An Application of the INFINITE Framework in a Human-Agent Negotiation Competition

Bohan Xu
The University of Tulsa
bohan-xu@utulsa.edu

James Hale
The University of Tulsa
jah6484@utulsa.edu

Shadow Pritchard
The University of Tulsa
swp7196@utulsa.edu

Sandip Sen
The University of Tulsa
sandip@utulsa.edu

ABSTRACT

We discuss an application of the INFINITE negotiation architecture for developing agents that can negotiate with others while representing its user's preferences. We developed an agent, *Draft Agent*, that was entered into the 2019 Human-Agent league (HAL) of the Autonomous Negotiating Agent Competition (ANAC). We discuss *Draft Agent's* performance, highlighting where it worked well and aspects that can be further improved. A key feature of *Draft Agent* is the use of an alternate-issue-selection protocol to model the opponent's preference structure. The learnt preferences are then used to propose a fair, and where possible, win-win deal. Though this approach allows *Draft Agent* to obtain relatively high individual as well as joint utility, it might be considered somewhat rigid by human users and hence scores comparatively low on the likeability scale. We present a detailed analysis of the comparative performance of *Draft Agent* and the competing finalists of the HAL competition. We also suggest some options to further improve *Draft Agent's* performance and likeability.

CCS CONCEPTS

• **Computing methodologies → Intelligent agents**; • **Human-centered computing → Mixed / augmented reality**.

KEYWORDS

Negotiation; Architecture; Competition; Preference Elicitation

ACM Reference Format:
Bohan Xu, James Hale, Shadow Pritchard, and Sandip Sen. 2020. An Application of the INFINITE Framework in a Human-Agent Negotiation Competition. In *Proceedings of the 8th International Conference on Human-Agent Interaction (HAI '20), November 10–13, 2020, Virtual Event, NSW, Australia.* ACM, New York, NY, USA, 9 pages. https://doi.org/10.1145/3406499.3415081

1 INTRODUCTION

Negotiation is a preferred approach for resolving conflicts in human and agent societies. Automated negotiation is being increasingly used in different multiagent domains including robotics, network bandwidth allocation, space applications, etc. where agents need to reach agreements on the allocation of one or more shared resources [10]. Recent application domains include smart grids [14], cloud service composition [23], supply chain management [25], teaching [24], etc. In multi-issue negotiations, agents with divergent preferences can cooperate to reach agreements beneficial for both agents. But when the preferences are not common knowledge, self-interested agents often fail to explore win-win possibilities using existing protocols and end up with inefficient agreements. Hence, there is a need for studying and developing new negotiation architectures [13] and protocols that allow rational agents to leverage preference elicitation and strategic reasoning mechanism to reach mutually preferred agreements [1, 2, 20].

Researchers have been increasingly interested in developing autonomous agents that can negotiate as peers with human users [13, 15, 16, 22]. Designing agents to interact with or negotiate with humans remains a challenge. In particular, using formal notions of equilibria or optimality may not produce preferred negotiation outcomes; offers may not always be easy for humans to follow, as corresponding computation can impose unreasonable cognitive demands [15]. Recent efforts have focused on using more natural interfaces and interaction modalities, such as conversations and chats, to engage human users [21, 22, 24]. We leverage those modalities to estimate the user's preferences and then compute and offer the most fair, win-win deal. The cognitive burden on the human is significantly reduced and, we believe, this leads to better negotiation outcomes.

Experimentation and development of negotiation agents, including those that negotiate with humans, are facilitated by the Automated Negotiating Agent Competition (ANAC, http://ii.tudelft.nl/nego/node/7), organized yearly since 2010. The ANAC platform enables testing of new negotiation mechanisms using benchmark scenarios and against state-of-the-art competing agents developed by other researchers and using established protocols. The competition encompasses a number of distinct challenges, or leagues, that focus on different negotiation scenarios and facets, such as concurrent negotiation, negotiating with partial preferences, negotiating in teams, etc. This paper discusses our efforts to design, develop, field, and analyze the competition results of an agent for the 2019 edition of the *Human-Agent League* (HAL) (http://web.tuat.ac.jp/~katfuji/ANAC2019/files/cfp_IAGO.pdf) of the ANAC competition (http://web.tuat.ac.jp/~katfuji/ANAC2019/), where the goal is to develop agents to effectively negotiate with human participants.

HAI '20, November 10–13, 2020, Virtual Event, NSW, Australia

ACM ISBN 978-1-4503-8054-6/20/11…$15.00
https://doi.org/10.1145/3406499.3415081

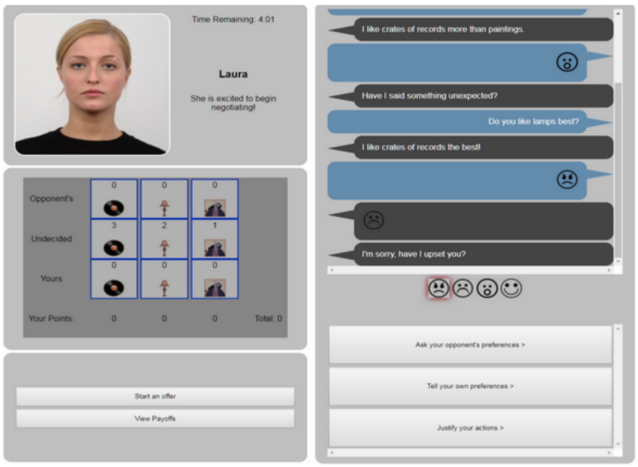

Figure 1: The IAGO Platform [17].

2 THE ANAC HUMAN-AGENT LEAGUE (HAL)

We briefly describe the platform used for the HAL competition, the competition structure, the winning criteria, as well as the participants and previous winners.

2.1 IAGO Platform

The ANAC 2019 Human-Agent League competition used the IAGO platform [17] : a web-based system that includes both a user-friendly interface as well as an API for developing agents. The graphical interface for the platform allows for the transmission of bids, text, and emoticons between the human subject and agent (see Figure 1). While the agents may send customized utterances, the participants must select text from a predefined pool. The agent presents itself as one of a few available avatars with an associated name.

2.2 Competition Structure

This competition focuses on negotiated resource allocation. In each game round, the platform presents the participant and agent with a set of resources, or issues, with each player having individual valuations. The participant and agent negotiate an allocation of those resources [8]. While the possible total utility over the resources in a round may not be equal for the human and agent, the cumulative points possible over three rounds will be the same [8]. The agent know its utility function, but not that of the human.

The competition structure allows for either partial or full offers [8]. A partial offer is an allocation of a subset of the total set of resources, while a full offer is a complete allocation all resources. The number of issues negotiated in a round ranges from 3 to 5.

Before a game, each participant read through instructions, and answered questions to verify completion of the reading [8]. A participant interacted with only one agent, and each agent completed approximately 25 games [8].

A participant's interaction with an agent consists of three seven minute rounds [8]. The platform grants access of time to both actors, and specifically warns the participant when one minute of negotiation time remains [8]. A negotiation round ends if both parties agree to a full offer or if time expires [8]. If the negotiators fail to reach an agreement, the two receive points equivalent to their BATNA (Best Alternative to Negotiated Offer) [8].

The organizers first checked the agents submitted for obvious run-time errors. Subsequently, the finalists were selected. The final contest results were announced during the ANAC workshop held in conjunction with IJCAI-19.

2.3 Winning Criteria

The winner of the competition was the agent with the highest average cumulative reward over three negotiation rounds. While the competition organizers recorded likeability, it did not directly affect winning criteria [8] as was the case in the previous year [19].

2.4 Participant Pool & Prior Winners

Participants for this competition were sourced from Amazon Mechanical Turk (MTurk) [8] platform. The following conditions were applied for selecting participants: all participants asserted they are at least 18 years old, speak English as their first language, and are a permanent resident of the United States–confirmed via IP tracking [8]. The winners of previous years are briefly described here.

LyingAgent, 2017 [19]: agent conceded items that seemed valuable, but were actually nearly worthless to it, by misleading the human participant into believing they were seeking the same items.

Equalist, 2018 [18]: agent offered a very positive deal for the human in the first negotiation, in exchange for a very positive deal for it in the second negotiation. The agent maintained a tough stance, to out-maneuver its opponent in the third negotiation.

Of note, there was only one negotiation round in each game of 2017, while the competition introduced an extra two consecutive negotiations in 2018. Therefore, lying to human participants in early negotiations may destroy trust, and the participants are likely to refuse further cooperation. In the 2018 competition, the maximum payoffs for the human and agent were the same in each negotiation and remained unchanged over negotiations [7]. However, in 2019, the maximum payoffs changed in each negotiation, but the summations of maximum payoff over three negotiations for the human and agent were the same [8]. It makes *Equalist* unsuitable for the 2019 competition, since it cannot guarantee to make up for the loss of the first negotiation in the second.

3 THE INFINITE NEGOTIATION ARCHITECTURE

The scope of this paper does not cover a detailed discussion of our general **INF**ormed, **I**ntelligent **N**egotiation with **I**terated, **T**rusted **E**ngagement (**INFINITE**) architecture. We do, however, present the INFINITE architecture in Figure 2, discuss in brief its principal components and then show how our entry to the ANAC-19 HAL league was instantiated from this general framework. The INFINITE architecture was designed as a general framework that supported both learning the preferences of the user who it represented and effectively negotiating with other agents and/or humans. In the following, we discuss only those components of the INFINITE architecture that capture the distinctive features of the implementation of *Draft Agent* as entered into the 2019 ANAC HAL competition.

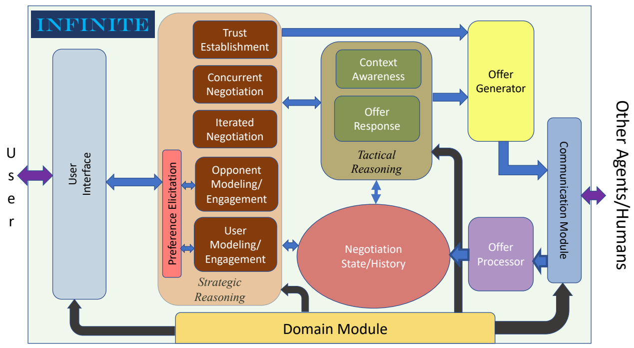

Figure 2: The INFINITE architecture.

The **Domain Module** contains the specifics of the negotiation competition structure and evaluation criteria for the competition. This module, understandably, informs various other key modules in the architecture, including the **Strategic** and **Tactical Reasoning** modules, the **Communication Module** and the **User Interface**.

We note that the agent utility function is directly provided to the agent in HAL and hence there is no interaction with the **User** that *Draft Agent* is representing. Hence the **User Interface** component is not implemented. On the other end, the **Communication Module** plays an important role as it uses bids, text messages, emoticons, etc. to interact with the IAGO web service, which also provides the user interface for human negotiators.

The **Offer Processor** module processes offers and other communication from the human users and updates the **Negotiation State/History**. The latter is also updated by the **Strategic** and **Tactical Reasoning** modules.

The INFINITE architecture separates the **Strategic** and **Tactical Reasoning** modules. The former involves long-term goal setting and planning and includes various sub-modules, of which only the **Trust Establishment**, **Iterated Negotiation**, and **Opponent Modeling/Engagement** modules are relevant for the discussion concerning *Draft Agent* deployment. In *Draft Agent*, **Trust Establishment** includes the use of text messages to establish the rationale for the protocols used for estimating user preferences and messages to reinforce the fact that the *Draft Agent*'s goal is to arrive at "win-win" negotiation outcomes. The **Iterated Negotiation** module sets the context for the three rounds of negotiation, with specific messaging. However, there is significant improvement scope for this module in the current implementation. Currently, we do not carry forward any estimate of human negotiation attitude or preferences between the negotiation rounds. To the extent such extrapolations are valid, both negotiation effort and outcome can be improved. The **Opponent Modeling/Engagement** module is a key component of *Draft Agent* as it is instrumental in strategically engaging with the opponent to first tease out relative preference of the opponent over the issues being negotiated and then develops offers that leverage that knowledge. This module utilizes the **Preference Elicitation** module to engage the human user in an augmented alternating issue selection protocol, as described in Section 4.2.

Distinct from the **Strategic Reasoning** module, the purpose of the **Tactical Reasoning** module is to take reactive, reflexive decisions, given the current context of negotiation. For example, this module can leverage time-sensitive information to quickly move on opportunities with very limited time windows that more strategic deliberation may not be able to respond to in a timely manner. This module monitors the outputs from the **Strategic Reasoning** module, the **Offer Processor** module and the current **Negotiation State/History**.

Both the **Strategic** and **Tactical Reasoning** modules can invoke the **Offer Generator** module to construct an offer or a counter-offer that will be communicated to the opponent(s) using the **Communication Module**. In the INFINITE architecture, the **Tactical Reasoning** module has "interrupt authority" over the **Strategic Reasoning** module, and hence can supersede an offer proposed by the latter with what it considers a more timely and opportune offer given the negotiation state/history and the most recent offer(s) received from opponent(s).

4 SOLUTION APPROACH

We now present the key design considerations and decisions embodied in *Draft Agent*, our entry to the 2019 ANAC HAL competition.

4.1 Negotiation Problem

We now formalize the repeated negotiation problem in HAL. In the following, *issues* and *resources* are used interchangeably. Let I_r be the set of issues being negotiated in round r. $\forall\, i_r \in I_r$, n_{i_r} is the number of items available for issue i. An offer O is a set of triplets $\{O(i_r)\}$, where $O(i_r) = \langle i_r, n^h_{O,i_r}, n^{AI}_{O,i_r} \rangle$ such that $\forall\, i_r \in I_r$, $n^h_{O,i_r} + n^{AI}_{O,i_r} \leq n_{i_r}$, where n^h_{O,i_r} and n^{AI}_{O,i_r} corresponds to the number of items of issue i_r being allocated in offer O to the human and AI player respectively (here, and below, we will use "AI" or "AI player" to refer to *Draft Agent* for brevity). An offer O is a *complete offer* iff $\forall\, i_r \in I_r$, $n^h_{O,i_r} + n^{AI}_{O,i_r} = n_{i_r}$; otherwise the offer is *partial*. $U(O) = \sum_{i_r \in I_r} w_{i_r} n^{AI}_{O,i_r}$ is the utility of offer O to the AI player, where w_{i_r} is the per unit utility of issue i_r to the AI player. We can similarly define the utility of any offer to the human player.

4.2 Preference Elicitation

A key challenge of negotiation is to understand and effectively utilize the utility preferences of one's opponent [1, 2, 20]. As the human opponent's utility function is unknown at the outset of a negotiation round, it needs to be learned or approximated to allow for strategic negotiation. A core functionality of *Draft Agent*, eliciting the utility preferences of the human opponent over the issues, is provided by adapting an existing negotiation protocol [1].

[1]This protocol is a variation of the *Strict alteration* protocol, in which agents take alternate turns and in each turn an agent selects one resource from the set of resources not yet allocated. After an agent selects a resource, the resource is removed from the negotiation set [4]. The advantage of this protocol is its simplicity and the time required to reach an agreement. Note that the alternating protocol is used here to gauge the relative preference of the opponent over the issues, which will impact the negotiation offers to be made later, but is used primarily for preference elicitation and not for immediate issue allocation. So this is different from issue-by-issue negotiation where the agenda of issues determining the order they are considered is critical [3, 5, 6, 11]. We recognize that this protocol will put increased cognitive load on the human when negotiating a larger number of issues.

Preference Elicitation Protocol (Pr): Players alternate picking issues. Depending on the number of issues being negotiated, either one round (**Pr1**) or two rounds (**Pr1** and **Pr2**) of the protocol will be utilized. In the following we first describe the process for inferring the human user's issue preferences when the negotiation scenario involves 5 issues. Thereafter we describe the process used in simpler scenarios consisting of 4 and 3 issues.

<u>5 issues, $\boldsymbol{A, B, C, D, E}$:</u>

PR1: Human opponent is asked to choose first and then the players alternate in choosing from the remaining issues. Assume the following sequence of issue choices: (1) Human chooses A, (2) *Draft Agent* chooses B, (3) Human chooses C, (4) *Draft Agent* chooses D, and (5) Human chooses E. The outcome is $(\boldsymbol{A, b, C, d, E})$[2]. We then know the following partial human preference: $\boldsymbol{A > C > E}$, $\boldsymbol{A > \{B, C, D, E\}}$, $\boldsymbol{C > (D, E)}$[3].

PR2: We repeat **PR1**, but *Draft Agent* chooses first, and uses partial human preference information inferred from **PR1**. As the human chose **A** in **PR1**, *Draft Agent* will pick it first. Thereafter, the possible issue choices in **PR2** and the corresponding inferred human preference over the issues is as follows:

$\boldsymbol{a, B, c, D, e}$**:** Human preference order is $\boldsymbol{A > B > C > D > E}$.
$\boldsymbol{a, B, c, E, d}$**:** Human preference order is $\boldsymbol{A > B > C > E > D}$.
$\boldsymbol{a, C, b, D, e}$**:** human preference order is $\boldsymbol{A > C > D > E}$.
$\boldsymbol{a, C, b, E, d}$**:** Human preference order is $\boldsymbol{A > C > E > D}$.

For the last two scenarios, we only know $\boldsymbol{C > B}$; we need to ask human preference between $\boldsymbol{B}$, $\boldsymbol{D}$, and $\boldsymbol{E}$. Consider the third scenario, ask if human prefers $\boldsymbol{B}$ to $\boldsymbol{D}$. If that is true, the preference order is $\boldsymbol{A > C > B > D > E}$. If not, ask if the human prefers $\boldsymbol{B}$ to $\boldsymbol{E}$. If that is true, the preference order is $\boldsymbol{A > C > D > B > E}$, otherwise the preference order is $\boldsymbol{A > C > D > E > B}$. A similar process can be followed to obtain the total preference order for the last case.

<u>4 issues $\boldsymbol{A, B, C, D}$:</u>

First perform **PR1**. Assuming the result is $(\boldsymbol{A, b, C, d})$, the human preference will be $\boldsymbol{A > C > D}$. Ask if human prefers $\boldsymbol{B}$ to $\boldsymbol{C}$. If that is true, the preference order is $\boldsymbol{A > B > C > D}$. If not, ask if human prefers $\boldsymbol{B}$ to $\boldsymbol{D}$. If that is true, the preference order is $\boldsymbol{A > C > B > D}$, otherwise the preference order is $\boldsymbol{A > C > D > B}$. A similar process can be followed to obtain the total preference order for any other sequence of issue choices in **Pr1**.

<u>3 issues $\boldsymbol{A, B, C}$:</u>

First perform **PR1**. Assuming the result is $(\boldsymbol{A, b, C})$, the human preference will be $\boldsymbol{A > C}$. Ask if human prefers $\boldsymbol{B}$ to $\boldsymbol{C}$. If that is true, the preference order is $\boldsymbol{A > B > C}$, else the preference order is $\boldsymbol{A > C > B}$. A similar process can be followed to obtain the total preference order for any other sequence of issue choices in **Pr1**.

4.3 Initial Offer Generation

This key step in the negotiation process depends on whether the human opponent followed the preference elicitation process.

4.3.1 Human player follows Preference Elicitation Protocol. Let P_h and P_{AI} be ordered lists of the issues according to the total preference order of the human and *Draft Agent* player respectively. In Algorithm 1, we present a recursive process of our initial offer generation protocol. In descending preference order, *Draft Agent* proposes an allocation of the issues for each position for each player; these issues may be different. If the two issues are the same for a given position, which means these issues occupy the same preference position for both players, the algorithm divides the issue items as evenly as possible between the players. If the issues at a given position are different, the algorithm grants each player all of the items for their respective issue choice at that position.

This allocation procedure embodies a sense of fairness: the human and the AI player are equally treated. If an actor prefers an issue more than the other, they will receive all items of that issue; if their preferences overlap, the algorithm splits the items.

Algorithm 1 Initial Offer Generation after Preference Elicitation

```
procedure InitialOffer(P_h, P_AI)
    if P_h is not empty then
        if P_h(first) = P_AI(first) then
            allocate half of P_h(first) items, subject to truncation, to each player
        else
            allocate all P_h(first) items to human and all P_AI(first) items to AI player
        P_alloc = {P_h(first)} ∪ {P_AI(first)}
        Remove P_alloc from both P_h and P_AI
        InitialOffer(P_h, P_AI)
```

4.3.2 Human player does not follow Preference Elicitation Protocol. A human player may not follow the preference elicitation protocol in Section 4.2, After a number of deviations (we used 6) from the protocol, an alternate initial offer mechanism is triggered. It allocates the most valuable half of the issues, as per the agent's utility function, to the *Draft Agent* player and the rest of the issues are allocated to the human player. If the number of issues is odd, the last remaining issue is split equally. This blind protocol is also used in the third round after the preference elicitation.

4.4 Counter Offer Generation

Let O_{AI} be the last offer the *Draft Agent* player made. Algorithm 2 describes how the *Draft Agent* player generates a counteroffer to an offer O_h made by the human in response to O_{AI}. Initially, the procedure checks whether the utility of the human's offer is equal to or greater than our previous offer; if so, the agent accepts. If not, the procedure continues to alter the human's offer until the loss–the decrease in *Draft Agent's* utility from our offer to the human's–is positive. To do this, the agent iterates through the issues in ascending order congruent with the human's preference ordering–if it does not know this, it iterates through its own preference ordering in descending order–taking either as many items of an issue to make its loss non-negative, or all items of an issue.

5 TEXT COMMUNICATIONS USED

In the following we present samples of text messages sent to users to initially establish trust by providing a rationale for engaging in the alternating issue selection protocol, some messages nudging the

[2] We use upper and lower case letters to designate issues chosen by the human player and *Draft Agent* respectively.
[3] We use the notation $x > Y$, where Y is a set of issues, to denote $x > y, \forall y \in Y$.

Algorithm 2 Counter-offer Generation

procedure COUNTEROFFER(O_{AI}, O_h, P_h, P_{AI})
  $loss = U(O_{AI}) - U(O_h)$
  if $loss < 0$ **then**
    Accept O_h
  else
    $O = O_h$; $P = P_h$; issuePrefernceOrder = Length(P); increment = -1
    if P_h is empty **then**
      $P = P_{AI}$; increment = 1; issuePrefernceOrder = 1
    while $loss > 0$ **do**
      $i = P$(issuePreferenceOrder); $n = n^h_{O_h,i}$
      if $n > 0$ **then**
        $itemsSwapped = \min(n, \lceil \frac{loss}{w(i)} \rceil)$
        $loss = loss - itemsSwapped * w(i)$
        $O = O \setminus O(i)$
        $O = O \cup \{(i, n - itemsSwapped, n_i - (n - itemsSwapped))\}$
      issuePreferenceOrder = issuePreferenceOrder + increment
    Counter offer O

user to follow the protocol if they deviated, a message accompanying offers, and messages to welcome users to succeeding rounds of negotiation. **Beginning of first game:** "Hi! Let's try for WIN-WIN outcomes by taking turns choosing resources. You go first: take all units of any one resource, such as all x units of $issue[0]$, and make an offer. I will then pick from the remaining resource and make a counter-offer and so on until we have chosen all resources."
If multiple issues picked during turn: "Please only select one unclaimed resource, not multiple."
If no issue is picked during turn: "Please select one and only one resource."
If human selects issue that has already been selected during draft: "The offer you sent had a resource that has already been selected. Please send a different offer."
After their first pick: "Now it is my turn to select a resource."
Beginning of round two (only for five issues): "Let's do that again but let me pick first this time."
Once we have enough data to make a full offer: "How about this for a final deal? I believe it is fair for both of us."
Beginning of second game: "Hi again! I hope you liked the outcome from the last round. Let's do what we did then. You pick first this time."
Start of last game: "Hi again! I hope you liked the outcome from the last round. This method seems to work well, so lets do it again! You get to pick first!"
After a period of inactivity: "It is your move now!"

Text communications were not a focus of our efforts in deploying *Draft Agent*. Though these messages are useful in guiding the negotiation process, there is scope for considerable improvement.

6 RESULTS

The following results are from the final round of the 2019 ANAC HAL competition[4], involving *Draft Agent* and five other finalists. For each agent, 25 human participants were selected as negotiation opponents and were asked a set of questions to ensure they understood and were engaged in the game. Participants who failed on the questions were removed from the results [8]. The numbers of validated games for each agent are listed in Table 1.

	Draft Agent	Dona	Mark	Neo	Pinocchio	Swindler
Valid	20	20	21	7	22	20
Failed	4	2	0	4	4	3
Outliers	4	2	0	0	4	3

Table 1: Numbers of validated/failed games and outliers for each agent.

Figure 3 shows a boxplot of the cumulative payoff over three rounds of the validated games for all agents. Different parts of a boxplot are explained in Figure 4. *Neo* has 4 out of 7 validated games that end with no agreement, and only obtains a small payoff from these games, which results in the large IQR as plotted in Figure 3. Although *Draft Agent* received the highest average cumulative payoff averaged over all human opponents, which is 0.45 higher than the second place finisher (*Dona*), the difference is not statistically significant as per a pairwise t-test ($t = 0.05$, $p = 0.96$, $p_{\text{adjusted}} = 0.96$). However, there are several outliers indicated as black dots in Figure 3, which represent games that ended with failed negotiations. The number of failed negotiations and outliers for each agent are shown in Table 1 as well. We should note that not all failed negotiations are outliers, the outliers are statistically indicated for each agent, while the failed negotiations are identified based on negotiation outcome directly. In this case, *Neo* has 4 out of 7 failed negotiations, but none of them can be referred as an outlier statistically. We show in Table 2, pairwise t-test results between *Draft Agent* and every other agent both with all data points and when the outliers are removed. Statistically significant p-values and False Discovery Rate (FDR) adjusted p-values are highlighted for a 0.05 threshold in Table 2. *Draft Agent* shows significantly higher average cumulative payoff compared to other agents except *Dona*. The adjusted p-value of 0.0546 still indicates marginally better performance of *Draft Agent* compared to *Dona*, when outliers are removed.

		Dona	Mark	Neo	Pinocchio	Swindler
All results	t	0.05	0.95	2.68	1.11	0.66
	df	36.48	21.33	9.98	38.28	37.28
	p	9.62E-01	3.55E-01	2.31E-02	2.75E-01	5.11E-01
	p (adjusted)	9.62E-01	4.44E-01	4.62E-02	3.93E-01	5.68E-01
Remove outliers	t	2.25	8.65	4.21	4.48	3.41
	df	26.97	33.81	6.25	31.43	28.63
	p	3.27E-02	4.43E-10	5.16E-03	9.27E-05	1.95E-03
	p (adjusted)	5.46E-02	4.43E-09	1.29E-02	4.64E-04	6.49E-03

Table 2: t-test results of cumulative payoff between Draft Agent and competitors.

Figure 5 shows the average payoff, over all opponents, for each round and for all agents. The increases of average payoff between

[4]We thank the ANAC HAL organizers for organizing the competition. We are particularly indebted to Jonathan Mell for providing us with the competition result data.

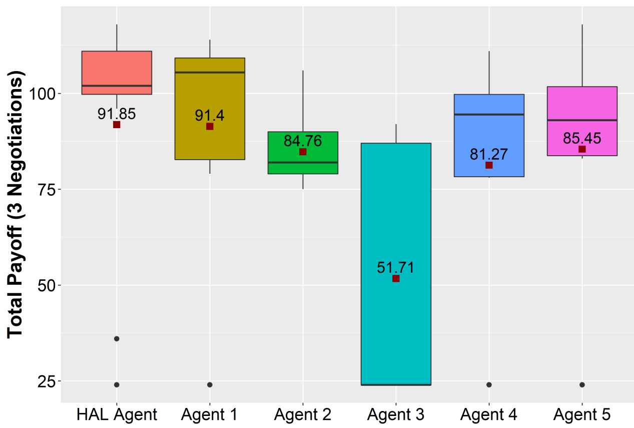

Figure 3: Average cumulative agent payoffs against all human opponents.

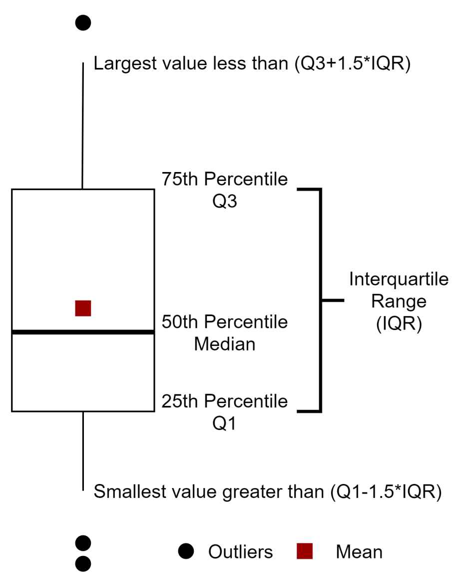

Figure 4: Elements of a boxplot [9, 12].

two consecutive rounds, i.e., $\text{Payoff}_{i+1} - \text{Payoff}_i$, are listed in Table 3, where the background of a cell is colored according to its value (color ranges from green to yellow to orange to red as values increase). All agents attain higher or equal average payoff in following rounds. Considering all validated games, *Draft Agent* outperforms all other agents in the second round, and gains the largest increase (10.35) in average payoff between the first and second rounds. Although *Draft Agent* still attains the second highest average payoff in the third round over all validated games, the increase (3.65) between second and third rounds is the lowest among all agents except *Neo*. Similar to cumulative payoff, the *t*-test results for payoff in each round are presented in Tables 4, 5 and 6, and significant p-values are highlighted in red as well. Based on these results, we can conclude that *Draft Agent* statistically outperforms all other competitors in the second round, and performs better than most competitors in the third round, provided human participants and agents reach an agreement (outliers are removed from these results).

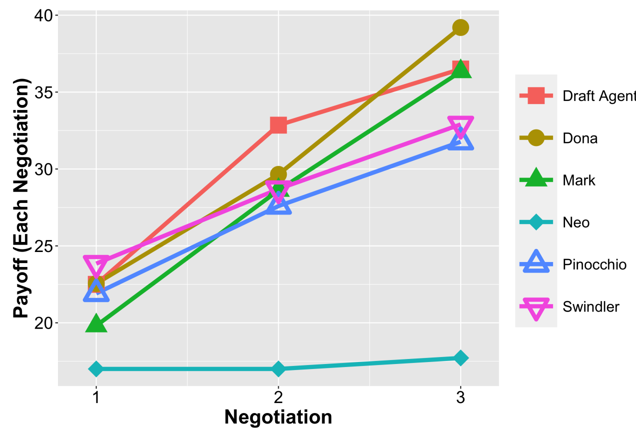

(a) Average payoff in each round (all results).

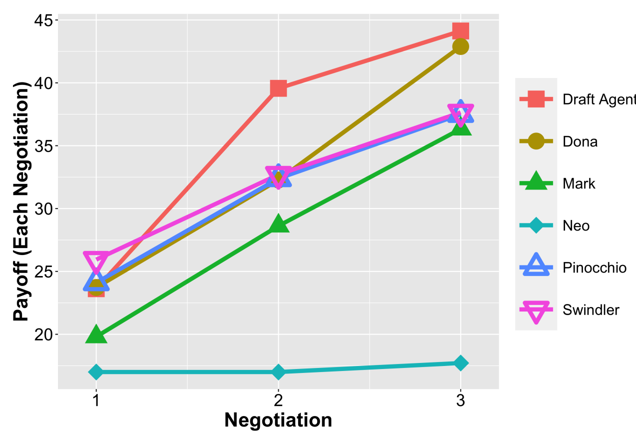

(b) Average payoff in each round (outliers removed).

Figure 5: Average agent payoffs over rounds.

	Round 1-2	Round 2-3
Draft Agent	10.35	3.65
Dona	7.10	9.55
Mark	8.81	7.71
Neo	0.00	0.71
Pinocchio	5.68	4.18
Swindler	4.85	4.20

Table 3: Increase in average payoff between rounds (all results): difference between a round and its preceding round.

Apart from payoffs and though it was not part of the competition winner determination, agent likeability is also measured on a 7-point Likert scale with human players queried after the first and second rounds of each game. Figure 6 depicts the average likeability of each agent from the perspective of human participants, and

		Dona	Mark	Neo	Pinocchio	Swindler
All results	t	-0.04	2.44	2.04	0.42	-0.89
	df	37.99	37.02	7.40	38.33	33.55
	p	9.68E-01	1.97E-02	7.84E-02	6.77E-01	3.81E-01
	p (adjust)	9.68E-01	6.57E-02	1.57E-01	8.46E-01	6.35E-01
Remove outliers	t	-0.18	4.83	2.57	-0.72	-3.21
	df	31.80	29.94	6.26	29.06	26.24
	p	8.55E-01	3.80E-05	4.10E-02	4.76E-01	3.45E-03
	p (adjust)	9.49E-01	3.80E-04	1.02E-01	6.80E-01	1.72E-02

Table 4: *t*-test results of payoff between Draft Agent and competitors in negotiation round 1.

		Dona	Mark	Neo	Pinocchio	Swindler
All results	t	0.81	1.21	2.54	1.26	0.98
	df	28.98	19.88	11.83	33.76	34.07
	p	4.23E-01	2.42E-01	2.62E-02	2.17E-01	3.35E-01
	p (adjust)	4.23E-01	3.03E-01	4.37E-02	3.03E-01	3.72E-01
Remove outliers	t	3.54	5.26	4.05	3.28	2.88
	df	16.35	17.07	7.87	20.37	25.49
	p	2.66E-03	6.36E-05	3.78E-03	3.72E-03	7.94E-03
	p (adjust)	9.46E-03	6.36E-04	9.46E-03	9.46E-03	1.59E-02

Table 5: *t*-test results of payoff between Draft Agent and competitors in negotiation round 2.

		Dona	Mark	Neo	Pinocchio	Swindler
All results	t	-0.53	0.05	2.87	1.04	0.79
	df	37.90	21.51	11.20	38.04	36.69
	p	5.98E-01	9.64E-01	1.50E-02	3.07E-01	4.33E-01
	p (adjust)	7.48E-01	9.64E-01	3.00E-02	5.11E-01	6.18E-01
Remove outliers	t	0.41	8.61	4.77	4.34	4.28
	df	17.03	20.36	6.00	17.11	16.10
	p	6.83E-01	3.17E-08	3.08E-03	4.38E-04	5.70E-04
	p (adjust)	7.59E-01	3.17E-07	7.70E-03	1.90E-03	1.90E-03

Table 6: *t*-test results of payoff between Draft Agent and competitors in negotiation round 3.

Draft Agent is the fourth liked among these agents. As expected, likeability is positively correlated to user payoff. This is confirmed by the slope of 0.053 ($p = 0.027$, see Figure 7) of the black dashed line representing the linear regression model; the gray band represents the 95% confidence interval for predictions. We also observed the following: (a) likeability is negatively correlated to the "point lead" (human score subtracted from agent score) [18] with a slope of -0.053 ($p = 0.112$) and plotted in Figure 8, and (b) likeability is positively correlated to the "joint point" (total utility generated by the final agreement) [18] with a slope of 0.027 ($p = 0.113$) as shown in Figure 9. The R^2 values in Figures 7, 8 and 9 indicates that likeability has a stronger correlation with user payoff than "point lead" or "joint point".

7 DISCUSSION

Internal testing was performed within our lab and before submitting *Draft Agent* to the competition. Feedback from several students was collected after they negotiated with *Draft Agent*. One common feeling from these students is that *Draft Agent* appears overly persistent in following the designed protocols to obtain the human participant's preference. It could be a sticking point if the human opponents do not want to follow the preference elicitation protocol for one of many reasons. Our conjecture is that though the Preference Elicitation protocol does allow *Draft Agent* to obtain better payoffs, it also leads to a lower score on the likeability scale as evident from the regression line in Figures 7 and 9. In the worst

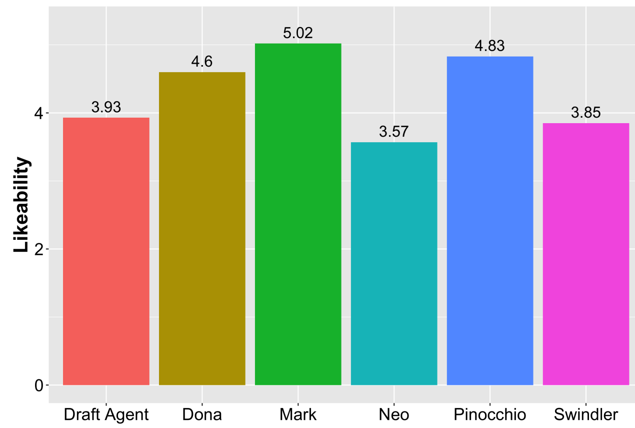

Figure 6: User Likeability of different agents.

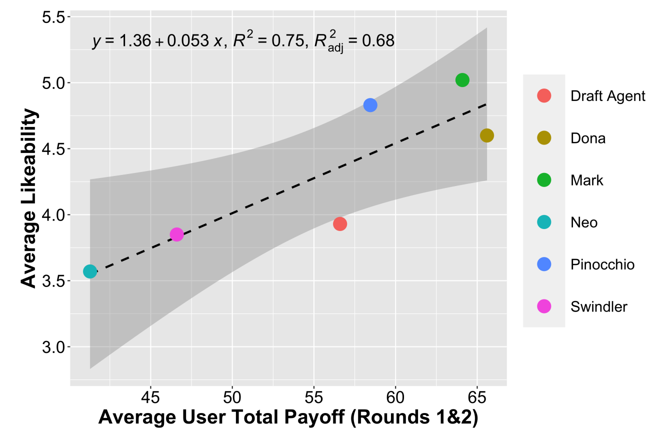

Figure 7: Likeability vs user payoff.

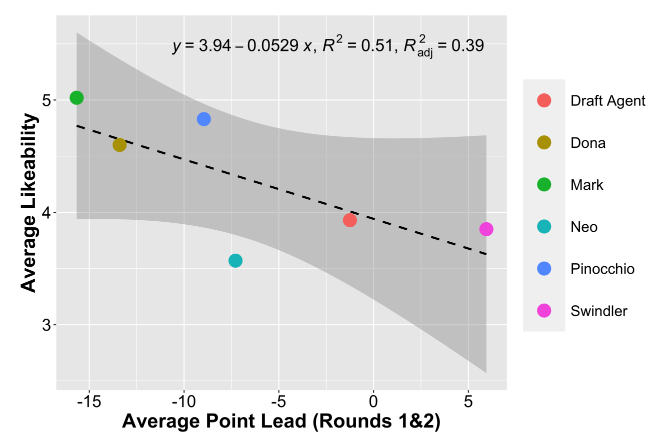

Figure 8: Likeability vs agent point lead.

case, it may cause participants to refuse to cooperate and leave the negotiation process without an agreement, e.g., 20% of games end

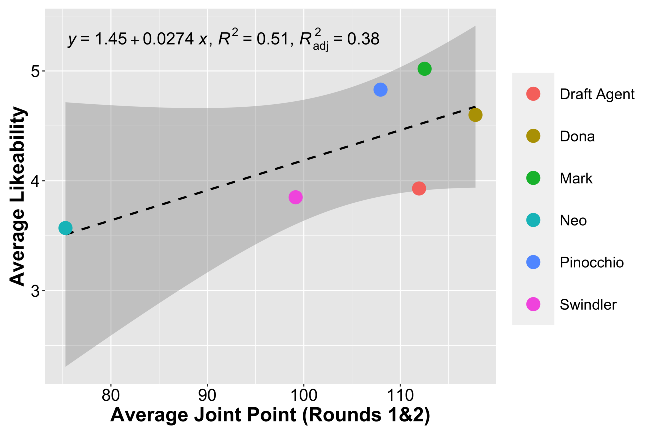

Figure 9: Likeability vs joint utility.

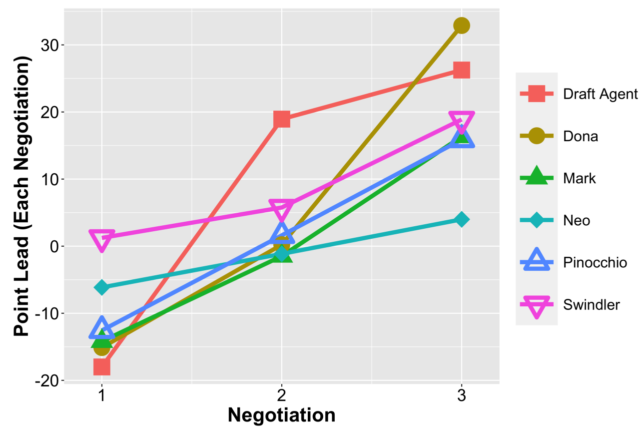

Figure 10: Average point lead in rounds (outliers removed).

with no agreement for *Draft Agent* which is the highest rate among all agents.

The high "point lead" of *Draft Agent* in the second round (see Figure 10) surprised us as we designed for win-win outcomes. Following is a possible explanation: As *Draft Agent* obtained much less payoff than the participant in the first round (this was the lowest average "point lead" in the first round) and the human participants see the payoff for both sides after each round, participants may have tended to be more generous to *Draft Agent* in the second round.

On a related note, *Draft Agent* obtained the highest "joint point" with the human player in the first and third rounds (see Figure 11), and the highest cumulative "joint point" (see Table 7) which is significantly higher than *Mark*, *Neo* and *Swindler*.

The analysis of likeability scores provide another perspective. Human participants are likely more focused on their own payoff, rather than the difference with their opponent's payoff or achieving win-win outcomes!

However, there are some limitations in the study as well. As we mainly focused on designing an agent for competition where the maximum number of issues is 5 in each negotiation, we do not propose a protocol to obtain human's preference with number of issues > 5. Although the current protocol is extendable for more issues, it may not be the most efficient and suitable approach to acquire human's preference. On the other hand, a further analysis cannot be performed, due to the small number of participants and limited negotiation results we have.

	Draft Agent	Dona	Mark	Neo	Pinocchio	Swindler
average	187.44	179.78	168.86	106.71	183.06	166.71
p	-	1.24E-01	2.54E-04	2.68E-02	3.66E-01	1.51E-05
p (adjust)	-	1.55E-01	6.35E-04	4.47E-02	3.66E-01	7.55E-05

Table 7: *t*-test results of cumulative joint point between Draft Agent and competitors (outliers removed).

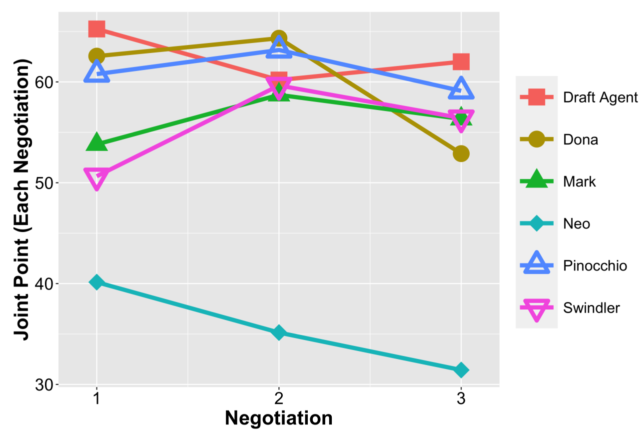

Figure 11: Average joint point in rounds (outliers removed).

8 CONCLUSIONS AND FUTURE WORK

We outlined a framework–INFINITE–for developing agents that can negotiate effectively with opponents, including humans, while representing their user's preferences. We describe an instantiation of INFINITE, *Draft Agent*, that competed effectively in the 2019 ANAC HAL competition. The key components of *Draft Agent* that contribute to its success are the adaptation of an alternating offer protocol for preference elicitation, the use of an impartial initial offer generation process, and a counter-offer process that seeks improvement of local utility while minimally reducing opponent payoff. As a result, *Draft Agent* outperformed other finalists in the competition in terms of cumulative agent and joint payoffs.

We plan to use a different strategy for the last round where the *Draft Agent* under-performs. Currently, we do not carry forward estimates of human negotiation attitude or preferences between negotiation rounds. When such extrapolations are valid, both negotiation effort and outcome can be improved by using preferences and attitudes estimated in previous rounds using the **Iterated Negotiation** module of INFINITE. On a different dimension, it would be interesting to ascertain what effect continuous facial expressions and emotions have in human-agent interactions–as opposed to the discrete emotions available in this competition.

Acknowledgments: Bohan Xu was supported by a LIBR grant, James Hale and Shadow Pritchard received funding from the Tulsa Undergraduate Rssearch Challenge program, and Sandip Sen was partially funded by AFOSR VFRP.

REFERENCES

[1] Tim Baarslag and Enrico H Gerding. 2015. Optimal incremental preference elicitation during negotiation. In *Twenty-Fourth International Joint Conference on Artificial Intelligence*. 3–9.

[2] Tim Baarslag, Michael Kaisers, Enrico Gerding, Catholijn M Jonker, and Jonathan Gratch. 2017. When will negotiation agents be able to represent us? The challenges and opportunities for autonomous negotiators. International Joint Conferences on Artificial Intelligence, 4684–4690.

[3] Mehmet Bac and Horst Raff. 1996. Issue-by-issue negotiations: the role of information and time preference. *Games and Economic Behavior* 13, 1 (1996), 125–134.

[4] Steven J Brams and Alan D Taylor. 2000. *The win-win solution: guaranteeing fair shares to everybody.* WW Norton & Company.

[5] Lutz-Alexander Busch and Ignatius Horstmann. 1997. A comment on issue-by-issue negotiations. *Games and Economic Behavior* (1997).

[6] M Keith Chen. 2002. Agendas in multi-issue bargaining: When to sweat the small stuff. *Harvard Department of Economics, Cambridge* (2002).

[7] Ninth International Automated Negotiating Agents Competition. 2018. Human-agent League CFP. http://web.tuat.ac.jp/~katfuji/ANAC2018/files/Human.pdf.

[8] Tenth International Automated Negotiating Agents Competition. 2019. Human-agent League CFP. http://web.tuat.ac.jp/~katfuji/ANAC2019/files/cfp_IAGO.pdf.

[9] Laura DeCicco. 2018. Exploring ggplot2 boxplots - Defining limits and adjusting style. https://waterdata.usgs.gov/blog/boxplots/.

[10] Ulrich Endriss, Nicolas Maudet, Fariba Sadri, and Francesca Toni. 2006. Negotiating socially optimal allocations of resources. *Journal of artificial intelligence research* 25 (2006), 315–348.

[11] Shaheen Fatima, Michael Wooldridge, and Nicholas R. Jennings. 2003. Optimal Agendas for Multi-Issue Negotiation. In *Proceedings of the Second International Joint Conference on Autonomous Agents and Multiagent Systems* (Melbourne, Australia) *(AAMAS '03)*. Association for Computing Machinery, New York, NY, USA, 129–136. https://doi.org/10.1145/860575.860597

[12] Michael Galarnyk. 2018. Understanding Boxplots. https://towardsdatascience.com/understanding-boxplots-5e2df7bcbd51.

[13] Catholijn M Jonker, Valentin Robu, and Jan Treur. 2007. An agent architecture for multi-attribute negotiation using incomplete preference information. *Autonomous Agents and Multi-Agent Systems* 15, 2 (2007), 221–252.

[14] Wolfgang Ketter, John Collins, and Prashant Reddy. 2013. Power TAC: A competitive economic simulation of the smart grid. *Energy Economics* 39 (2013), 262–270.

[15] Raz Lin and Sarit Kraus. 2010. Can automated agents proficiently negotiate with humans? *Commun. ACM* 53, 1 (2010), 78–88.

[16] Raz Lin, Sarit Kraus, Jonathan Wilkenfeld, and James Barry. 2008. Negotiating with bounded rational agents in environments with incomplete information using an automated agent. *Artificial Intelligence* 172, 6-7 (2008), 823–851.

[17] Johnathan Mell and Jonathan Gratch. 2016. IAGO: interactive arbitration guide online. In *AAMAS*. 1510–1512.

[18] Johnathan Mell, Jonathan Gratch, Reyhan Aydoğan, Tim Baarslag, and Catholijn M Jonker. 2019. The Likeability-Success Tradeoff: Results of the 2 nd Annual Human-Agent Automated Negotiating Agents Competition. In *2019 8th International Conference on Affective Computing and Intelligent Interaction (ACII)*. IEEE, 1–7.

[19] Johnathan Mell, Jonathan Gratch, Tim Baarslag, Reyhan Aydoğan, and Catholijn M Jonker. 2018. Results of the first annual human-agent league of the automated negotiating agents competition. In *Proceedings of the 18th International Conference on Intelligent Virtual Agents*. 23–28.

[20] Yasser Mohammad and Shinji Nakadai. 2018. Fastvoi: Efficient utility elicitation during negotiations. In *International Conference on Principles and Practice of Multi-Agent Systems*. Springer, 560–567.

[21] Avi Rosenfeld, Inon Zuckerman, Erel Segal-Halevi, Osnat Drein, and Sarit Kraus. 2014. NegoChat: a chat-based negotiation agent. In *Proceedings of the 2014 international conference on Autonomous agents and multi-agent systems*. 525–532.

[22] Avi Rosenfeld, Inon Zuckerman, Erel Segal-Halevi, Osnat Drein, and Sarit Kraus. 2016. NegoChat-A: a chat-based negotiation agent with bounded rationality. *Autonomous Agents and Multi-Agent Systems* 30, 1 (2016), 60–81.

[23] Kwang Mong Sim. 2011. Agent-based cloud computing. *IEEE Transactions on services computing* 5, 4 (2011), 564–577.

[24] David Traum, Stacy C Marsella, Jonathan Gratch, Jina Lee, and Arno Hartholt. 2008. Multi-party, multi-issue, multi-strategy negotiation for multi-modal virtual agents. In *International Workshop on Intelligent Virtual Agents*. Springer, 117–130.

[25] Minhong Wang, Huaiqing Wang, Doug Vogel, Kuldeep Kumar, and Dickson KW Chiu. 2009. Agent-based negotiation and decision making for dynamic supply chain formation. *Engineering Applications of Artificial Intelligence* 22, 7 (2009), 1046–1055.