Identifying Culture and Leveraging Cultural Differences for Negotiation Agents

Brian Ziebart[1], Miroslav Dudík[2], Geoff Gordon[1], Katia Sycara[1], Wendi Adair[3], and Jeanne Brett[4]

[1]Carnegie Mellon University; 5000 Forbes Ave. Pittsburgh, PA 15213, USA
[2]Yahoo! Research; 111 West 40th Street, 17th Floor New York, NY 10018, USA
[3]Univ. of Waterloo; 200 University Avenue West; Waterloo, ON N2L 3G1, Canada
[4]Northwestern University; 2001 Sheridan Rd, Evanston, IL 60208, USA

bziebart@cs.cmu.edu, mdudik@yahoo-inc.com, ggordon@cs.cmu.edu, katia@cs.cmu.edu, wladair@uwaterloo.ca, jmbrett@kellogg.northwestern.edu

Enabling computational agents to efficiently aid or automate negotiations with humans requires a recognition and understanding of cultural differences in negotiation behavior. In this work, we investigate two tasks needed to enable such agents: (1) the identification of negotiators' cultures from observed negotiation transcripts; and (2) the prediction of culturally-based differences in negotiation outcomes.

Keywords: negotiation, culture recognition, outcome prediction, Markov models, maximum entropy regularization

1 Introduction

Negotiators have mixed motives: they must cooperate to reach jointly-acceptable agreements, while also competing to maximize personal gains (Walton and McKersie 1965). Multiple stages of negotiation have been proposed to explain aspects of the transitions from competitive persuasion and posturing interactions to information-sharing and cooperative deal-making behavior over the course of the negotiation process (Pruitt 1971, 1981, Putnam and Jones 1982, Morley and Stephenson 1977, Adair and Brett 2005). Cultural differences within these stages have been identified based on e.g., high- and low-context (Hall, 1976), differences in information exchange (Brett and Okumura, 1998) and the effects of early offers on final negotiation outcomes (Adair et al., 2007).

We are motivated to understand these cultural differences in negotiations from a computational agent's perspective. Namely, we are interested in whether and to what degree a computational agent can identify the cultures (and, more generally, other latent attributes) of negotiators based on negotiation patterns to assist in ultimately arriving at better negotiation outcomes. This problem is very different from the traditional negotiation theory-building research objectives. From a statistical estimation perspective, the task of identifying negotiators' cultures from a single negotiation transcript is significantly more challenging than establishing cultural differences from a large pool of negotiation transcripts due to in-group variability. Thus, we focus on applying more sophisticated statistical estimation and machine learning techniques to aid in this task.

Though challenging, progress in this direction could be fruitful. For instance, automated negotiation agents could play a role in enabling automated customer care phone systems to more efficiently and amicably negotiate resolutions to consumer complaints, or help computer systems for translation also provide culturally-sensitive guidance for negotiators in international business deals. Successful inter-cultural negotiations are often characterized by the negotiators' abilities to adapt to the norms of other cultures (Adair et al. 2001). Computational agents could, in general, help assist in this adaptation to cultural norms in negotiation settings.

In this paper, we analyze a collected data set of sequential negotiation interactions between dyads of different cultural combination. We focus our analysis on two tasks: (1) identifying the culture of negotiators based on their interactions; and (2) understanding cultural differences in the outcomes of issues being negotiated. Together, these tasks are important components for empowering agents to reach better mutually agreeable outcomes based on cultural differences. While our investigation in this work focuses solely on the cultural dimensions of negotiation differences, the techniques can be generally applied to infer more general latent information about negotiators, such as personality types, to help efficiently reach better negotiation outcomes.

2 Negotiations and Cultural Difference Analysis

2.1 Collected Negotiation Data and Encoding Scheme

We employ a cartoon negotiation scenario developed by Brett and Okumura (1998), which was extended from an earlier scenario by Tenbrunsel and Bazerman (1995). In this scenario, a film production company sales representative and a television station manager negotiate the sale and licensing of a 100-episode children's cartoon series, *Ultra Rangers*. There are four primary issues in this scenario:

- The number of broadcasts per episode (runs);
- The down-payment amount and schedule of payments (financing);
- The inclusion of an additional cartoon, *Strums* (strums); and
- The price per episode (price).

The sales representative and the station manager have different utilities for combinations of agreed-upon outcomes for these four issues as shown in Table 1.

Table 1 -- Utilities for the negotiating parties for each issue

Issue	TV Station	Film Company
Revenue ($)	8,400,000	
Price/episode ($)		
Limit	60,000	35,000
Aspiration	30,000	70,000
Runs Per Episode Adjustment ($)		

4	-1,680,000	500,000
5	-840,000	250,000
6	0	0
7	840,000	-250,000
8	1,680,000	-500,000
Financing Savings or Cost (%)		
Year 1	10	-20
Year 2	20	-35
Year 3	30	-50
Year 4	40	-60
Year 5	50	-70
Strums (second comic)		
Reservation ($)	20,000	10,000
Rating – estimated likelihood (%)		
6-7	20	10
7-8	50	10
8-9	10	10
9-10	10	50
10-11	10	20
Alternative Deal ($)		
Value	3,000,000	2,500,000

Additionally, each negotiator has a different belief in what ratings *Ultra Rangers* will receive upon airing. Negotiators can incorporate the differences of their beliefs into their deal by e.g., providing bonus compensation if the ratings exceed a specified threshold. Lastly, each negotiator has an alternative deal value. Any potential deal with a lower value than the alternative deal's value should be avoided by rational negotiators.

Transcripts from 211 negotiations were collected from this cartoon negotiation scenario between 12 different dyad culture combinations. Agreed-upon issue outcomes and the step-by-step interactions leading to those outcomes were encoded from these transcripts. The interaction encoding scheme contains 34 types of interactions that are categorized into 9 higher-level interaction types. These are detailed in Table 2.

Table 2 - Transcript encoding schema

INFORMATION

Code	Definition
I-1 (101)	Preference for a negotiable issue, option, relative importance of issues; assertion of interest
I-2 (102)	Reference to minimum acceptable price or conditions (reservation price) [implicit offer here]
I-3 (103)	Reference to BATNA (Best-Alternative to Negotiation Agreement - what we do if we don't reach an agreement)
I-4 (104)	Reference to or preference for Multiple Issues with or without Trade-offs; usually "if...then..." or "..and..."
I-6 (106)	Information about product or non-negotiable issue (ratings, # episodes, story, characters, sponsor, tie-ins, Strums, profitability)
I-7 (107)	Information about competitors (other stations, other cartoons or shows, other suppliers)
I-8 (108)	Information about own company (strategic plan, profitability, long-term relationships, reputation, power)
I-10 (110)	Reference to personal stake of negotiator in transaction
I-11 (111)	Other information (ratings of current show in time slot)

QUESTIONS

Code	Definition
Q-1 (201)	About preferences for negotiable issue, option, relative importance of issues, offers
Q-2 (202)	About minimum/maximum acceptable price or conditions (reservation price) [will often also be coded as offer]
Q-3 (203)	About BATNA
Q-6 (206)	About product or non-negotiable issue (ratings, # episodes, story, characters, sponsor, tie-ins, profitability)
Q-7 (207)	About competitors (other stations, other cartoons or shows, other suppliers)
Q-8 (208)	About company (strategic plan, profitability, long-term relationships, reputation, power)
Q-10 (210)	About personal stake of negotiator in transaction
Q-11 (211)	Question about external information

SUBSTANTIATION/ARGUMENT/PERSUASION

Code	Definition
S-1 (301)	Substantiation (you do this/good for you/because how affects you, your company) or Argument/Persuasion (we need/because why; informational persuasion) [reference to something positive for either me, you or us]
S-2 (302)	Sympathy (you do this/good for you/because how affects me, my company) [reference to something negative for me]
S-3 (303)	Argument (you don't need/because..); threats [reference to something negative for you]

OFFERS

Code	Definition
O-1 (401)	Single issue offer or counter-offer
O-2 (402)	Multiple issue offer without trade-off (often phrased "...and...")
O-3 (403)	Multiple issue offer, with trade-off (often phrased "if...then...")

REACTIONS

Code	Definition
R-1 (501)	Positive or neutral reaction (vague, ideas, arguments); Positive acceptance of offer; affirming what other said.
R-3 (503)	Negative reaction

MUTUALITY

Code	Definition
M-1 (601)	Noting common or mutual interests (this is a tactic, likely at beginning of negotiation; positive)
M-2 (602)	Noting differences (negative)

PROCEDURAL COMMENTS

Code	Definition
P-1 (701)	Comments regarding procedures to be used, or in use EXCEPT reciprocity (issue by issue packaging, moving on without resolving issue, reopening issue already resolved, making compromises)
P-2 (702)	Comments regarding process of reciprocity (not used very often)
P-4 (704)	Positive expectations about negotiation process or outcome
P-6 (706)	Limits of case information; sources of case information
P-7 (707)	Time out to calculate, think or break

JUNK

Code	Definition
J-1 (801)	Junk; uncodable

CONFIRMATION/QUALIFICATION

Code	Definition
C-1 (901)	Question or response to question for clarification, repetition; clarification of offer; summarizing or paraphrasing; stating "that's not what I meant"

Additionally, each interaction has a more compact encoding that uses a prioritized "OPRAX" scheme based on whether the interaction is an offer (O), priority (P), rational (R), affective (A), or other categories (X).

2.2 Identifying Culture from a Single Negotiation

Our initial two analyses illustrate the relative difficulties of recognizing the culture of negotiators based on the frequencies of the different types of interactions. On average, there are 287 interactions per negotiation transcript and 17.6 transcripts per dyad culture combination. The sample means and 95% confidence intervals for those means of the relative frequencies of each interaction type (i.e., the percentage of the

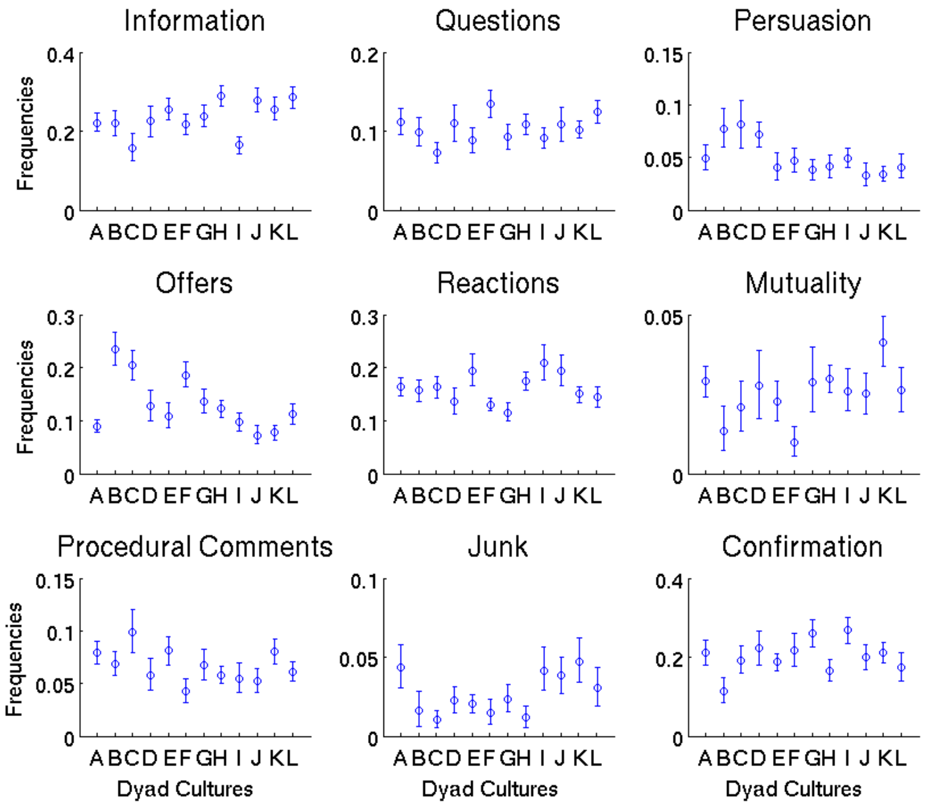

Figure 1: The group mean and 95% confidence intervals for nine interaction type frequencies for each of the 12 dyad culture combinations: American-American (A), Japanese-Japanese (B), Russian-Russian (C), Chinese-Chinese (D), Swedish-Swedish (E), Thai-Thai (F), Israeli-Israeli (G), German-German (H), American-Japanese (I), American-Chinese (J), American-Israeli (K), and American-German (L).

interactions that belong to that type) for each dyad culture combination are shown in Figure 1. We employ a 5000-simulation bootstrap procedure to establish these confidence intervals for the measures we analyze. For each simulation, the procedure samples transcripts with replacement from the available dyad culture combination transcripts and the 95% range is reported based on the distribution from the combined simulations' mean frequency values.

Each point and confidence interval in the plots of Figure 1 corresponds to one particular dyad culture combination. We note that a number of statistically significant differences exist in these frequencies between cultures. For example, in the first plot, Japanese-Japanese and American-Japanese negotiations include significantly fewer information-based interactions than e.g., American-Chinese, American-Israeli, and American-German negotiations.

While there are significant cultural differences in the means of the frequencies of these types of interactions across a population of dyads, identifying the culture combination of a particular transcript is difficult due to the within-group variability that exists. Figure 1 illustrates this by showing the middle 95% range of the interaction frequency types.

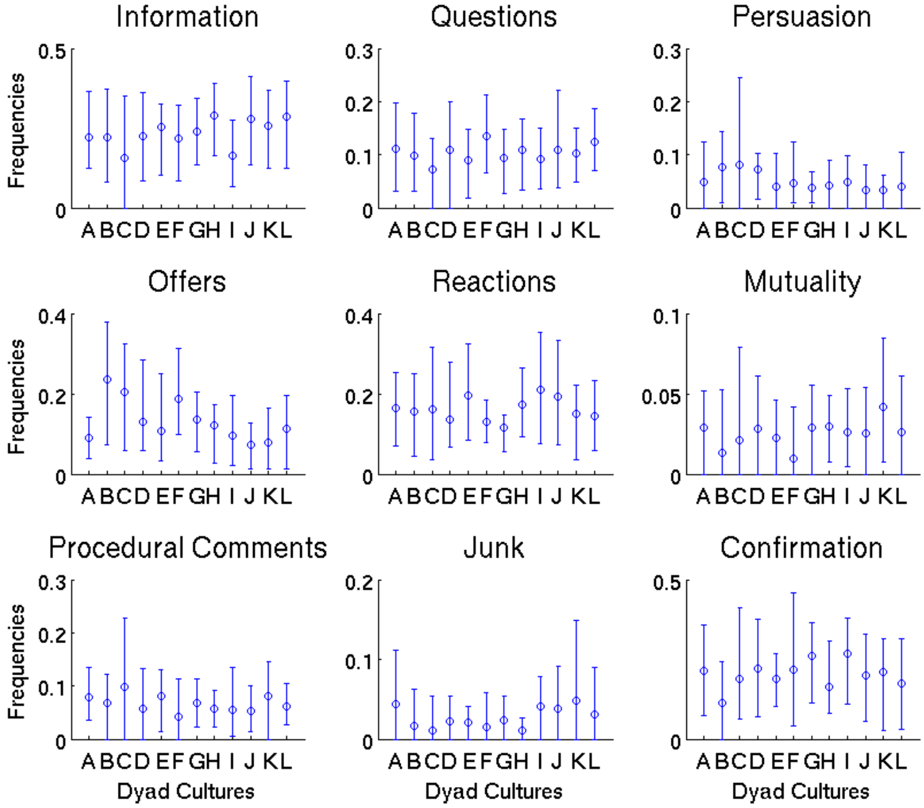

Figure 2: Range of in-group interaction type frequencies (middle 95%) for each of the 9 interaction types and 12 dyad culture combinations: American-American (A),

Japanese-Japanese (B), Russian-Russian (C), Chinese-Chinese (D), Swedish-Swedish (E), Thai-Thai (F), Israeli-Israeli (G), German-German (H), American-Japanese (I), American-Chinese (J), American-Israeli (K), and American-German (L).

As shown by Figure 2, significant overlap in these in-group ranges exists between the different cultures. This makes distinguishing the culture combination of the participants in single transcript much more difficult than distinguishing the culture combination of participants from a large number of transcripts.

More sophisticated models for recognizing the culture of negotiators leverage not only the frequencies of interaction types, but also the sequential process by which they are generated. A Markov model of K-th order represents the statistical relationships between the K previous interactions (and the cultures of negotiators) and the next interaction in the negotiation based on the conditional probability distribution:

$$P(\mathrm{act}_T|\mathrm{act}_{T-1}, \ldots, \mathrm{act}_{T-K}, \text{dyad cultures C}). \qquad (\mathbf{1})$$

Rather than dividing the transcript into phases (Pruitt 1971, 1981, Putnam and Jones 1982, Morley and Stephenson 1977, Adair and Brett 2005), this approach considers transcript sub-sequences of size K to base cultural reasoning. This difference can be advantageous in multi-issue negotiations where negotiation phases are sequentially repeated for multiple issues.

Given a portion of a transcript (i.e., a temporal sequence of interactions), a belief of negotiators' cultures can be obtained from this model using Bayes' rule:

$$P(\text{dyad cultures C}|\mathrm{act}_1, \ldots, \mathrm{act}_T) \qquad (\mathbf{2})$$

$$= \frac{\prod_{t=1}^{T} P(\mathrm{act}_t | \mathrm{act}_{t-1}, \ldots, \mathrm{act}_{t-K}, \text{dyad cultures C})}{\sum_{C'} \prod_{t=1}^{T} P(\mathrm{act}_t | \mathrm{act}_{t-1}, \ldots, \mathrm{act}_{t-K}, \text{dyad cultures C'})}.$$

Given the conditional distribution of actions (Equation 1), this Bayesian belief of dyad cultures can be directly calculated for any amount of transcript data.

However, the naïve approach for constructing this model – directly estimating the conditional probability distribution (Equation 1) from previously observed negotiation interactions – runs the risk of overfitting the model to the available negotiation data. This is especially true when the number of possible interactions, the number of cultures, and/or the order of the Markov model grow large relative to the amount of available data. Exponential family distributions based on the principle of maximum entropy (Jaynes, 1957) reduce the parameterization of the conditional probability distribution by using the following form:

$$P(\mathrm{act}_T|\mathrm{act}_{T-1}, \ldots, \mathrm{act}_{T-K}, \text{dyad cultures C}, \theta) \qquad (\mathbf{3})$$

$$= Z^{-1} \exp\{\theta^T F(\mathrm{act}_T, \mathrm{act}_{T-1}, \ldots, \mathrm{act}_{T-K}, \text{dyad cultures C})\},$$

where Z is a normalization constant that ensures that the action probabilities sum to 1.0:

$$Z = \sum_{act_T} \exp\{\, \theta^T \mathrm{F}(\mathrm{act}_{\mathrm{T}}, \mathrm{act}_{\mathrm{T\text{-}1}}, \ldots, \mathrm{act}_{\mathrm{T\text{-}K}}, \mathrm{dyad\,cultures\,C})\}. \qquad \textbf{(4)}$$

The parameters of this model, θ, are typically estimated by maximizing the (log) likelihood of the available data.

$$\theta^* = \mathrm{argmax}_\theta \sum \log \mathrm{P}(\mathrm{act}_1, \ldots, \mathrm{act}_\mathrm{T} \,|\mathrm{dyad\ cultures\ C}, \theta) - \mathrm{R}(\theta), \qquad \textbf{(5)}$$

where the sum is over training example transcripts: act_1, …, act_T, and dyad cultures combinations C.

To avoid overfitting to a small amount of available data, i.e., fitting available data well, but fitting data drawn from the same underlying distribution poorly, a *regularization* penalty, R(θ), can be employed. This penalty term can prevent the values of θ from growing too large, which corresponds to overfitting. We additionally employ a form of regularization that learns general culture-independent parameters and then penalizes the differences of individual culture combinations from those culture-independent values. We employ a maximum entropy-based regularization technique (Dudík, Phillips, and Schapire, 2007) using various combinations of available interaction characteristics with results shown in Figure 3.

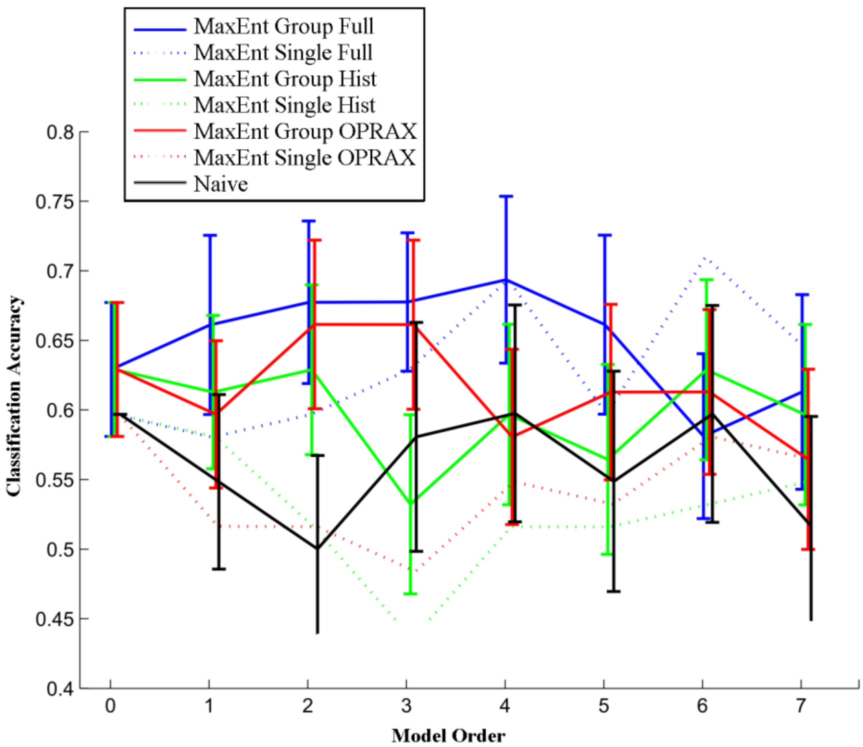

Figure 3: Cross-validated culture classification accuracy rate of K-order Markov models with Japanese, Chinese, and German negotiations.

We evaluate a number of different models for various order sizes using entire negotiation transcripts. The resulting classification accuracy is shown in Figure 3 with 95% confidence intervals using 10-fold cross-validation (i.e., splitting the available data into 10 parts and evaluating each part using the remaining 9 parts to fit the model parameters). The difference models employed are:

- **Maxent Group full:** maximum entropy distribution based on culture, low- and high-level interaction types and OPRAX codes with group regularization
- **Maxent Single full:** maximum entropy distribution based on culture, low- and high-level interaction types and OPRAX codes with independent model for each culture
- **Maxent Group hist:** maximum entropy distribution based on culture and low-level interaction types with group regularization
- **Maxent Single hist:** maximum entropy distribution based on culture and low-level interaction types with independent model for each culture
- **Maxent Group OPRAX:** maximum entropy distribution based on culture and OPRAX encoding with group regularization
- **Maxent Single OPRAX:** maximum entropy distribution based on culture and OPRAX encoding with independent model for each culture
- **Naïve Cult:** frequency count distribution estimated based on culture and OPRAX encoding

By incorporating maximum entropy regularization, there is a general trend: a larger amount of available information can be employed to more accurately predict culture. Additionally, we find that 2-order through 4-order Markov models with this regularization tend to perform the best. However, because of the small sample sizes, establishing statistical significance between these different models for any particular order size is difficult.

2.3 Negotiation Outcome Analysis

Differences in negotiation outcomes based on culture (and other unobserved negotiator characteristics) make the dyad culture recognition task of Section 2.2 important for efficiently reaching acceptable agreements. For example, if a negotiation assistance agent is able to recognize the cultures of negotiating dyads, it can guide them towards negotiation outcomes that have been reached in the past for that combination of negotiator cultures. This guidance can avoid prolonged negotiations spent on particular issue outcomes that are unlikely to be reached.

Our second set of analyses investigates these cultural differences in issue outcomes to illustrate the benefits of recognizing dyad cultures for assistive agents to aid in reaching culturally-compatible issue outcomes. We employ the same bootstrap procedure to establish 95% confidence intervals for the measures we analyze. We first investigate how the total value of the negotiated outcomes differs by culture.

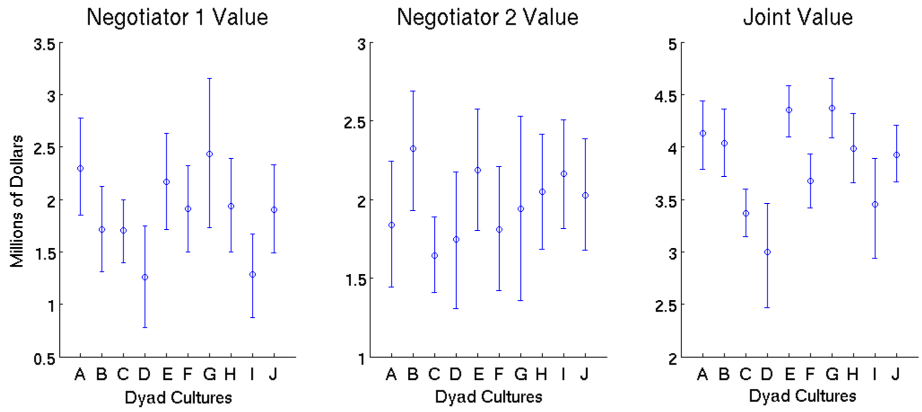

Figure 4: Individual and joint values of negotiated outcomes based on culture: American-American (A), Japanese-Japanese (B), Russian-Russian (C), Chinese-Chinese (D), Swedish-Swedish (E), Thai-Thai (F), Israeli-Israeli (G), German-German (H), American-Japanese (I), American-Chinese (J).

The results of this analysis, shown in Figure 4, illustrate that there are significant cultural differences, especially in the joint values of negotiated outcomes. This information is useful for an agent capable of recognizing negotiator cultures to either help avoid sub-optimal negotiation outcomes or exploit the culture-based disparities that tend to be realized.

We next investigate the mean and covariance of each of the four issue outcomes for each of the 10 dyad culture combinations.

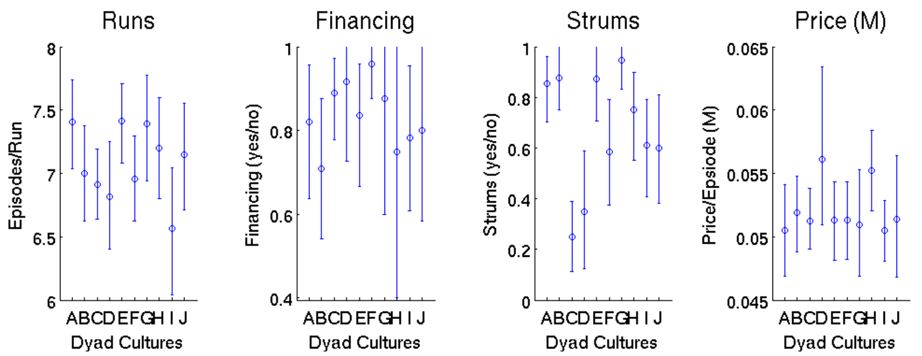

Figure 5: Four main issue outcomes for each of 10 culture combinations: American-American (A), Japanese-Japanese (B), Russian-Russian (C), Chinese-Chinese (D), Swedish-Swedish (E), Thai-Thai (F), Israeli-Israeli (G), German-German (H), American-Japanese (I), American-Chinese (J).

Statistical analysis of the four issue outcomes is presented in Figure 5 for each of the 10 dyad culture combinations. The greatest negotiation outcome differences between dyad types occur in the issue of whether to include the additional cartoon, *Strums*, as shown in the third plot of Figure 5. This suggests that e.g., when negotiating with Chinese or Swedish participants, including the *Strums* cartoon in the deal is less likely to occur and that effort may be saved by not heavily pursuing that issue.

We additionally investigate to what degree these outcomes are related and how those relations differ by dyad culture combination.

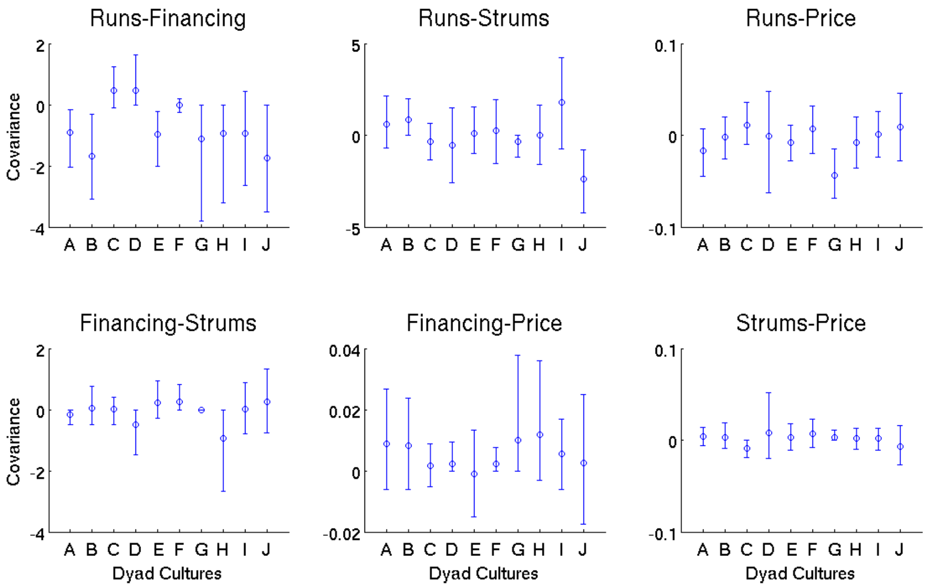

Figure 6: Covariance of the runs-financing, runs-strums, runs-price, financing-strums, financing-price, and strums-price issue outcomes by culture combinations: American-American (A), Japanese-Japanese (B), Russian-Russian (C), Chinese-Chinese (D), Swedish-Swedish (E), Thai-Thai (F), Israeli-Israeli (G), German-German (H), American-Japanese (I), American-Chinese (J).

Statistical analysis of the covariance between outcomes is shown in Figure 6. These relationships between issue outcomes provide additional statistically significant differences between dyad culture combinations where none were evident from the individual outcomes. For example, the covariance is positive for the Russian-Russian and Chinese-Chinese negotiators for the number of runs and financing issues, but negative for American-American and Japanese-Japanese negotiators. Understanding the correlations between issues can inform an agent's long-term negotiation strategy or help identify more likely negotiation outcomes given what issues have already been resolved. The results of both the means and covariance analyses are useful for an autonomous agent to guide interactions towards outcomes that may be more culturally acceptable to the negotiation partner.

3 Discussion

Recognizing the culture of negotiators and understanding the cultural differences in negotiation outcomes are both important for enabling a computational agent to guide or lead successful negotiations. We employ passive learning of cultural indicators and relationships to negotiation outcomes from observed negotiation transcripts in this work. However, the degree to which these indicators and relationships may change when an agent intervenes upon the negotiation process is not clear and the subject of our future work. Such automated negotiation agents have been studied by Kraus et al. (2007), but with available interactions that are much more restricted than those of the human negotiation schema identified in Table 1. We believe that our results for identifying cultures and recognizing cultural outcome differences in more natural interaction settings will lead to negotiation agents able to assist negotiators in broader negotiation domains.

References

Adair, W.L. and Brett, J.M. (2005), The Negotiation Dance: Time, Culture, and Behavioral Sequences in Negotiation." *Organization Science*, 16:1, 33-51.

Adair, W.L., Okumura, T., and Brett, J.M. (2001), Negotiation Behavior When Cultures Collide: The United States and Japan, *Journal of Applied Psychology*, 86:3, 371-385.

Adair, W.L. and Weingart, L. and Brett, J.M. (2007), The Timing and Function of Offers in US and Japanese Negotiations. *Journal of Applied Psychology*, 92:4, 1056-1068.

Brett, J.M. and Okumura, T. (1998), Inter- and Intracultural Negotiation: US and Japanese Negotiators. *Academy of Management Journal*, 41:5, 495-510.

Dudík, M., Phillips, S.J., and Schapire, R.E. (2007), Maximum entropy density estimation with generalized regularization and an application to species distribution modeling. *Journal of Machine Learning Research*, 8, 1217-1260.

Hall, E.T. (1976), *Beyond Culture*. New York: Anchor Press.

Jaynes, E. T. (1957), Information Theory and Statistical Mechanics. II. *Physical Review*, *108*(2), 620-630. APS.

Kraus, S., Hoz-Weiss, P., Wilkenfeld, J., Anderson, D.R., and Pate, A. (2007), Resolving crises through automated bilateral negotiations. *Artificial Intelligence*, 172:1, 1-18.

Morley, I.E. and Stephensen, J.M. (1977), *The Social Psychology of Bargaining*. Allen & Unwin, London, U.K.

Pruitt, D. G. (1971), Indirect Communication and the Search for Agreement in Negotiation. *J. Appl. Soc. Psych.,* 1, 205-239.

Pruitt, D.G. (1981), *Negotiation Behavior*. Academic Press, New York.

Putnam, L.L., and Jones, T.S. (1982), The role of communication in bargaining. *Human Comm. Res.,* 8:3, 262-280.

Tenbrunsel, A.E., and Bazerman, M.H. (1995), Working Women. J. M. Brett, ed. *Teaching Manual*. Dispute Resolution Research Center, Northwestern University, Evanston, IL.

Walton, R.E. and McKersie, R.B. (1965), *A Behavioral Theory of Labor Negotiations: An Analysis of a Social Interaction System*. McGraw-Hill, New York.