A Cultural Decision-Making Model for Virtual Agents Playing Negotiation Games

Elnaz Nouri and David Traum

Institute for Creative Technologies, University of Southern California
12015 Waterfront Drive, Playa Vista, CA 90292 USA

Abstract. We present a novel model of decision-making in social tasks for virtual humans. The model considers multiple valuations of the available choices in a decision set according to individual and social factors such as own utility, total group utility, and relative utility. Cultural differences are incorporated using Hofstede's dimensional model of culture and affect the decision making process by changing the different weightings of the factors. We have integrated the decision model into the dialogue manager of a virtual human system, and developed protocols and dialogue capabilities to support virtual humans in playing a simple negotiation game (Ultimatum Game). We present evaluations between both a culturally oriented virtual human and a person and between two virtual humans (with different culture models).

1 Introduction

Decision-making is an important part of social interaction. In the most general case, a decider must consider not only the impact on his own utility, but also the impact on others, including individuals, groups and society as a whole. There are also differences in how individuals value the options in a decision-space as well as broad similarities in outlook between similar individuals. In this paper we develop a model of decision-making that can be used by virtual humans to make decisions in a social setting, such as game play or negotiation.

Most classical economic game-theory accounts of decision-making (e.g. [10] look at a monolithic notion of utility and maximizing expected utility as the key to rationality. While such models are often used to good effect, they do not explain a number of aspects of actual human behavior in social situations [2, 7]. Behavioral game theory focuses more on modeling the way people actually make decisions [2]. Multi-attribute decision-making applies techniques for considering different sources of value and means for comparing them [3]. We use these techniques, but applied to social-decision-making, in which agents are concerned with not just their own utility, but also those of others, as well as relative and total utility. Such a model can naturally capture differences in individual or group behavior, by assigning different weights to the factors.

[9] uses Hofstede's model of culture [8] to inform the decision-making of virtual agents. We find this work inspiring, though limited in several respects, notably that the model only covers two dimensions, it is limited to only very

specialized types of decisions, and seems hard to generalize the utility equations. Moreover, it has not been implemented within virtual humans that can interact with people.

In this paper, we present a model of decision-making that can be used by virtual humans to make a wide variety of decisions in different contexts, including interacting with both other virtual agents, such as in [9], but also human interactors. In Section 2, we describe the general model, that can be specialized to different kinds of situations. In Section 3, we describe how this model can be adapted to cover cultural differences Using Hofstede's multi-dimensional model of culture [8]. In Section 4, we present a concrete example - the Ultimatum game, in which a fair amount of data has been gathered on human decision-making, and to which we adapt our model, calculating relevant aspects of cultural differences for valuing the situation. In Section 5, we show how this model has been incorporated into a virtual human dialogue model, allowing agents to play the ultimatum game with each other and with humans, using natural language dialogue interaction. Finally, in Section 6, we present a first evaluation of the model, looking at results of different culture-specific models used to drive virtual human interaction with other virtual humans and real people, in a dialogue context. We conclude with some directions for future work.

2 The General Social Decision-Making Model

Our basic model considers a number of different metrics for evaluating a given situation, even for something as simple as a division of money in an economic game such as the prisoner's dilemma [2] or the ultimatum game [5]. Each of the metrics can be calculated from a basic payoff matrix. The metrics we have considered include:

1. Self Interest (the agent's own utility)
2. Other Interest (the utility of another)
3. Total Utility (sum of individual utilities of all participants)
4. Average Utility (may not be derivable from Total Utility when the number of participants is variable)
5. Relative Utilities (viewed in several ways, such as self/total, self−other, self/other, self/average)
6. Minimum Utility (lower bound for any participant - the aim of Rawls' theory of justice [12])
7. Uncertainty (variation among possible outcomes)
8. For each of the above, minimum outcome versions, denoting the worst case rather than expected utility (in cases where payoffs are probabilistic rather than guaranteed for a particular choice).

Each of these metrics can be given one or more valuations, choosing an optimum point and scale. For example, the optimum for self/average might be infinity, when considering relative self-interest, 1 when considering fairness, or below 1, if trying to satisfy a vow of poverty. For a given set of possible choices, each valuation is considered and the values are brought into one scale by being centered by the mean and scaled down by the standard deviation. Each individual agent has a vector of weights, one per valuation, indicating the relative importance of that valuation. If the weight is zero, then the valuation is not considered in the overall utility computation for that individual. The total value for each choice is the sum of the product of values and weights for each valuation, as shown in (1). For every decision, our agent calculates the utility of all of its possible choices and selects the one that has the highest overall valuation (according to the agent's knowledge and ability to calculate or estimate these values).

(1) $Value(Choice_{\mathrm{i}}) = \sum_{j=1}^{n}(W_{\mathrm{j}} * V_{\mathrm{j}}(Choice_{\mathrm{i}}))$

The vector for a simple economic model would have weight one for self-interest zero for all other weights. A simple cooperative model would have all weight on total utility. An advantage of the multi-valuation approach is that it can model an agent who cares (possibly to different extents) about different aspects of the situation, such as self-interest, collective interest, and fairness. Our belief is that such a model can better reflect the kinds of decisions that people make considering multiple factors, which can also account for systematic differences in decisions made by different individuals (who would have a different weight vector), as well as subtle changes in the context of a decision, where changes are only present on some dimensions but not others.

3 Modeling Cultural Differences in Decision Making

The model in (1) can capture cultural differences among populations by setting the weights for different valuations to be congruent with the norms of specific cultures for those valuations. E.g. the relative value of utility given to self, a group, or the whole population. We use Hofstede's model of culture [8] as our basis for cultural modeling of decision-making because it has the following advantages:

- explicit dimensions of cultural norms that can be tied to valuation
- multiple ways in which cultures can be similar or differ
- data on dimension values for a large range of (national) cultures

Hofstede's model has five dimensions: Individuality (IDV), Power Distance (PDI), Long Term Orientation (LTO), Masculinity (MAS), Uncertainty Avoidance (UAI). In theory, each of these dimensions could contribute to the relative weight of any of the valuations. Thus our generalized model, shown in (2), breaks down the elements of the weight vector into one component per dimension, and thus an overall matrix of n valuations and m (= 5) dimensions.

(2) $Value(Choice_{\mathrm{i}}) = \sum_{j=1}^{n}((\prod_{d=\mathrm{IDV}}^{\mathrm{UAI}} W_{\mathrm{j,d}}) * V_{\mathrm{j}}(Choice_{\mathrm{i}}))$

In practice, however, not every dimension value will impact every valuation, so a number of these weights will be 1, and not modify the resulting valuation.

The issue then becomes how to assign the weights for individuals from particular cultures, in specific circumstances. Individuals could vary quite a bit from the cultural norms, however we would expect a weighted average of individuals from the culture to roughly match the norms for that culture. If we have no other information about the preferences of an individual agent, we can use the cultural norms of the society that the agent belongs to as an approximation of his cultural profile. If we have information about the personality, power position and the status of the membership of the agent to different groups, we can also take into account the variations it imposes on the weights of the valuation functions. The variation among agents also introduces uncertainty in prediction of how another agent may choose, even when following the above deterministic decision process.

Our initial weights are chosen by trying to match intuitions about the meanings of Hofstede's dimensions to the values that high and low points in the dimension would place on each of the metrics, under different circumstances. We outline those intuitions below. We attempt to evaluate and refine these intuitions by comparison of our predictions based on the model to observations of tendencies of different cultures to make different choices for each decision. In the future we intend to refine these valuations by using machine learning techniques to optimize the fit to known outcomes, and then evaluating against held out cultural data.

3.1 The Individualism-Collectivism Dimension (IDV)

The assumption is that in collectivistic societies where people have low individualism scores, people tend to care more about other individuals and give higher weight to group rather than self utility. The distinction between in-group and out-group is essential in collectivistic cultures [8], with value being placed on utility of others only within the group.

In our model, we define group IDs for individuals and each individual considers his in-group or out-group relationships when he wants to calculate the utility function of the decision. High individualism cultures put high weight on self-utility, while low individualism puts less weight on self utility. Low individualism cultures put high weight on other utility for group members, and high weight on total utility, when everyone considered is in the group.

3.2 Power Distance Index (PDI)

Power distance is the tendency to accept that more powerful individuals should have more resources. In a low power-distance culture, more weight is put on fairness (minimum utility and average utility), regardless of the status of the participants. For high power distance cultures, the value of self or other is proportional to the power or status of the individual, and the ideal value for relative utility is in accordance with the relative power.

3.3 Masculinity (MAS)

The masculinity index in general refers to differentiation among gender roles. However it is also correlated with higher assertiveness and competitiveness, especially for male members of the culture. For low masculinity cultures, the values are more similar across gender roles, and tend more toward caring and general well-being. In terms of our valuations, high masculinity cultures have higher weights for relative utility (self/average) and self utility, while low masculinity cultures have higher weights for other utility, average utility and minimum utility.

3.4 Uncertainty Avoidance Index (UAI)

The Uncertainty avoidance indices in societies show how much people do not want to tolerate uncertainty and ambiguity. This dimension of the culture brings up another example of violation of classic game theory measurement of utility. [1] showed that many people would prefer a risk free decision over a possible decision with higher pay off but uncertainty involved. We capture this aspect by looking not at expected utility of each valuation, but at a more qualitative representation of the decision-space, looking at the range of different possible outcomes that are not under the control of the player. In the simplest case, this degenerates to looking at the worst-case options.

3.5 Long-Term Orientation (LTO)

Long term orientation can be viewed as different ways of assigning utility, depending on the time scale of the utility and discount rate for future payoff. It can also be viewed as the tendency to view a set of decisions together as a policy toward an ultimate goal, rather than looking at each decision in its own right. We defer modeling LTO for the time being, focusing first on individual decisions with simple utility values, without taking time into account.

4 Example Application: The Ultimatum Game

In the ultimatum game two people have the opportunity to split a certain amount of money. The first player makes an offer and second player can accept or reject this offer. If the offer is accepted the money is split according to the proposal. If the offer is rejected, both get nothing. This stylized negotiation was first studied in [5]. Analytical game theory predicts that the responder should accept any offer greater than 0. In reality many people from most modern cultures would turn down an offer as low as \$10 (out of \$100) in order to punish the "unfair" proposer [7] and [2].

In the case of the ultimatum game, given the simplicity of the task and symmetries coming from splitting a fixed amount with one other player, we use only the following four metrics: Self Utility, Other Utility, Self/Other Utility, and Lower Bound. Decisions are made according to calculations of (2), with one weight for each combination of valuation and cultural factor

In the rest of this section we give initial weights for low and high values of the dimensions, under different circumstances. In all cases, for a given dimension d, weights are shown as a vector for that dimension: ($W_{self,d}$, $W_{other,d}$, $W_{self/others,d}$, $W_{lowerbound,d}$).

4.1 Impact of Individualism-Collectivism Dimension

(3) shows our initial weights for High individualism and for collectivism, considering both in-group and out-group partners.

(3)

Low IDV In-group	Low IDV Out-group	High IDV
(2, 2, 1, 4)	(1, 1, 1, 1)	(2, 1, 1, 1)

4.2 Impact of Power Distance Dimension

In the context of the ultimatum game, the assumption is that in societies with low PDI, both parties would expect a more even split, while in high power distance societies, it would be natural to allocate more to the more powerful party. Our initial weights for this dimension are showin in (4), considering the power relationship as well as the cultural dimension value.

(4)

PDI score	Low PDI	High PDI
High to Low power position	(2, 1, 2, 2)	(2, 1, 2, 1)
Low to High power position	(1, 2, 1, 4)	(2, 4, 1, 1)

4.3 Impact of Masculinity Dimension

Our model of the masculinity dimension contrasts fairness and care for everyone vs competitive self-interest. It is summarized (5).

(5)

Low MAS	High MAS
(2, 2, 1, 4)	W = (4, 1, 4, 1)

4.4 Impact of Uncertainty Avoidance Dimension

The main source of uncertainty in the Ultimatum game is only for the proposer, considering how likely the responder is to accept. Proposers with a high uncertainty avoidance score, would try to come up with decisions that would minimize the chance of rejection. The probability of acceptance is calculated by looking at the likelihood of a random responder accepting, considering the decision that each cultural model would make and the frequency of each type of culture assumed to be in the population of players. For low uncertainty avoidance, expected utility is used. For high uncertainty avoidance, we assume a deterministic decision based on how the majority would decide.

5 Integration with Virtual Humans

We use the model developed above to control virtual human decision-making in playing the ultimatum game. We integrated an implementation of the decision-making protocol described above within the tactical questioning architecture [4]. The authoring environment was used to construct domain knowledge and textual realizations for natural language understanding and generation of a range of speech acts.

Two agents, were also developed, each of which can play roles of proposer and the responder in the context of the ultimatum game. The dialogue model was extended so that for certain situations, the dialogue manager would consult the decision-making module before deciding on which offer to give (for proposer) or whether or not to accept an offer (for responder). Other dialogue moves (e.g., greetings, closings, explanations) were handled by the existing dialogue networks in [4]. For each character we can use any of the possible cultural models, and instantiate relationship variables, such as whether the players are from the same group, the relative power status, and the population set considered for calculating probability of refusals.

6 Evaluation and Conclusions

We evaluate the culture models by testing virtual humans playing against humans and virtual humans. In our experiments we ran the single shot ultimatum game between two players, splitting $100 with all offers being multiples of $10. We also played the games in rounds of 10 games each, so that a context of total earnings can be built up between games to be used for calculating relative utility.

6.1 Evaluation: Virtual Humans versus Humans

We first evaluate the virtual human dialogue integration by having the virtual humans play against human players. Our virtual humans were tested against humans playing both proposer and responder roles. The goal was to demonstrate that the system can play the game successfully, including recognizing a range of human utterances as needed to participate. An example dialogue is shown in (6). Three different users acted as both proposer and responder against a range of cultural models. The average dialogue was composed of 3 dialogue acts and 3 exchanges of turns. An average of 2.5 offers were made by the user proposer before they were accepted by the virtual human responder. Our initial results are successful in that humans can play the game by engaging in unrestricted (text) natural language for the domain. In future work we plan a more controlled study to quantify performance levels.

(6)

Player	Sample Utterances
VH Proposer	"Would you accept and offer of ten dollars?"
H Responder	"I reject your offer of ten dollars."
H Proposer	"I offer you thirty dollars."
VH Responder	"I accept thirty dollars."

6.2 Evaluation of Culture Models

We evaluate the culture models by testing the virtual humans playing against other virtual humans. We included all possible configurations (low, high) of the cultural dimensions and under four different conditions considering in group and out-group status crossed with lower and higher power. The goal was to compare cultural models with data from human players from those cultures.

In order to evaluate the result of the experiments done so far and the performance of the model we referred to the extensive data available from literature. [2] provides a detailed of the history and data of the different ultimatum bargaining games experiments.[6] reveals more variability across cultures and along with [11] provides the baseline data for our model. The average amount of money offered in all the experiments done so far is between 0.26 (for Machiguenga in Peru) and 0.58 (in Lamelara in Indonesia). Our model's VH proposed values also fall into this range of data. According to [2] the rejection rate and rejection data vary from 0 percent (in Tsimane', Bolivia) to 67 percent (in Mapuche in Chile). A summary of some of the results is shown in (7), along with data that has been reported for human players of these cultures.

(7)

Culture	Hofstede: PDI,IDV, MAS,UAI	VH mean offer	Human mean offer	VH rejection rate	Human rejection rate
Israel	13,54,47,81	$21.66	$41.71	25.0%	17.7%
Japan	54,46,95,92	$32.5	$44.73	1.0%	19.3%
Chile	63,23,28,86	$33.13	$34.0	1.0%	6.7%
Austria	11,55,79,70	$33.13	$39.21	9.1%	16.1%
Ecuador	78,8,63,67	$33.13	$34.5	1.0%	7.5%
Germany	31,64,61,60	$36.88	$36.7	9.1%	9.5%
US	40,91,62,46	$41.88	$42.25	12.0%	17.2%
Spain	57,51,42,86	$28.75	$26.66	25.0%	29.2%
Sweden	34,70,4,26	$33.13	$35.33	10.2%	18.2%

These initial results are encouraging. With the exceptions of Israel and Japan, the VH mean offers are quite close to the human means, and differences are tending in the right directions (higher for US, lower for Sweden). Similarly for rejection rates, the rough magnitudes and orderings are consistent for VH and observed humans.

Clearly there is still some more work to be done, however, to better account for the data. There are many ways in which the model can be changed, while still keeping the basic idea of different weights on different valuation functions in an overall decision-making process. First, we can look at improving the individual weights used, from those reported in Section 4. These were based mostly on intuition, applying on reasoning about the descriptions of the dimensions. However in many cases this is just guesswork - particularly when it comes to questions of magnitudes of the weights rather than the qualitative differences. We intend to apply machine learning techniques to try to learn optimal weights for a set of cultures, and then evaluate on held-out cultures, to see if this may provide more robust weights. Second, it may be the case that a different set of valuation functions must be considered for the Ultimatum game - we may be missing some important elements, at least for some cultures. Third, we intend to refine the model of Hofstede's dimensions, taking into account not just whether the culture scores low or high, but making the weight dependent on the actual values, to capture finer distinctions. Finally, it may be that the Hofstede dimensions themselves are not adequate for capturing the kind of cross-cultural variation that has been observed.

7 Acknowledgements

This research was funded by a MURI award through ARO grant number W911NF-08-1-0301.

References

1. Allais, M.: Le comportement de l'homme rationnel devant le risque: Critique des postulats et axiomes de l'ecole americaine. Econometrica 21(4), pp. 503–546 (1953)
2. Camerer, C.F.: Behavioral game theory - Experiments in strategic interaction. Princeton University Press (2003)
3. Fishburn, P.C.: Utility theory. MANAGEMENT SCIENCE 14(5), 335–378 (1968)
4. Gandhe, S., DeVault, D., Roque, A., Martinovski, B., Artstein, R., Leuski, A., Gerten, J., Traum, D.R.: From domain specification to virtual humans: an integrated approach to authoring tactical questioning characters. In: INTERSPEECH. pp. 2486–2489 (2008)
5. Güth, W., Schmittberger, R., Schwarze, B.: An experimental analysis of ultimatum bargaining. Journal of Economic Behavior & Organization 3(4), 367 – 388 (1982)
6. Henrich, J., Boyd, R., Bowles, S., Camerer, C., Fehr, E., Gintis, H., McElreath, R., Alvard, M., Barr, A., Ensminger, J., Henrich, N.S., Hill, K., Gil-White, F., Gurven, M., Marlowe, F.W., Patton, J.Q., Tracer, D.: in cross-cultural perspective: Behavioral experiments in 15 small-scale societies. Behavioral and Brain Sciences 28(06), 795–815 (2005), http://dx.doi.org/10.1017/S0140525X05000142
7. Hoffman, E., McCabe, K., Smith, V.: On expectations and the monetary stakes in ultimatum games. International Journal of Game Theory 25, 289–301 (1996)
8. Hofstede, G.H.: Culture's consequences:comparing values, behaviors, institutions, and organizations across nations. SAGE, Thousand Oaks, CA (2001)
9. Mascarenhas, S., Dias, J., Afonso, N., Enz, S., Paiva, A.: Using rituals to express cultural differences in synthetic characters (2009)
10. Neumann, J.V., Morgenstern, O.: Theory of Games and Economic Behavior. Princeton University Press (1944)
11. Oosterbeek, H., Sloof, R., van de Kuilen, G.: Cultural differences in ultimatum game experiments: Evidence from a meta-analysis. Experimental Economics 7, 171–188 (06 2004), http://dx.doi.org/10.1023/B:EXEC.0000026978.14316.74
12. Rawls, J.: Some reasons for the maximin criterion. The American Economic Review 64(2), 141–146 (1974), http://www.jstor.org/stable/1816033