Ethical Views, Religious Views, and Acceptance of Robotic Applications: A Pilot Study

Ted Metzler[1] and Lundy Lewis[2]

[1]Wimberly School of Religion, Oklahoma City University, 2501 N. Blackwelder, Oklahoma City, OK 73106-1493
[2]Dept. of Information Technology, Southern New Hampshire University, 2500 N. River Road, Manchester, NH 03106-1043
tmetzler@okcu.edu l.lewis@snhu.edu

Abstract

This paper reports a pilot study investigating relations of people's moral and religious beliefs with their acceptance of specific suggested scenarios for human-robot interaction (HRI) involving life-like personal robots. Data collected via a multiple-choice survey that was focused upon these three classes of variables were subjected to k-means cluster analysis, disclosing some interesting prototype patterns of responses that recommend specific hypotheses for follow-on research.

Introduction

Currently available examples of fairly life-like personal robots include *animaloids* (such as the Sony Corporation's robot dog AIBO), *humanoids* (illustrated by the Honda Motor Corporation's ASIMO), and *androids* (e.g., the female android EveR-1, developed at the Korea Institute of Industrial Technology). The pilot study reported in this paper is motivated, in part, by indications that people's *acceptance* of various scenarios for HRI with life-like personal robots might be influenced by their *religious* beliefs. The opinion that such robots tend to be accepted differently in cultures endorsing different religious faith traditions has been voiced repeatedly in recent years. Writing in 2006 about Japan, *The Chicago Tribune*'s Mark Jacob observed, for example, "The nation's chief religions – Shintoism and Buddhism – are extraordinarily open to robots in people's lives." Professor Masahiro Mori, who happens to combine practice of Buddhism with his work in robotics engineering, supplies us with corroboration for this observation with the following comments from his book, *The Buddha in the Robot*:

> From the Buddha's viewpoint, there is no master-slave relationship between human beings and machines. The two are fused together in an interlocking entity. Man achieves dignity not by subjugating his mechanical inventions, but by recognizing in machines and robots the same Buddha-nature that pervades his own inner self. (180)

Again, Anthony Faiola directs attention to possible links of religious differences with attitudes toward robotics in a 2005 essay for *The Washington Post*:

> "In Western countries, humanoid robots are still not very accepted, but they are in Japan," said Norihiro Hagita, director of the ATR Intelligent Robotics and Communication Laboratories […]. "One reason is religion. In Japanese (Shinto) religion, we believe that all things have gods within them. But in Western countries, most people believe in only one God." (3)

Indeed, the late Christian theologian Paul Tillich – in Volume Three of his *Systematic Theology* – displayed evidence of a much different "Western" understanding of the relation between humans and the machines that they create:

> There are no 'things' in nature, that is, no objects which are nothing but objects which have no element of subjectivity. But objects that are produced by the technical act *are* things. It belongs to man's freedom in the technical act that he can transform natural objects into things: trees into wood, horses into horsepower, men into quantities of workpower. In transforming objects into things, he destroys their natural structures and relations. But something also happens to man when he does this, as it happens to the objects which he transforms. He himself becomes a thing among things. (74)

Some systematic empirical study of actual relations between people's *religious* views and their acceptance of certain robotic technology applications, then, seems to be recommended by the foregoing kinds of opinions and academic observations.

In a similar manner, we have noted indications that people's acceptance of certain interactions with life-like personal robots might well be related to positions that they endorse in *ethical* theory. The theological ethics articulated by Paul Tillich in *Morality and Beyond*, for example, assigns the crucial task of guiding our applications of moral principles to nothing less than unlimited love (*agape*)! Implementing this prescription in computer software – even with provision for machine learning in the robotic system – could be a formidable obstacle to equipping such systems with authentic moral

status. Worse yet, it also has been suggested that reluctance to *ascribe* full moral status to robots could plausibly generate *additional* problems. Reflecting on their studies of HRI involving children and the robotic dog AIBO, Gail Melson and colleagues voiced the following concerns in 2006:

> If it is the case, for example, that people will develop rather robust social relationships with social robots, but not accord them full moral status, then we may be creating one-sided interactions, not unlike a person might have with a therapist, a servant, or even a slave. In turn, is it possible that increasing one's interactions with social robots will lead people to see other humans or animals as "robot-like"? That is as equally problematic. Both concerns, we suggest, should be part of the research and design space of the communities in artificial intelligence, machine learning and adaptive systems, robotics, and android science. (41)

For the animaloid, humanoid, and android personal robots of focal concern in the present paper, therefore, we have considered it a reasonable plan to include sampling of respondents' *ethical* views in our study.

In addition, we note results of prior research indicating that people's acceptance of robots can be sensitive, *inter alia*, to specification of *particular application scenarios* – i.e., it appears to matter just what kind of *task* the robot is suggested to be performing. Specifically, Professor Tatsuya Nomura and colleagues, using the "Negative Attitudes Toward Robots Scale" (NARS), have reported results their social research revealing "attitudes toward robots differ depending on assumptions about robots such as their type and task …" (29). Accordingly, the survey form used in our study probes acceptance of personal robot types explicitly in terms of specific "application scenarios" that suggest, in each case, a particular task and/or role for the robot.

Methodology

Experimental Protocol

Prior to its execution, the research protocol for this study was approved by the Institutional Review Board of Oklahoma City University (OCU). The protocol required each participant in the survey to be at least eighteen years old, and it included initial provision of an Informed Consent Form (ICF). Pending his or her signing and dating of the ICF, each participant was then supplied with a ten-page paper Survey Form. All participation was voluntary, and all Survey Forms remained anonymous to protect privacy of participants. The study was named "HRI Study #1" and was advertised on the OCU campus via posters and an announcement in the campus newspaper. Most of the data collection was conducted on campus, although several off-campus groups also chose to participate. HRI Study # 1 was sponsored by OCU's Darrell W. Hughes Program for Religion and Science Dialogue, which has posted results of the study on its website.

Data Collection

The Survey Form used for this study contained five parts. The first part supplied an introduction to the subject technology of life-like personal robots, defining terms such as "humanoid" and "android" that were used elsewhere in the form. The second part collected several items of "demographic" data that might have future value (age range, sex, and affiliation with OCU), although they were not part of the data analysis planned for this study. The remaining three parts of the Survey Form were comprised of multiple-choice items in the categories of Ethics, Religion, and Acceptance, respectively. For the Ethics and Religion parts of the form, these items typically consisted of statements (e.g., "Some actions are morally right actions regardless of the consequences that follow from them" or "There are gods, or spirits, present everywhere in the world that we experience"), followed by five choices:

1. I definitely agree.
2. I'm inclined to agree.
3. I neither agree nor disagree.
4. I'm inclined to disagree.
5. I definitely disagree.

In the final Acceptance part of the form, the items were brief descriptions of robotic application scenarios (e.g., "Humanoid or android robots work as receptionists at business offices" or "Humanoid robots teach ethics classes at universities"). These were, again, followed by five choices, albeit with "approve" replacing the "agree" language of the preceding parts of the form. The numbers of items in the Ethics, Religion, and Acceptance parts of the form were 8, 12, and 9, respectively – yielding a row vector of data from each completed Survey Form that contained 29 elements, each numerically valued 1-5. The Survey Form is available for inspection under "HRI Study 1 Survey Form" at < http://starport.okcu.edu/SI/GS/ >.

The participants who completed Survey Forms for this study principally were OCU-affiliated (59% undergraduate or graduate students; 10% staff, faculty, or OCU alumni). Items of the Survey Form were designed to focus upon specific ethical and religious *views* (rather than calling for explicit declarations by participants of their identification, say, with the faith tradition of Islam or with deontological ethical theory). Nevertheless, the experimenters who administered the survey are confident that the population for the study was religiously diverse, including participants representing Christianity, Islam, Judaism, Native American spirituality, Hinduism, and Buddhism. Nearly equal numbers of men and women comprised the group of 75

participants who completed Survey Forms. Of these, 51% were in the age range 18-30, 16% in the range 30-50, and 33% were over age 50.

Data Analysis

Four of the 75 Survey Forms submitted by participants were not used during data analysis because they contained unmarked item choices. Accordingly, the input to our data analysis was a matrix consisting of 71 row vectors, each comprised of 29 elements. Numerical values for all of these elements were integers between 1 and 5, inclusive (representing, of course, choices made by a participant for each of the 29 items on the Survey Form).

The mathematical method that we selected for data analysis was consistent with our intention to conduct a *pilot study*. Indeed, this study was motivated by at least one *general* hypothesis of the kind already mentioned (i.e., there may be significant relations of ethical and religious views with acceptance of the subject technology), and this expectation did guide the design and content of our Survey Form. Prior to applying statistical inference procedures to test *specific* hypotheses, however, we considered it important to address the logically prior task of discerning *patterns* in participant responses that plausibly *recommend* specific hypotheses for testing in *follow-on* research. Accordingly, we believe it is quite consistent with the intentions of the present study that our analytic method of choice was to perform *cluster analysis* on the input data matrix. For that purpose we used MATLAB® 7.0 to perform K-means analysis.

Cluster analysis of this kind allows the user to specify in advance of execution a particular number, K, of "centroid" (or "prototype") vectors. The algorithm randomly generates K of these initial centroids, or (optionally) may randomly select K starting centroids from the input data matrix. It then compares, in iterative cycles, each of the centroids with all of the input data vectors. These comparisons employ a distance measure that also is specified by the user (e.g., Euclidean distance, City Block distance, etc.). The final output is a partitioning of the input data matrix into K clusters, for each of which the sum of the distances from its centroid to the input data vectors in that cluster is minimized. As its iterative execution proceeds, the algorithm adjusts the values of elements in the centroid vectors until it has established the optimal partitioning of the data just described. Inasmuch as the result can be sensitive to values that were randomly selected for the initial centroid vectors, it also is prudent for the user to execute the algorithm more than once to guard against reaching sub-optimal outcomes reflecting so-called "local minima." Run-time control parameters that we used for the analysis implemented this latter precaution, setting the 'replicates' parameter to 10. We chose MATLAB®'s default distance measure (Squared Euclidean distance) because it is faster than Euclidean

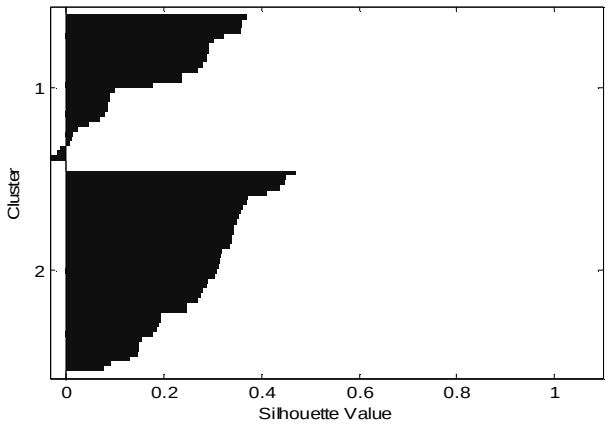

Figure 1. Cluster silhouette plots for K = 2.

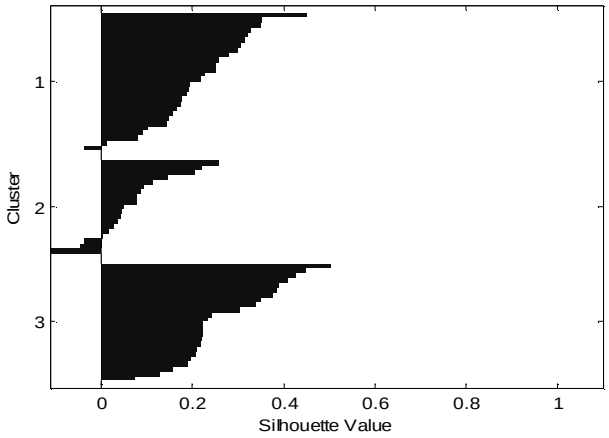

Figure 2. Cluster silhouette plots for K = 3**.**

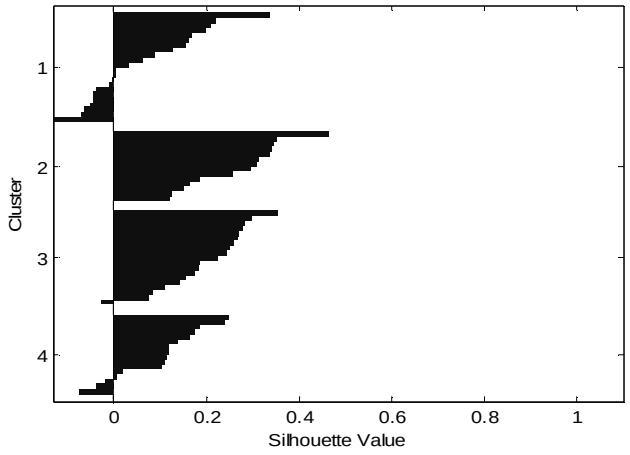

Figure 3. Cluster silhouette plots for K = 4**.**

distance and the substitution does not affect the output for K-means clustering; in addition, we allowed the default procedure of randomly selecting starting centroids from the input data matrix. Following common practice with this method of analysis, we began at K=2 and explored the effects of progressive increases in the K value.

Results

Data Clusters

Figures 1-3 display MATLAB®-generated "silhouette plots" for the first three of fifteen exploratory executions of the algorithm that we conducted (i.e, for K=2, K=3, and K=4). The vertical axis in each plot simply identifies the K clusters obtained for the value of K shown; hence, for example, we may refer to the silhouette plots of "Cluster 1" and "Cluster 2" in Figure 1. The computations determining shapes of the silhouette plots are explained concisely in MATLAB® documentation:

> The silhouette plot displays a measure of how close each point in one cluster is to points in the neighboring clusters. This measure ranges from +1, indicating points that are very distant from neighboring clusters, through 0, indicating points that are not distinctly in one cluster or another, to –1, indicating points that are probably assigned to the wrong cluster.

With this explanation of the silhouette plots in mind, inspection of Figures 1-3 discloses a pattern that we found as we increased values of K beyond K=2. Increasing values for K consistently resulted in deterioration of quality for the clusters, strongly recommending that our

TABLE I

CENTROID VECTOR VALUES FOR K=2. (Element values for successive parts of the Survey Form shown separately; column numbers at top correspond to those in the Figure 4 histogram below.)

Ethics	1	2	3	4	5	6	7	8							
Cluster 1	2.93	2.70	2.17	2.60	2.90	3.93	2.23	4.60							
Cluster 2	2.10	2.00	1.73	1.80	2.17	2.85	2.37	3.10							
Religion		9	10	11	12	13	14	15	16	17	18	19	20		
Cluster 1		4.30	4.10	4.20	3.30	3.97	4.07	3.23	3.77	3.80	2.47	1.90	3.90		
Cluster 2		3.12	2.29	3.98	2.02	2.27	2.12	1.59	2.71	3.05	1.85	1.98	3.32		
					Acceptance		21	22	23	24	25	26	27	28	29
Cluster 1							2.40	4.17	2.43	3.70	3.17	2.17	3.73	3.27	3.30
Cluster 2							1.93	3.41	2.73	3.71	2.78	2.22	4.24	3.85	3.98

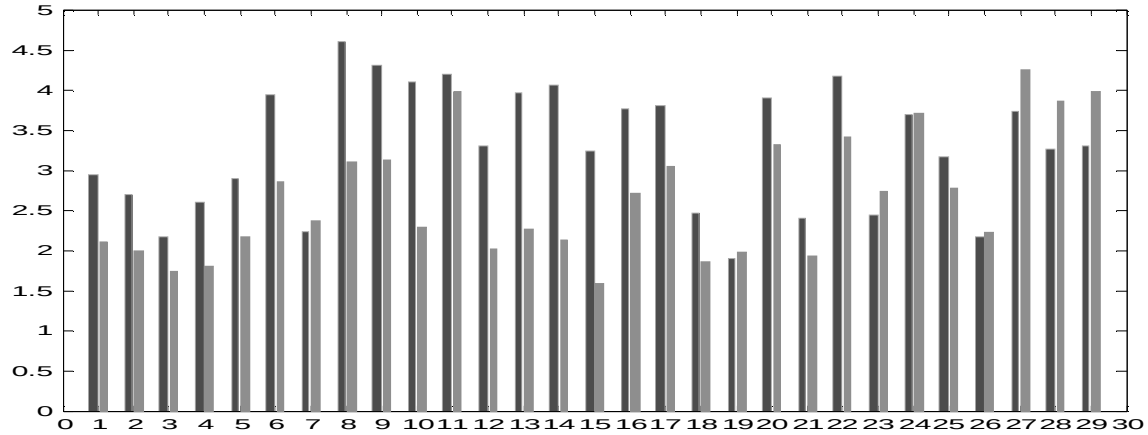

Figure 4. MATLAB®-generated histogram of centroid values for K=2. (Cluster 1 on left – Cluster 2 on right)

data set most appropriately is partitioned at K=2.

Data Cluster Centroids

For K=2, the MATLAB® software produced a listing of the 71-row x 29-column input data matrix, assigning each of the 71 row vectors either to Cluster 1 or to Cluster 2. As the cluster silhouettes in Figure 1 suggest, the number of row vectors (i.e., Survey Forms) assigned to Cluster 1 was somewhat less than the number assigned to Cluster 2. In fact, the ratio was 30/41, giving Cluster 1 about 42% of the 71 data vectors, and Cluster 2 approximately 58%.

The particular utility of cluster analysis for a pilot study such as ours is evident in its ability to help us discern *patterns* in the data – and an important part of this assistance is furnished by the resulting *centroid* vectors, for each cluster, that we may examine as mathematically appropriate representatives of the data vectors in their clusters. Data reduction of this kind makes the analytic task considerably more manageable. Table I and Figure 4 display the values of the centroids for Cluster 1 and Cluster 2, both numerically and graphically; together, they can help us discern patterns of interest.

The histogram in Figure 4, for example, immediately helps one notice an overall tendency of Cluster 1 survey responses to be "higher" in their values than Cluster 2 responses. In terms of the Survey Form, of course, this translates as a general tendency of Cluster 1 participants to be less affirming or approving in their responses – and, indeed, the numerical mean for the complete Cluster 1 centroid is 3.29 (vs 2.66 for the Cluster 2 centroid).

Recalling, however, that the Survey Form contained three main parts (concerning Ethics, Religion, and Acceptance), a closer inspection of *parts* of the centroids readily reveals some features of interest. First, if one calculates the means only for the *first 9 elements* of each centroid (representing the Ethics part of the survey) – and compares these "Ethics-means" with the overall means of the respective centroids – it turns out that this *part* of the centroid values is below the centroid mean for both clusters (9% below for Cluster 1 and 15% below for Cluster 2). However, if one performs a similar computation for elements 9 through 20 of both centroids (elements representing the Religion part of the survey) a very different pattern emerges (9% *above* for Cluster 1, and 5% *below* for Cluster 2). It seems, in translation, that Cluster 1 and Cluster 2 participants both tended to treat the Ethics items of the survey relatively more affirmatively – but their relative trend of responses in the Religion section diverged. The plot then thickens, as it were, when one performs the same calculations for the Acceptance part of the survey; here, we find an even stronger divergence – but in directions *opposite* to those in the Religion section (4% *below* for Cluster 1, and *20% above* for Cluster 2). In sum, it appears that, while participants in both clusters shared a tendency toward more affirmative answers in the Ethics part of the survey, Cluster 1 participants moved relatively toward more negative responses in the Religion part, and then toward more affirmative responses in the Acceptance part; Cluster 2 participants executed a similar flip-flop, but in the opposite directions. Thus far, of course, these are only patterns of interest in the general *kinds of responses* that participants in the two clusters displayed (particularly, in the Religion and Acceptance parts of the survey); an obvious next step in the analysis would be to examine *specific item responses* within each part of the survey.

For this purpose, we simply marked two extra copies of the Survey Form with answers representing those indicated (with rounding to integers) by the numerical values in both centroids for each of the 29 items in the survey (Table I). For the eight items in the Ethics part of the Survey Form, this straightforward procedure disclosed two that were the same in both clusters (columns 3 and 7, in Figure 4), five that differed between clusters by one choice option (columns 1, 2, 4, 5, and 6), and one item that differed by *two* choice options (column 8). Repeating this comparison for the Religion section displayed the greater divergence noted previously: three items the same in both clusters (columns 11, 18, and 19), six items differing by one choice option (columns 9, 12, 15, 16, 17, and 20), and three differing by *two* choice options (columns 10, 13, and 14). Notice that in moving from Ethics to the Religion part of the survey, the ratio of items differing by two choice options increased from 1/8 to 3/12 – a 100% increase. For the Acceptance part of the survey numbers of items displaying the same choices in both clusters increased (from 25% of the items in both Ethics and Religion to 56% of the items in the Acceptance part). However, of the remaining four items differing between clusters by one choice option in this last part of the survey, 75% (columns 23, 28, and 29 in Figure 4) showed participants from Cluster 2 less inclined to approve of the suggested robotic applications than participants from Cluster 1.

Although the foregoing computations and quantitative comparisons establish a useful "syntactic skeleton" of patterned inter-cluster differences, we must turn to the actual *content* of the numbered survey items to add the "semantic flesh" needed to disclose outlines of meaningful hypotheses. The following Discussion section will engage this task.

Discussion

The Ethics part of our Survey Form systematically probed agreement by study participants with a number of standard positions that one encounters in moral philosophy. Specifically, the Survey Form statements associated with columns 1 through 7 of Figure 4 simply furnished assertions, in layman's terms, of the main theses

of utilitarianism, deontological ethical theory, virtue ethics, prima facie ethical theory, intuitionism, absolutism, and noncognitivism. Item 8 probed the participant's agreement with the following claim: "Any authentic moral code must have its foundation in religious belief and/or experience." Item 8, one recalls, is precisely the one item in this part of the survey for which Cluster 1 responses differed from Cluster 2 responses by *two* choice options (Cluster 2 centroid approximating "I neither agree nor disagree, " and Cluster 1 falling closest to "I definitely disagree"). This result might also be conceptually consistent with one of the items for which inter-cluster centroid differences approximated *one* choice level. In particular, item 6 of the Survey Form offers the assertion "Some moral claims, such as 'It is wrong to tell a lie,' are true in an absolute sense – regardless of the historical period or culture in which they're asserted." If one inspects the first eight columns of the Figure 4 histogram, it is striking that, beyond the item 8 just mentioned, item 6 is the *only* other item in the Ethics part of the survey for which a response in *either* centroid rises notably above 3 – and it is, again, the *Cluster 1* centroid that approximates "I'm inclined to disagree." Whether, as these results at least suggest, an individual who strongly rejects resting moral codes upon religious foundations will also be inclined to reject ethical absolutism may be an hypothesis of interest for the context of this study.

The Religion part of our Survey Form resembled the Ethics part of the survey at least to the extent that it also probed specific and fairly common views associated with its topic. Formulation of its items, however, presented some special challenges, and probably some controversial choices. First, for example, we chose explicitly to avoid items dealing with identification of particular religious traditions or sects. The principal reason for this choice did not concern privacy issues (the Survey Forms, after all, were completely anonymous instruments); rather, it reflected the judgment that actual variance within nearly any religious tradition or sect appears to render highly problematic any presumed association of, say, identification with Christianity and acceptance of a specific doctrine or practice. On the other hand, the opening remarks of the present essay clearly betray our general hypothesis that there may, indeed, be some differences in the prevailing religious views of different populations that condition acceptance of robotic technology. Accordingly, we deliberately conducted most of our data gathering in a campus environment characterized by rich religious diversity – but we concurrently designed items of the survey instrument to focus upon specific views germane to religious life (some of which might well cross-cut different stated religious affiliations). Language also presents a problem in a project of this sort. Even among speakers of English, for example, the word "God" apparently can display a remarkably wide range of meanings. Recognizing these difficulties, we nevertheless chose to formulate a Religion part of the Survey Form, hoping that we shall learn from our mistakes (and believing that some empirical study of this class of variables in relation to attitudes toward robotic technology needs to be cultivated at this time).

The foregoing disclaimers having been issued, we can observe that the twelve items in the Religion part of the Survey Form (columns 9 through 20 in Figure 4) explored the participant's agreement with:

9 - the concept of an omnipotent God
10 - the concept of an engaged but not omnipotent God
11 - deism
12 - animism
13 - the concept of *imago Dei*
14 - divine creation of individual human souls
15 - existence of human souls that survive death
16 - reincarnation
17 - a behavioristic understanding of personhood
18 - human categorical freedom, relative to natural laws
19 - human categorical freedom, relative to divine will
20 – Laplacian determinism, relative to natural laws

Beyond the prior mention of a relative displacement of the centroid mean, in the Religion part of the survey, toward negative responses (for Cluster 1) and toward positive responses (for Cluster 2), we can now add "flesh" to the differences. In particular, a Cluster 1 response of "I'm inclined to disagree" appears for 8 of the 12 items in this part of the survey (specifically, items 9, 10, 11, 13, 14, 16, 17, and 20, in the above list). In contrast, a Cluster 2 response of "I'm inclined to agree" appears for 7 of the 12 items (viz., 10, 12, 13, 14, 15, 18. and 19).

Attending only to the column numbers in this contrast that are common to both clusters, we see a Cluster 1 disposition to reject (and a Cluster 2 disposition to accept): the concept of an engaged but not omnipotent God, the concept of *imago Dei*, and divine creation of human souls.

Contrasting response to the first of these (item 10) may seem to be ambiguous; the item wording ("God interacts with the world that we experience, but does not completely control every aspect of it") might be accepted or rejected for various reasons since it (probably unwisely) suggests two divine properties. However, item 10 also intentionally interacts with item 9, which asserts "God completely controls every aspect of the world that we experience." Noting that the Cluster 1 response to item 9 also is "I'm inclined to disagree," it would appear that the Cluster 1 rejection of item 10 is not dependent upon the question of the level of God's control. More plausibly, it turns on a rejection of the notion of a God interacting with the world *in any fashion*. This, of course, would be compatible with the attendant Cluster 1 disposition to reject both the concept of *imago Dei* and divine creation of human souls.

It should also be compatible with the Cluster 1 disposition to reject, as noted, deism and reincarnation (items 11 and 16). More fruitfully, though, this Cluster 1 pattern of limited rejection for a God interacting with the world, *imago Dei*, divine creation of human souls, and deism could *also* be reflecting the presence of practicing Buddhists in our sample. Indeed, some Buddhists could be expected to reject the term "reincarnation," as well, (favoring "rebirth"), allowing this interpretation to span all of the differences just noted.

The suggestion that Cluster 1 patterns in the Religion part of the survey could be reflecting Buddhist (and possibly Hindu) presence in our sample becomes more attractive when one considers a complementary interpretation for Cluster 2 – viz., that it may be reflecting the general profile of participants in our sample who identify with the Abrahamic religious traditions. Indeed, a plausible "Western" pattern of "I'm inclined to agree" responses may be detected in Cluster 2 for items 10, 13, 14, and 15 ("God interacts with the world that we experience, but does not completely control every aspect of it," "God created humanity in God's image," "God creates individual human souls," and "Each living human possesses a soul that survives death of the physical body"). On the other hand, the Cluster 2 centroid also shows an inclination to agree with "There are gods, or spirits, present everywhere in the world that we experience" (item 12), which might reasonably be viewed as an anomaly for this interpretation. However, some level of hospitality to animist beliefs probably can be found in nearly all of the major faith traditions.

Although we see plausibility in the foregoing "East – West" interpretation of cluster differences within the Religion part of the survey, some of the opinions cited at the outset of this essay should encourage us to expect a strong corresponding Cluster 1 acceptance of suggested robotic applications in the third part of the survey, relative to Cluster 2 – but our pilot survey data do not at least immediately support this expectation. Indeed, none of the cluster differences in this part involve changes of more than one level in the choice options, and we have observed that 56% of the "robotic application scenarios" suggested in the Acceptance part of the form actually exhibit identical choices for both clusters.

Nevertheless, our prior examination of mean centroid values for the three parts of the survey, when compared with the *overall* centroid means, displayed shifts *upward* (9% toward disagreement) in the Religion part but *downward* (4% toward approval) in the Acceptance part, for Cluster 1 responses. In sharp contrast, the movement that we noted for Centroid 2 was *downward* (5% toward agreement) in the Religion part and *upward* (20% toward disapproval) in the Acceptance part of the survey! Moreover, this computational evidence of differences between the cluster responses becomes interestingly more meaningful when we attend to the actual content of items in the Acceptance part of the survey. In particular, the last three items in the Acceptance part (columns 27, 28, and 29, Figure 4) present the following application scenarios:

27 Humanoid robots teach ethics classes at universities.
28 Android robotic children are leased to childless couples.
29 Android robotic spouses are leased to single humans.

Inspection of the Figure 4 histogram for these three columns shows a pattern of stronger disapproval of their scenarios among the Cluster 2 responses, relative to Cluster 1. We may recall, at this point, the prior observation that Cluster 1 differed by approximately *two choice options* from Cluster 2 in its rejection of the item 8 statement "Any authentic moral code must have its foundation in religious belief and/or experience." The correspondingly greater agreement among Cluster 2 participants with this suggested link between religion and moral codes may be viewed as compatible with Cluster 2's relatively stronger disposition to *disapprove* application of humanoid robots in college ethical instruction. And the coherence of the Cluster 2 choices with their prior survey answers becomes increasingly evident when we consider the semantics of items 28 and 29 in relation to the *preceding* Acceptance items. In fact, the proposed robotic application scenarios for items 21 through 27 uniformly involve robots essentially performing "jobs"; specifically, these preceding items propose robots serving as: companions for elderly people, soldiers, receptionists in business offices, nannies for children, caregivers for infirm people, butlers and/or housekeepers, and (finally) college ethics instructors. Items 28 and 29 suggest qualitatively different roles for the robots . Beyond simply performing tasks for us, these ending scenarios describe *intimate family relations* – the robots become accepted as our *children* and our *spouses*. Recalling that Cluster 2 participants endorsed, in the preceding Religion part of the survey, propositions such as "God created humanity in God's image," and "God creates individual human souls," (propositions with which Cluster 1 participants were inclined to *disagree*), it is understandable that Cluster 2 responses more strongly reject the intimate human-robot interactions proposed in items 28 and 29.

Conclusions and Future Work

In view of the prior explanation that we have intended HRI Study # 1 to be a pilot study from which hypotheses might be derived for testing in follow-on research, it should be appropriate now to state our conclusions from the study *in the form of* examples of hypotheses that we believe it recommends for investigation in future work:

1. Disapproval of HRI with life-like personal robots that requires human acceptance of the robots at intimate levels, such as family membership, will tend to be displayed significantly more often by individuals whose religious beliefs endorse the concept of a God that interacts with the world but does not control every aspect of it, is responsible for creating humanity *imago Dei*, and is responsible for creating individual human souls.

2. Disapproval of HRI with life-like personal robots performing jobs that presume the robot is capable of giving ethical instruction will tend to be displayed significantly more often by individuals whose religious beliefs endorse the concept of God described in Hypothesis (1), *and* who are not inclined to reject the claim that any authentic moral code must have its foundation in religious belief and/or experience.

3. Approval of HRI with life-like personal robots performing jobs serving human needs (except military applications and/or jobs that presume the robot is capable of giving ethical instruction) will not be significantly affected by whether the individual's religious beliefs endorse the concept of God described in Hypothesis (1).

We believe that the foregoing hypotheses warrant the investment needed to test them for at least two reasons. First, the hypotheses collectively cover a very wide range of potential commercial applications for life-like personal robots (many of which already have been undertaken, or are being suggested); determining population features that could either foster or discourage acceptance of these applications should have practical value for many business enterprises. Second, the hypotheses – particularly, if future research scientifically supports them – should contribute to improving interfaith understanding over a wide range of religious traditions; knowing clearly how differences in religious belief can affect attitudes toward emerging technologies of the kind we are considering may be expected to have genuine social value in our age of globalization. The co-authors of the present paper look forward to contributing to the future work, having collaborated on a number of related projects such as organization of a special "Robotics and Religion" session at the 2006 Metanexus Conference (see Lewis, Metzler).

Acknowledgments

We especially wish to thank John Branch, Molisa Derk, Kris Ladusau, and Karen McKellips for their kind assistance during data collection for this study.

References

Faiola A. 2005. "Japan embraces new generation of robots." *The Washington Post* 12 Mar. 2005. 31 Mar.2008 <http://www.detnews.com/2005/technology/0503/12/tech-114940.htm>.

Lewis, L. 2006. "Representations of the Brain and the Mind in Computer Science: Is There Room for Religious Concepts?" *Continuity + Change: Perspectives on Science and Religion*, Metanexus Conference, Jun. 3-7, 2006. 31 Mar. 2008 <http://www.metanexus.net/conferences/pdf/conference2006/Lewis.pdf>

Melson, G., P. Kahn, Jr., A. Beck, and B. Friedman. 2006. "Toward Understanding Children's and Adults' Encounters with Social Robots." *Technical Report WS-06-09*. AAAI Press, 2006: 36-42.

Metzler, T. 2006. "Don't Blame Me: I'm Just a Robot!" *Continuity + Change: Perspectives on Science and Religion*, Metanexus Conference, Jun. 3-7, 2006. 31 Mar. 2008 <http://www.metanexus.net/conferences/pdf/conference2006/Metzler.pdf>

Mori, M. 1999. *The Buddha in the Robot: A Robot Engineer's Thoughts on Science and Religion.* Tokyo: Kosei Publishing Company.

Nomura, T., T. Suzuki, T. Kanada, and K. Kato. 2006. "Altered Attitudes of People toward Robots: Investigation through the Negative Attitudes toward Robots Scale." *Technical Report WS-06-09*. AAAI Press, 2006: 29-35.

Tillich, P. 1995. *Morality and Beyond*. Louisville, KY: Westminster John Knox Press.

Tillich, P. 1971. *Systematic Theology.* Chicago, IL: The University of Chicago Press.