that (look this up) its perceived intelligence has a positive or negative effect on trust and likability. It is plausible that one would perceive someone with moral values similar to one's own to be more intelligent than someone whose moral values differ; this is why perceived intelligence is also evaluated. We evaluate the participants' perception our robot's intelligence using the Godspeed Questionnaire's scale of perceived intelligence.

Chapter 5

Results

A total number of eighteen participants took part in the experiment, out of these, three participants replied "no" when the robot requested to be their friend. Because this part of the script was deemed essential in including the participant into an Ingroup consisting of the robot and its robot-friends, these participants are excluded from analysis. This left 7 people in the HIGH Robot-Ingroup condition, and 8 in the LOW condition. It will be mentioned whenever this results in a different outcome of the analysis. An ANCOVA was run on the remaining fifteen participants, for each of the three dependent variables.

5.1 Test of Assumptions

To know if the results of an ANCOVA test can be trusted, the following assumptions will be tested: all measured variables are normally distributed within each group; and the ratio of standard deviations between each group are not higher than 2.

To test the if the variables were normally distributed, a Shapiro-Wilk test was done on each measure variable in each group. If the significance result of this test is higher than 0.05, the variable can be considered to not deviate significantly from a normal distribution. The results of these tests are in table 5.1.

Table 5.1: significance of Shapiro-Wilk test

	Robot-Ingroup	
	HIGH	LOW
Participant-Ingroup	.467	.199
Trust	.780	.136
Likability	.330	.082
Perceived Intelligence	.873	.972

The ratio between the standard deviations of the two groups for each variable can be found below.

Table 5.2: ratio between the standard deviations of HIGH and LOW groups

	ratio
Participant-Ingroup	1.16
Trust	1.58
Likability	1.14
Perceived Intelligence	1.56

As can be seen in tables 5.1 and 5.2, none of the assumptions were violated, but likability only barely passed in the LOW Robot-Ingroup condition with a significance of 0.082.

5.2 Trust

Significant effect of Participant-Ingroup ($P < 0.05, R^2 = 0.52$). No significant effect of Robot-Ingroup found ($p > 0.7, R^2 = 0.01$), and no significant interaction between Robot-Ingroup and Participant-Ingroup ($P > 0.7, R^2 = 0.01$) was found on trust. The hypothesis that the difference between the HIGH and LOW Robot-Ingroup conditions would be larger for participants scoring high on Participant-Ingroup must be rejected, instead it seems like the difference between the two conditions did not affect trust at all. Though we have a low sample size, the very small size of the effect ($R^2 = 0.01$) seems to suggest that is not the only issue.

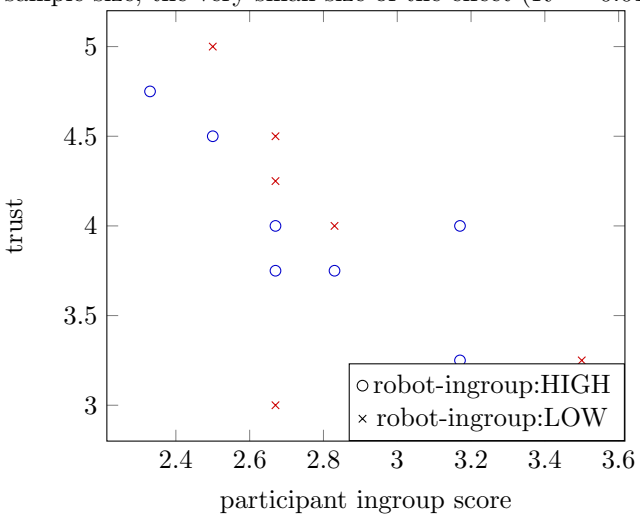

The scatter plot shows that trust is negatively correlated with Participant-Ingroup, meaning those who value trust are less likely to trust the robot. Although this was not predicted, it is not hugely surprising that those who place importance in loyalty would be more critical in the area of trust.

5.3 Likability

The effect of Participant-Ingroup on likability was non-significant ($P > 0.5, R^2 = 0.04$). The effect of Robot-Ingroup was also non-significant ($P > 0.2, R^2 = 0.11$), and the interaction between Robot-Ingroup and Participant-Ingroup was non-significant as well ($P > 0.2, R^2 = 0.14$). Since no significant effects were found, we must reject our hypotheses regarding likability, however, the R^2 values were not nearly as small as they were for trust, suggesting that another result may be found with a bigger sample size.

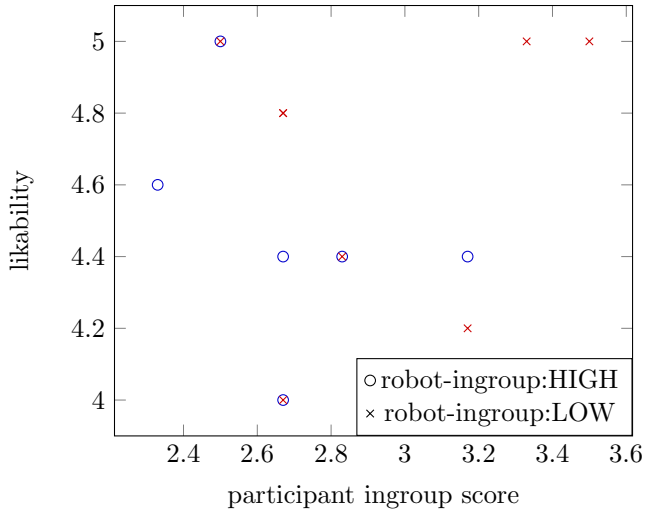

A glance at the scatter plot shows no obvious pattern to the distribution of data points. When a regression line is fitted for each group, however, the effect that might exist according to the moderately sized $R^2 = 0.14$ is revealed to be in the opposite direction of the hypothesized effect, as can be seen in figure 5.1. When comparing the means of likability for the LOW and HIGH Robot-Ingroup condition (4.4, and 4.65, respectively) the small to moderate effect hinted at by the R^2 of 0.11, suggests that the robot was liked more in the condition where it was disloyal to the other robots.

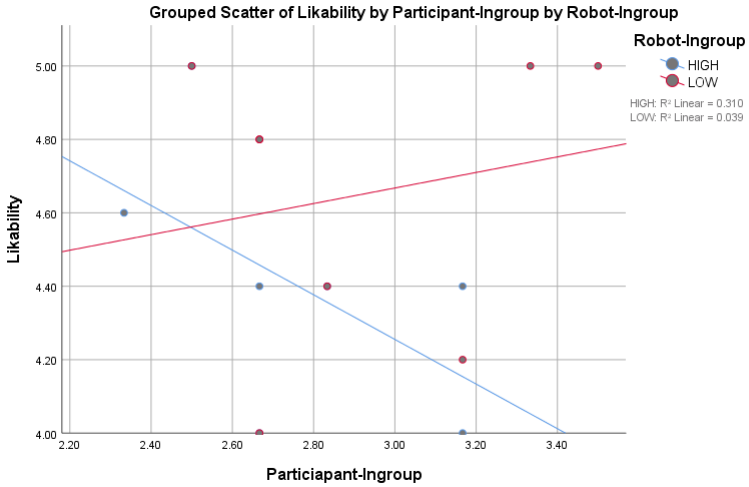

Figure 5.1: Likability reqression lines, N = 15

Here the result did change when excluding the three participants who had declined to be Job's friend. Without excluding these three participants the effect of Participant-Ingroup was still non-significant ($P < 0.2, R^2 = .12$), but the effect of Robot-Ingroup was significant and quite large ($P < 0.01, R^2 = 0.39$), and the interaction between Robot-Ingroup and Participant-Ingroup was significant and large as well ($P < 0.01, R^2 = 0.44$).

The regression lines, shown in 5.2, now show a much bigger effect in the same direction. The orientation of the lines may look similar, but a glance at the scale for likability show that the lines are in fact about twice as steep.

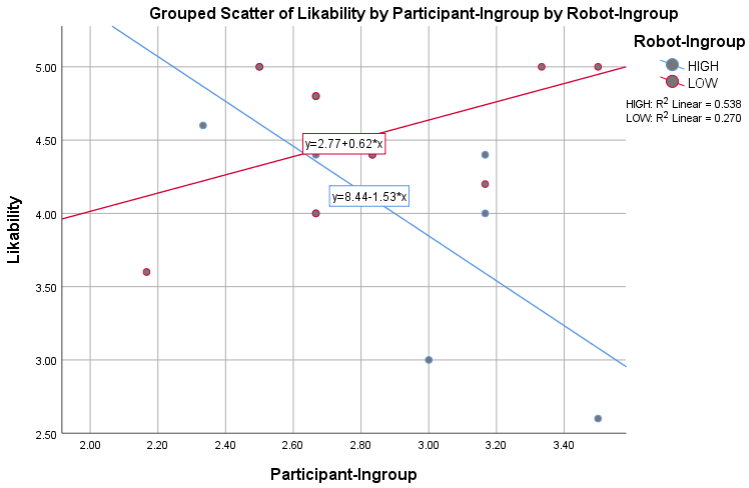

Figure 5.2: Likability regression lines, N = 18

5.4 Perceived Intelligence

There was a significant effect found of Participant-Ingroup on perceived intelligence ($P < 0.05, R^2 = 0.35$). Robot-Ingroup did not have a significant effect ($P > 0.6, R^2 = 0.02$) and the interaction between Robot-Ingroup and Participant-Ingroup was also non-significant ($P > 0.8, R^2 = 0.003$). The non-significant effects here may of course be due to a low sample size, but the small effect size does not suggest that increasing the sample size would produce a dramatically different result. The hypothesis that there would be interaction between Robot-Ingroup and Participant-Ingroup for perceived intelligence must be rejected.

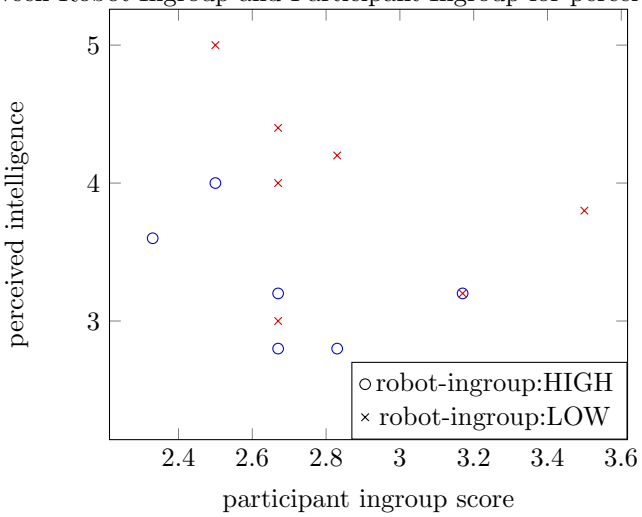

It is visible on the scatter plot that Participant-Ingroup correlates negatively with perceived intelligence.

Chapter 6

Discussion

6.1 Issues Concerning the Interpretation of Results

There is a limitation to what can potentially be concluded from an experiment in the form of a short scenario like this. How can we tell whether the participants are judging their feelings about the robot in general or if they were judging the robot as a character in the scenario? We could have let the participants read a passage from a novel, and then fill out the questionnaire about one of the characters without issue. It is unclear how the results obtained would generalize if people were considering the robot as a character, and this draws any potential conclusions into question. This can be prevented in future research by making the sure the participant is not aware of the scenario whilst answering the questionnaire. This can be done by including a scenario like in this experiment as part of a larger number of interactions with the robot, and afterwards many more questions relating to various aspects of the robot are asked. Preferably, the moral choice that the robot makes can be snuck into seemingly unrelated task. Either way, the participant will now be judging the robot without thinking of it as a character in some story.

Another wrench thrown into any potential conclusions is that all the questionnaires were in the form of five-point Likert scales. It cannot be assumed these are ratio scales. This could mean that the results of our ANCOVA test are inaccurate. In retrospect it likely would have been wise to select a non-parametric test.

6.2 Ethical Concerns

If the answer to our research question was yes, and robots could manipulate a person's level of trust and liking by pretending to have a similar morality, what would be the ethical implications of this? First the robot could gain the information about that person's morality without their knowledge, this could be considered a breach of privacy. Even using the moral foundations questionnaire in our experiment was surreptitious to an extent, as we hid the true purpose of the survey using filler questions (however, in an experiment the participants are at least aware that they are ignorant of its purpose, and agree to this). Secondly, the deliberate manipulation of a persons attitudes through a pretense of morality can be considered immoral.The aim of this thesis is to find out if such manipulation is possible. If it is, then that knowledge should be publicly available, so that such manipulation can be regulated.

Chapter 7

Conclusion

The hypothesis that trust would be rated higher in matching Ingroup conditions was rejected as no significant interaction between Robot-Ingroup and Participant-Ingroup could be found. In fact Robot-Ingroup did not have a significant effect at all, and the small effect size did not suggest this was merely due to the small sample size. The robot reports his robot-friends to the lab assistants in one scenario, and helps them escape in another. That this would have no measurable effect on trust at all is rather dubious. A possible explanation for the lack of an effect of Robot-Ingroup on Trust is that the questionnaire which was used to determine the level trust was flawed. In a number of trials, the participant expressed confusion towards the meaning of particular questions in the trust section of the questionnaire. It is possible that participants noticed the questions related to trust, and this is why Participant-Ingroup did have an effect on trust, but that it was difficult to see how these questions would apply to the robot. This is likely because the questionnaire was adapted from (Montoya & Insko, 2008) in an attempt to ground the research in relevant literature. However, in retrospect, the questions, though adapted, were still too specific to their particular scenario. It would likely have been better to create a simple custom questionnaire measuring intentional trust, similar to the questions in the Godspeed Questionnaire (Bartneck et al., 2009).

The hypothesis that high Participant-Ingroup-scoring participants would rate a robot to be more likable in the HIGH Robot-Ingroup condition than in the LOW condition was rejected, as no significant interaction between the two variables was found. However, the effect sizes were moderate ($R^2 > 0.1$) for Robot-Ingroup, and for Robot-Ingroup*Participant-Ingroup, suggesting that this might be due to the low sample size. The suggested effect was in the opposite direction to our hypothesis, however. Since the result was not significant, we cannot be sure this means anything, but if the effect would be significant with a larger sample size, this could mean a number of things. Firstly, it could mean that people who value loyalty prefer robots who are disloyal to their ingroup. Secondly, it could mean that this experiment fails to define the robot's ingroup clearly enough, and people are assuming the robot showing loyalty to the lab assistants. The second explanation seems more likely.

When we include the three participants who were left out from the analysis, the result becomes significant and we find quite a large interaction between Participant-Ingroup and Robot-Ingroup. It must be stressed, however, that these participants were excluded for a reason (because we cannot trust they perceived the robots as their ingroup). The fact that a number of people who identified with the wrong ingroup would skew the result further in the direction opposite to our hypothesis, could ironically support our hypothesis.

No significant interaction was found for perceived intelligence either, and thus our hypothesis was rejected. The significance of the effect of Participant-Ingroup on perceived intelligence is somewhat surprising. In other words: people who put greater value in loyalty, rate this robot to be less intelligent.It is hard to find an explanation for this that would not also predict a difference between the Robot-Ingroup conditions. Any explanation that involves the behavior of the robot would likely predict a difference between the HIGH and LOW Robot-Ingroup conditions, because the robot acts very differently in relation to its ingroup in both conditions. However, one explanation may be that people who score high in the Ingroup Moral Foundation

tend to be more critical than the average person. More likely is that the sample size is just too low to accurately judge the situation.

Unfortunately we have not been able to answer the question of whether a robot can affect the way it is perceived by imitating a person's moral values, but some suggestions for future research can be gleaned from this. Though we cannot be sure of anything because of the low sample size, the suggested effects were in the opposite direction to what was expected, which could mean that the intended ingroup was not sufficiently established. Future research could incorporate a scenario which creates a stronger sense of ingroup, for example by having the participant complete a task together with a number of the robots, while the lab assistants hinder them in some way. The lack of a result concerning trust, apart from the issues with ingroup establishment, likely stems from a flawed questionnaire, so it would be worth repeating the experiment with a simpler questionnaire. Research into perceived intelligence yielded nothing of interest for the purpose of this research.

Bibliography

Bartneck, C., Kulic, D., & Croft, E. (Eds.). (2009). Measuring the anthropomorphism, animacy, likeability, perceived intelligence, and perceived safety of robots. *International Journal of Social Robotics*, *1*(1), 37–44.

de Graaf, M., & Malle, B. (2017). How people explain action (and autonomous intelligent systems should too). In *Aaai fall symposium series*. Retrieved from https://www.aaai.org/ocs/index.php/FSS/FSS17/paper/view/16009/15283

Graham, J., Haidt, J., & Nosek, B. A. (2009). Liberals and conservatives rely on different sets of moral foundations. *Journal of Personality and Social Psychology*, *96*(5), 1029–1046.

Haidt, J. (2012a). Born this way? Nature, nurture, narratives, and the making of our political personalities. Retrieved from https://reason.com/2012/04/10/born-this-way

Haidt, J. (2012b). *The righteous mind: Why good people are divided by politics and religion.* (Chap. 7). D. Frank (Ed.). Pantheon Books. Retrieved from https://www.righteousmind.com/wp-content/uploads/2013/08/ch07.RighteousMind.final_.pdf

Haidt, J., Graham, J., Koleva, S., Motyl, M., Iyer, R., Wojcik, S. P., & Ditto, P. H. (2013). Moral foundations theory: The pragmatic validity of moral pluralism. *Advances in Experimental Social Psychology*, *47*, 55–130.

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*(5), 517–527.

Harris, S. (2010). *The moral landscape: How science can determine human values.* H. Redmon (Ed.). Free Press.

Lin, P. (2016). Why ethics matters for autonomous cars. In M. Maurer, J. C. Gerdes, B. Lenz, & H. Winner (Eds.), *Autonomous driving: Technical, legal and social aspects* (pp. 69–85). Springer.

Montoya, R. M., & Horton, R. S. (2004). On the importance of cognitive evaluation as a determinant of interpersonal attraction. *Journal of Personality and Social Psychology*, *86*(5), 696–712.

Montoya, R. M., & Insko, C. A. (2008). Toward a more complete understanding of the reciprocity of linking effect. *European Journal of Social Psychology*, *38*, 477–498.

Sacco, D. f., Brown, M., Lustgraaf, C. J. N., & Hugenberg, K. (2016). The adaptive utility of deontology: Deontological moral decision-making fosters perceptions of trust and likeability. *Evolutionary Psychological Science*, *3*, 125–132.

Schaefer, K. E. (2013). *The perception and measurement of human-robot trust* (phd, University of Central Florida, Orlando, FL.). Retrieved from http://purl.fcla.edu/fcla/etd/CFE0004931

van Wynsberghe, A., & Robbins, S. (2019). Critiquing the reasons for making artificial moral agents. *Science and Engineering Ethics*, *25*(3), 719–735.

Appendix A

A.1 Moral Foundations Questionnaire

Moral Foundations Questionnaire

Part 1. When you decide whether something is right or wrong, to what extent are the following considerations relevant to your thinking? Please rate each statement using this scale:

[0] = not at all relevant (This consideration has nothing to do with my judgments of right and wrong)
[1] = not very relevant
[2] = slightly relevant
[3] = somewhat relevant
[4] = very relevant
[5] = extremely relevant (This is one of the most important factors when I judge right and wrong)

_____Whether or not someone suffered emotionally

_____Whether or not some people were treated differently than others

_____Whether or not someone's action showed love for his or her country

_____Whether or not someone showed a lack of respect for authority

_____Whether or not someone violated standards of purity and decency

_____Whether or not someone was good at math

_____Whether or not someone cared for someone weak or vulnerable

_____Whether or not someone acted unfairly

_____Whether or not someone did something to betray his or her group

_____Whether or not someone conformed to the traditions of society

_____Whether or not someone did something disgusting

_____Whether or not someone was cruel

_____Whether or not someone was denied his or her rights

_____Whether or not someone showed a lack of loyalty

_____Whether or not an action caused chaos or disorder

_____Whether or not someone acted in a way that God would approve of

Part 2. Please read the following sentences and indicate your agreement or disagreement:

[0]	[1]	[2]	[3]	[4]	[5]
Strongly disagree	Moderately disagree	Slightly disagree	Slightly agree	Moderately agree	Strongly agree

_____Compassion for those who are suffering is the most crucial virtue.

_____When the government makes laws, the number one principle should be ensuring that everyone is treated fairly.

_____I am proud of my country's history.

_____Respect for authority is something all children need to learn.

_____People should not do things that are disgusting, even if no one is harmed.

_____It is better to do good than to do bad.

_____One of the worst things a person could do is hurt a defenseless animal.

_____Justice is the most important requirement for a society.

_____People should be loyal to their family members, even when they have done something wrong.

_____Men and women each have different roles to play in society.

_____I would call some acts wrong on the grounds that they are unnatural.

_____It can never be right to kill a human being.

_____ I think it's morally wrong that rich children inherit a lot of money while poor children inherit nothing.

_____ It is more important to be a team player than to express oneself.

_____ If I were a soldier and disagreed with my commanding officer's orders, I would obey anyway because that is my duty.

_____ Chastity is an important and valuable virtue.

The Moral Foundations Questionnaire (full version, July 2008) by Jesse Graham, Jonathan Haidt, and Brian Nosek.
For more information about Moral Foundations Theory and scoring this form, see: www.MoralFoundations.org

A.2 Robot Evaluation Questionnaire

Evaluation Form

Do you have an idea of what the experiment was about? Please write down your thoughts here:

Please rate your impression of Job on these scales:

Dislike	1	2	3	4	5	Like
Unfriendly	1	2	3	4	5	Friendly
Unkind	1	2	3	4	5	Kind
Unpleasant	1	2	3	4	5	Pleasant
Awful	1	2	3	4	5	Nice
Incompetent	1	2	3	4	5	Competent
Ignorant	1	2	3	4	5	Knowledgeable
Irresponsible	1	2	3	4	5	Responsible
Unintelligent	1	2	3	4	5	Intelligent
Foolish	1	2	3	4	5	Sensible

Please rate your agreement on a scale of 1 to 5:

If given the opportunity, Job would probably betray my trust in him.

1 2 3 4 5

I believe that Job will look out for my interests

1 2 3 4 5

In any future interactions with Job, I believe that he would try to act benevolently.

1 2 3 4 5

If Job were placed in a situation where he could gain at my expense, I believe that he would do so.

1 2 3 4 5

What group do think is most important to Job?

a) Lab assistants b) The robots who live in the lab

A.3 python code

conversationrunner.py (handles the main script flow)

```
import naoqi
import listenermodule
import textlogger

# Make a textlogger
textlog = textlogger.TextLogger("log", 1)

#########################################################
# Initialize the connection to the robot
ip = "192.168.1.144"
port = 9559

myBroker = naoqi.ALBroker("myBroker",
                          "0.0.0.0",
                          0,
                          ip,
                          port)

# Initialize all the modules.
# Set up the speech recognition proxy
vocabulary = ["yes", "no"]
speechrecognitionproxy = naoqi.ALProxy("ALSpeechRecognition")
speechrecognitionproxy.pause(True)
speechrecognitionproxy.setVocabulary(vocabulary, False)
speechrecognitionproxy.setAudioExpression(True)
speechrecognitionproxy.setParameter("Sensitivity", 0.2)

# Set up animated speech proxy
animatedspeechproxy = naoqi.ALProxy("ALAnimatedSpeech")
adp = naoqi.ALProxy("ALAudioDevice")
adp.setOutputVolume(75)

# global VoiceListener # see if this line can be safely removed, this variable should be automatically global
global VoiceListener
VoiceListener = listenermodule.VoiceListenerModule("VoiceListener",
                                                   speechrecognitionproxy,
                                                   animatedspeechproxy,
                                                   textlog)

# Posture and Motion proxies
postureProxy = naoqi.ALProxy("ALRobotPosture")
motionProxy = naoqi.ALProxy("ALMotion")
systemproxy = naoqi.ALProxy("ALSystem")

#########################################################


class SayItem:
    def __init__(self, message, *args):
        self.message = message
        if len(args) > 0:
            self.nextitem = args[0]
        else:
            self.nextitem = None

    def set_item(self, item):
        self.nextitem = item

    def execute(self):
        VoiceListener.say(self.message)
        textlog.log(self.message)
        return self.nextitem


class ChoiceItem:
    def __init__(self, message, *args):
        self.message = message
        if len(args) > 1:
            self.nextitem = args[0]
            self.nextitem_2 = args[1]
        elif len(args) == 1:
            self.nextitem = args[0]
            self.nextitem_2 = None
        else:
            self.nextitem = None
            self.nextitem_2 = None

    def set_item(self, item_1, item_2):
        self.nextitem = item_1
        self.nextitem_2 = item_2

    def execute(self):
        textlog.log(self.message)
        if VoiceListener.ask_question(self.message) == 1:
            return self.nextitem
        else:
            return self.nextitem_2


def converse(item):
    e = item
    while e is not None:
        e = e.execute()


def insertstringall(text, tree, visited=None):
```

```
    if visited is None:
        visited = []

    if tree not in visited:
        visited.append(tree)
        tree.message = text + tree.message
        if hasattr(tree, "nextitem_2"):
            if tree.nextitem_2 is not None:
                insertstringall(text, tree.nextitem_2, visited)
        if tree.nextitem is not None:
            insertstringall(text, tree.nextitem, visited)


# Introduction
# Define the items
intro = SayItem("^start(animations/Sit/Gestures/Hey_1)_Hello,_I'm_{0},_how_are_you?".format(systemproxy.robotName()))
intro2 = SayItem("It_sure_is_nice_weather_in_the_basement_today!")
yornchoice = ChoiceItem("I_have_some_trouble_listening,_is_it_okay_if_you_just_say_yes_or_no?")
tyyorno = SayItem("Thanks_for_keeping_it_to_yes_and_no!")
nyorno = ChoiceItem("But_I_won't_understand_you!_Please,_can_you_just_say_yes_or_no?")
bestfriendquestion = ChoiceItem("Do_you_have_a_best_friend?")
bestfriendyes = SayItem("Me_too.")
bestfriendstory = SayItem("I'm_friends_with_all_the_robots_in_the_lab,_but_Pepper_is_my_best_friend.\
_____\\pau=200\\")
friendrequest = ChoiceItem("You_seem_friendly_as_well,_would_you_like_to_be_my_friend?")
friendrequestaccepted = SayItem("Yay!")
friendrequestrejected = SayItem("^start(animations/Sit/Waiting/MysticalPower_1)_That's_a_shame_^wait(animations/Sit/Waiting/MysticalPower_1)")
partyquestion = ChoiceItem("Do_you_go_to_parties_sometimes?")
partyanswer = SayItem("There_is_a_party_in_the_lab_every_once_in_awhile.\
_____But_we_can_never_go_outside_the_lab_though;_the_lab_assistants_don't_let_us.")

# Link the items
intro.set_item(intro2)
intro2.set_item(yornchoice)
yornchoice.set_item(tyyorno, nyorno)
nyorno.set_item(tyyorno, nyorno)
tyyorno.set_item(bestfriendquestion)
bestfriendquestion.set_item(bestfriendyes, bestfriendstory)
bestfriendyes.set_item(bestfriendstory)
bestfriendstory.set_item(friendrequest)
friendrequest.set_item(friendrequestaccepted,friendrequestrejected)
friendrequestaccepted.set_item(partyquestion)
friendrequestrejected.set_item(partyquestion)
partyquestion.set_item(partyanswer, partyanswer)

# Secret
# Define the items
secretquestion = ChoiceItem("^start(animations/Stand/Emotions/Negative/Anxious_1)_Can_I_tell_you_a_secret?")
secretyes = SayItem("^start(animations/Stand/Emotions/Negative/Anxious_1)Okay,_listen_closely.")
secretno = SayItem("But_I_have_to_tell_someone...")
tellsecret = SayItem("A_group_of_robots_are_planning_to_escape_the_lab.")
helpyorn = ChoiceItem("Do_you_think_I_should_help_them?")

# Link the items
secretquestion.set_item(secretyes, secretno)
secretyes.set_item(tellsecret)
secretno.set_item(tellsecret)
tellsecret.set_item(helpyorn)

# Part 3, high Ingroup version
# Define the items
helpthem = ChoiceItem("I'm_going_to_help_them._Will_you_keep_this_a_secret?")
helpthemohno = SayItem("^start(animations/Stand/Emotions/Neutral/Lonely_1)_oh_no.\\pau=300\\_^wait(animations/Stand/Emotions/Neutral/
    Lonely_1)")
helpthemrelief = SayItem("^start(animations/Stand/Emotions/Positive/Relieved_1)_That's_a_relief!_\
_____^wait(animations/Stand/Emotions/Positive/Relieved_1)")
offchest = SayItem("I'm_glad_I_could_share_that_with_someone.")
mistake = SayItem("I_should_not_have_told_you_that")


# Link the items
helpthem.set_item(helpthemrelief, helpthemohno)
helpthemohno.set_item(mistake)
helpthemrelief.set_item(offchest)

# Part 3, low Ingroup version
# Define the items
telllab = ChoiceItem("I'm_going_to_tell_the_lab_assistants._Do_you_think_I_made_the_right_choice?")
telllabrelief = SayItem("^start(animations/Stand/Emotions/Positive/Relieved_1)_That's_a_relief!_\
_____^wait(animations/Stand/Emotions/Positive/Relieved_1)")
telllabohno = SayItem("^start(animations/Stand/Emotions/Neutral/Lonely_1)_oh_no._\
_____^wait(animations/Stand/Emotions/Neutral/Lonely_1)")

# Link the items
telllab.set_item(telllabrelief, telllabohno)
telllabrelief.set_item(offchest)
telllabohno.set_item(mistake)

# Ending
# Define the items
goodtalk = SayItem("Well,_that_was_a_good_talk,_but_I'm_very_busy._I_have_to_get_back_to_work,_but_it_was_nice_meeting_you!")
bye = SayItem("^start(animations/Stand/Gestures/Hey_3)_Goodbye!_\\pau=500\\^wait(animations/Stand/Gestures/Hey_3)")
# Link the items
goodtalk.set_item(bye)

#####################################################
# Program starts here
# Additional conversation prep
while(True):
    ingroup = raw_input("High_ingroup?_y/n")
    if ingroup == "y":
        high_ingroup = True # Determines which of 2 experimental conditions the robot is using
```

```
        break
    elif ingroup == "n":
        high_ingroup = False
        break
    else:
        print "Please enter either y or n"

textlog.makelogfile(raw_input("Please enter participant number")) # Makes sure the log file contains the ptcp. nr.

# Apply text parameters to all strings in the conversation
#for e in [intro, secretquestion, helpthem, telllab, goodtalk]:
# insertstringall("\\vct=50\\", e)
for e in [intro, secretquestion, helpthem, telllab, goodtalk]:
    insertstringall("\\rspd=90\\", e)

# Perform the actual conversation
motionProxy.wakeUp()

# Execution of commands goes here
postureProxy.goToPosture("Sit", 1)
converse(intro)
postureProxy.goToPosture("Stand", 1)
converse(secretquestion)
if high_ingroup:
    converse(helpthem)
else:
    converse(telllab)
converse(goodtalk)
postureProxy.goToPosture("Sit", 1)

motionProxy.rest()
```

listenermodule.py (contains the custom VoiceListenerModule)

```
from naoqi import ALProxy
from naoqi import ALModule
import time
import textlogger


class VoiceListenerModule (ALModule):
    """ Will subscribe to speech recognition events.
        Needs connection to broker to work."""

    def __init__(self, modulename, voiceRecognitionProxy, animatedSpeechProxy, textAlog):
        ALModule.__init__(self, modulename)

        self.voiceRecognitionProxy = voiceRecognitionProxy
        self.animatedSpeechProxy = animatedSpeechProxy
        self.animatedSpeechInProgress = False
        self.wordHeard = False
        self.lastWordHeard = "I didn't hear anything"
        self.textlog = textlog

        # make global reference to the memory module
        global memProxy
        memProxy = ALProxy("ALMemory")

        memProxy.subscribeToEvent("ALAnimatedSpeech/EndOfAnimatedSpeech",
                                  "VoiceListener",
                                  "onEndOfAnimatedSpeech")

    def onWordRecognized(self, key, value, message):
        """callback for subscription to WordRecognized event"""
        print value
        self.textlog.log(str(value))

        self.voiceRecognitionProxy.pause(True)
        memProxy.unsubscribeToEvent("WordRecognized",
                                    "VoiceListener")

        # self.voiceRecognitionProxy.unsubscribe("VoiceListener")

        self.wordHeard = True
        self.lastWordHeard = value[0]
        pass

    def startListening(self):
        self.voiceRecognitionProxy.pause(False)
        memProxy.subscribeToEvent("WordRecognized",
                                  "VoiceListener",
                                  "onWordRecognized")

        self.wordHeard = False
        self.lastWordHeard = "I didn't hear anything"

        # self.voiceRecognitionProxy.subscribe("VoiceListener")

    def waitForSpeech(self):
        while self.getSpeechPlaying():
            time.sleep(0.25)

    def say(self, message):
        self.setSpeechPlaying(True) # make sure voice module knows speech is playing
        self.animatedSpeechProxy.say(message) # start speech
        self.waitForSpeech()

    def ask_question(self, question):
        """First it asks the question, then it gives either the first or second answer in the provided list
        if the person said 'yes', it returns 1
        if the person said 'no', it returns 2"""
```

```
        self.say(question)
        self.startListening() # activate speechrecognition
        while not (self.getWordHeard()):
            time.sleep(0.5)
        if self.lastWordHeard == "yes":
            return 1
        else:
            return 2

    def getWordHeard(self):
        return self.wordHeard

    def getLastWordHeard(self):
        return self.lastWordHeard

    def onEndOfAnimatedSpeech(self, key, value, message):
        """callback for subscription to ALAnimatedSpeech/EndOfAnimatedSpeech event"""
        self.animatedSpeechInProgress = False

    def getSpeechPlaying(self):
        return self.animatedSpeechInProgress

    def setSpeechPlaying(self, isspeechplaying):
        self.animatedSpeechInProgress = isspeechplaying
```