

The Significance of Bidding, Accepting and Opponent Modeling in Automated Negotiation

Tim Baarslag¹ and Alexander Dirkzwager and Koen V. Hindriks and Catholijn M. Jonker

Abstract.

Given the growing interest in automated negotiation, the search for effective strategies has produced a variety of different negotiation agents. Despite their diversity, there is a common structure to their design. A negotiation agent comprises three key components: the *bidding strategy*, the *opponent model* and the *acceptance criteria*. We show that this three-component view of a negotiating architecture not only provides a useful basis for developing such agents but also provides a useful analytical tool. By combining these components in varying ways, we are able to demonstrate the contribution of each component to the overall negotiation result, and thus determine the key contributing components. Moreover, we study the interaction between components and present detailed interaction effects. Furthermore, we find that the bidding strategy in particular is of critical importance to the negotiator's success and far exceeds the importance of opponent preference modeling techniques. Our results contribute to the shaping of a research agenda for negotiating agent design by providing guidelines on how agent developers can spend their time most effectively.

1 Introduction

Negotiation is a frequently used process by which different parties communicate with one another in an attempt to reach a mutually acceptable agreement [18]. At the same time, negotiation can also be time consuming and costly. With the advances in modern computing, software agents are becoming more capable of automating the negotiation process.

The last two decades have seen a growing interest in the automation of negotiation, and the search for an effective negotiating agent has produced ever more effective negotiation strategies. A number of studies put emphasis on the *bidding strategy*; focusing on what kind of offers the negotiation strategy should choose, and in particular, what kind of concessions should be made [9, 19, 20]. Other research puts more emphasis on *learning* techniques to model various opponent attributes, such as the (partial) preference profile [2, 23] or the opponent's next move [2, 6]. Finally, some studies focus more on when to *accept*; i.e., under which circumstances an offer by the opponent should be agreed to [4, 5].

Automated negotiation research has been successful in finding increasingly better components in each of these aspects, either by focusing on one particular component, or by creating a new negotiation strategy altogether. However, the interactions and relative importance of the individual components have not been studied in detail before. Some components may have a stronger impact on the performance than others, and there could be strong interdependencies between the

components (e.g., a very competitive strategy hampering a learning technique by not exploring enough options). Thus, the main question of this paper is: *How do the components of a negotiating agent influence its overall performance, and which components are the most important?*

Answering this question requires us to bring together so-far unconnected research on various elements of negotiating agent design, and to research whether combining effective key components from different agents improves an agent's overall performance. If this is the case, then we have demonstrated that our component-based view on the architecture of negotiating agents forms the basis of a useful tool which allows an analytical approach towards the individual components of a negotiating agent. This process will give us a better understanding of how we can improve negotiation agents, as well as provide guidelines for the design of negotiating agents.

This paper investigates the importance and relations between three key components of negotiation strategies: the bidding strategy, the opponent model and the acceptance criteria. The question then becomes: *what is more important for a negotiation strategy to do well: to bid, to learn about the opponent, or to accept?* Or, more formally, how do each of the three components contribute to the effectiveness of a negotiation agent, and what are the interaction effects between them? We make these questions precise by formulating them in quantifiable terms of predictability of variance, and we determine the contribution of each component using the statistical measure of effect size. Once the individual contribution of each key component is established, we focus on the effects of combining components.

Our findings indicate that the bidding strategy is by far the most important component, while the significance of learning about the opponent's preferences is rather small in comparison. Given the current focus on opponent learning techniques in automated negotiation research, we argue that more effort needs to be made to formulate effective bidding strategies.

The paper is structured as follows. Section 2 discusses the distinctions between the different components within a negotiation agent in more detail. Section 3 presents our experimental setup, followed by Section 4, which looks at the contribution of each component. We present work relating to ours in Section 5 and finally, Section 6 presents our conclusions and recommendations for future work.

2 The Components of a Negotiating Agent

In this work we focus on bilateral automated negotiations, where agents take turns in exchanging offers using the alternating offers protocol. A negotiation scenario consists of the negotiation domain, which specifies all possible bids, together with a privately-known preference profile for each party. This means that the players do not

¹ Delft University of Technology, Netherlands, email: T.Baarslag@tudelft.nl

have access to the preferences of the opponent; however, the players can attempt to learn them during the negotiation encounter. The agents seek to reach an agreement while at the same time aiming to maximize their own utility. We assume a common discrete time line as in [21], with a deadline after a specified number of rounds. Both agents receive utility 0 if they do not succeed in reaching an agreement in time.

There is a wide variety of currently existing agent strategies for the above setting, such as the time dependent tactics [9], the behavior dependent tactics such as *Absolute* and *Relative Tit for Tat* [9], and strategies from the Automated Negotiating Agents Competition (ANAC) [2], which is a yearly international competition for negotiation agents. Despite this diversity, there is some common structure to the overall design of the agents. For example, every agent decides whether the opponent's offer is acceptable, and if not, what offer should be proposed instead. In addition, when the agent decides on the counter-offer, it considers its own utility, but it usually also takes the opponent's utility into account. We use the classification given in [3] to distinguish three different types of behavior, which can be considered to be the basic components of a negotiation strategy:

1. **Bidding strategy (BS).** At each turn, the *bidding strategy* determines the counter offer by first generating a set of bids B , depending on factors such as the negotiation history (i.e., the previous offers by the opponent), a target threshold, time, and so on. Note that during this stage, the agent only considers what concessions it deems appropriate given its own preferences. The *bidding strategy* may use the *opponent model* (if present) to select a bid from B by taking the opponent's utility into account.
Input: *opponent utility of bids, negotiation history.*
Output: *provisional upcoming bid ω .*
2. **Opponent model (OM).** An *opponent model* uses learning techniques to learn the opponent's attributes. We focus primarily on preference learning techniques (i.e., models of the opponent's preferences), but the agent may use this information to learn other attributes as well (e.g., predicting the opponent's strategy).
Input: *set of bids B , negotiation history.*
Output: *estimated opponent utility of the bids in B .*
3. **Acceptance Criteria (AC).** The *acceptance criteria* decide whether the opponent's offer should be accepted. If the opponent's bid is not accepted, the bid generated by the bidding strategy is offered instead.
Input: *provisional upcoming bid ω , negotiation history.*
Output: *accept, or send out the upcoming bid ω .*

To better understand how the different components work together, we might view the negotiation process as a search problem, where the negotiation strategy explores the outcome space for a contract that both parties are willing to agree upon (Figure 1). The bidding strategy controls the rate of concession by setting the *target utility range*, which determines the general location of the offer in the outcome space according to the agent's own utility. The opponent model can restrict this area even further, by refining the possible offers to bids that are near the Pareto frontier, and hence are the *best outcomes for the opponent*. Finally, the acceptance criteria define the area that consists of all *acceptable outcomes*, depending on the jump the agent is willing to make towards the opponent in order to reach an agreement.

3 Experiments

To analyze the relative importance and interactions of each strategy component, we require a wide array of tactics and techniques for

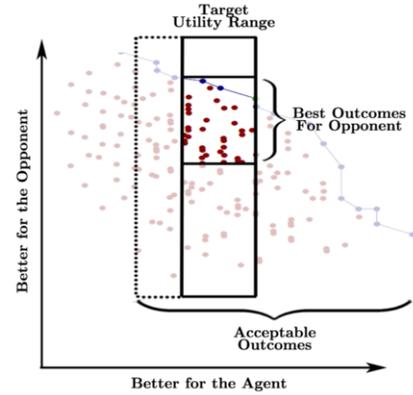


Figure 1. The bidding strategy sets a certain target utility range, which is a subset of all acceptable outcomes. From these outcomes, the opponent model selects the offers that are also good for the opponent.

every component. As we cannot possibly explore the entire space of all possible components, we need a representative selection for every component, ranging from baseline techniques to state of the art techniques. For every component, we aimed to select many variants, given that they were designed for a negotiation setting consistent with ours, with publicly available code, and generic enough so that we could freely interchange them with arbitrary other components. To ensure as much variety as possible, we aimed to include components designed by various research teams.

Our range of components are shown in Table 2. We selected 11 different bidding strategies in total. We selected state of the art bidding strategies from the top three strategies of the ANAC 2011 and 2012 competitions, and as baseline bidding tactics, we included the time-dependent tactics *Boulware* (with concession rate $e = 0.2$), *Conceder Linear* ($e = 1$), and *Conceder* ($e = 2$) taken from [9] and behavior-dependent strategies such as *Nice Tit for Tat*, *Absolute Tit for Tat* [2] and *Relative Tit for Tat* [9].

All bidding strategies were combined with six different opponent models. We selected five state of the art opponent models from ANAC that we could freely combine with arbitrary bidding strategies. To ensure a fair representation of OM's, we selected from two main categories; *Bayesian models*, which use Bayesian learning techniques to create and update a model of the opponent's preference profile [23] and *Frequency models*, which keep track of how often certain items are requested by the opponent. As a baseline, we chose *No Model*, which selects bids for the opponent at random.

For the acceptance criteria, we selected a number of sophisticated strategies from top ANAC agents, together with simple baseline acceptance policies, such as the very simple acceptance criterion *Constant* (c) ($c \in \{0.6, 0.7, 0.8, 0.9\}$), which accepts exactly when the utility of the opponent's offer is higher than a constant threshold c , and *Next* (α, β) ($\alpha \in \{1.0, 1.1, 1.2\}$, $\beta \in \{0, 0.1, 0.2\}$), which accepts when $\alpha \cdot u' + \beta \geq u$, where u is the utility of the bid that is ready to be sent out and u' is the utility of the opponent's offer.

For our opponent pool, we needed to make a selection of agents to make running the experiment feasible. We selected a representative set of 7 agents, ranging from baseline strategies (*Boulware* and *Conceder Linear*) to some of the top performing agents from the ANAC competitions (the top 2 agents from ANAC 2011: *Hard-Headed* and *Gahboninho* [2], and number 1 and 3 from ANAC 2012: *CUHKAgent* [12] and *The Negotiator Reloaded*).

The negotiation scenarios were chosen on the basis of the following characteristics: *domain size* (number of possible bids), *bid distribution* (average distance of all bids to the nearest Pareto-optimal bid), and the

opposition of the domain (the distance from the Kalai-Smorodinsky point to complete satisfaction). We picked 5 well-known negotiation scenarios used in ANAC such that every characteristic varies between *high*, *medium* and *low* (see Table 1).

Scenario name	Size	Bid distrib.	Opposition
ADG [2]	15625 (<i>med.</i>)	0.136 (<i>low</i>)	0.095 (<i>low</i>)
Grocery [2]	1600 (<i>med.</i>)	0.492 (<i>high</i>)	0.191 (<i>med.</i>)
Itex-Cypress [4]	180 (<i>low</i>)	0.222 (<i>med.</i>)	0.431 (<i>high</i>)
Laptop [2]	27 (<i>low</i>)	0.295 (<i>med.</i>)	0.178 (<i>med.</i>)
Travel [4]	188160 (<i>high</i>)	0.416 (<i>high</i>)	0.230 (<i>med.</i>)

Table 1. Characteristics of the negotiation scenarios.

We created a large number of negotiation agents by combining all components into full negotiation agents; i.e., we created a pool of 11 bidding strategies \times 6 opponent models \times 24 acceptance criteria, which amounts to 1584 negotiation agents in total.² We let all of them play against the 7 opponents in each scenario listed in Table 1, and we kept track of their obtained utility in every negotiation session. To run our tournament we used the GENIUS framework [17], which is a well-established negotiation environment where automated agents can negotiate in a bilateral multi-issue negotiation setting. Since not all the negotiation strategies are deterministic, the tournament setup was run 5 times in order to reduce the amount of variance in the data, resulting in 277200 negotiation sessions in total.

Bidding, Opponent Modeling, and Accepting Components

Bidding Strategy	Acceptance Criteria
CUHK Agent	Agent K2 [2]
HardHeaded [2]	AgentLG
The Negotiator Reloaded	AgentMR
IAMhaggler2012	BOA Constrictor
AgentLG	BRAMAgent2 [2]
Nice Tit For Tat [2]	Constant(<i>c</i>)
Time-dependent tactics [9]	HardHeaded [2]
Absolute Tit For Tat [9]	IAMhaggler2012
Relative Tit For Tat [9]	Next(α, β)
	Nice Tit For Tat [2]
Opponent Model	OMAC Agent
Agent Smith (Frequency) [4]	The Fawkes
HardHeaded (Frequency) [2]	The Negotiator [2]
IAMhaggler (Bayesian) [2]	The Negotiator Reloaded
NASH Agent (Frequency)	
The Negotiator Reloaded (Bayesian)	
No Model	

Table 2. All negotiation strategy components used in experimental setup.

4 Measuring the Contribution of Strategy Components

To determine the importance of a particular agent component in relation to another, and to measure the interactions between them, we use the notion of *effect size*. The effect size is a statistical measure to quantify the amount of difference between various controlled variables and measures their effect on the dependent variable, which in our case is the average outcome utility of the agent. In this way, the effect size expresses *how* significant or important a variable is [7]. This goes beyond the question of *whether* a specific variable is significant or important, which can be answered with standard statistical significance testing. We use the relative proportion of variation, known as the η^2

² Note that this set also includes already existing agents such as *HardHeaded* and *The Negotiator Reloaded*, since their components occur in all three groups.

measure. The η^2 measure rates the variance that a variable introduces using the *Sum of Squares*. The *Total Sum of Squares* (SS_{Total}) expresses the total dispersion of all data points; i.e., the total variance found within the data, which can be calculated from the differences between the data points x_i and the total mean \bar{x} :

$$SS_{Total} = \sum_{i=1}^n (x_i - \bar{x})^2. \quad (1)$$

The *Sum of Squares Between Groups* (SS_b) determines the variance between the different *groups* and is calculated as follows:

$$SS_b = \sum_{j=1}^{|G|} n_j (\bar{x}_j - \bar{x})^2, \quad (2)$$

where the \bar{x}_j represents the group mean, $|G|$ represents the number of groups and n_j is the number of data points in group j .

To calculate the η^2 value for each variable i , we use equation (3), where SS_{bi} is the sum of squares between groups for variable i :

$$\eta_i^2 = \frac{SS_{bi}}{SS_{Total}} \quad (3)$$

To determine the contribution of each component, we calculated the η^2 of our three components, using as input the utilities of all component combinations averaged over all the runs, domains and opponents. The calculated η^2 values are presented in Table 3 and visualized in Figure 2.

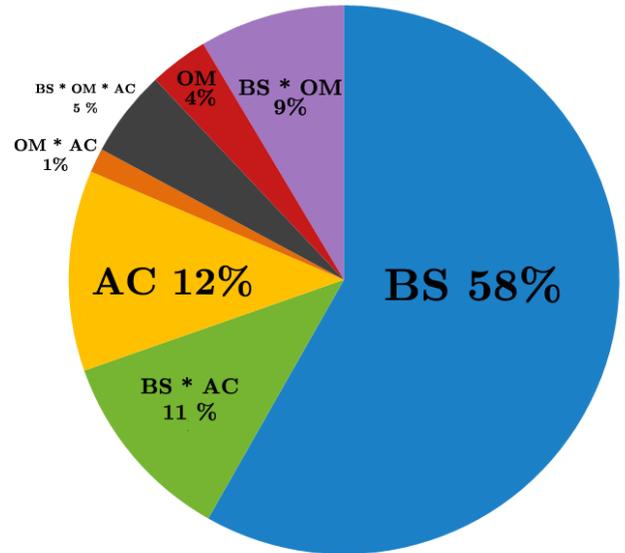


Figure 2. Visual representation of the contribution of components

Component	η^2	Standard Deviation		
		Runs	Domains	Opponents
Bidding Strategy (BS)	0.582	0.003	0.118	0.163
Opponent Model (OM)	0.035	0.002	0.020	0.037
Acceptance Criteria (AC)	0.118	0.003	0.121	0.071
BS * AC	0.114	0.003	0.023	0.082
OM * AC	0.014	0.001	0.011	0.016
BS * OM	0.085	0.004	0.040	0.040
BS * OM * AC	0.051	0.002	0.037	0.051

Table 3. The η^2 measure for every component and the standard deviation of η^2 over the different runs, domains and opponents.

From the calculated contributions, we can observe that the BS is by far the most important component, accounting for 58% of the variation in the negotiation strategy's performance, which is significantly higher (*one-tailed t-test*, $p < 0.01$) than the other components. Recall that the BS controls the concession speed of the agent, thereby managing the rate with which it moves towards the opponent according to its own utility. This is expected to have a huge impact on the final outcome, as it determines the subset of the outcome space an agreement can possibly be made in. We can also observe that the variance of the BS contribution is rather high, especially with respect to the opponent pool; this is explored further in Section 4.1.

Note that our method allows us to express the effectiveness of a specific choice for any given strategy component. For example, we can fix the bidding strategy to *CUHK Agent*, and calculate its average obtained utility when combined with all OM's and AC's and then compare this with the average utility of another choice of the BS, for example *Conceder*. In fact, the difference in this case is particularly large, because using this method, *CUHK Agent* (bidding strategy of winner of ANAC 2012) obtains the highest score of 0.79, while *Conceder* (a baseline bidding strategy) has the lowest score of 0.63. Such a utility difference means the difference between place 1 and place 8 in ANAC 2012, which gives a good indication of how important the bidding strategy is.

The second most important component is the AC, with a contribution of 12%, which is significantly (*one-tailed t-test*, $p < 0.01$) higher than the OM. Also in this case, the baseline AC's (e.g., *Constant(c)* and *Next(α, β)*) obtain the lowest scores (0.65), while state of the art AC's from the top ANAC agents score highest (0.75). The contribution of the AC's still comes out low because these are only the extreme scores, and most AC scores are clustered around the group average (0.71).

Finally, the OM makes a surprisingly small contribution to the performance of a negotiation strategy, accounting for only 4% of the explained variance. There is only a small difference between using *No Model* (0.69) and any of the state of the art models (0.72). This is still a significant difference and could make the difference between place 1 and place 5 in ANAC 2012. However, there is almost no difference between the learning methods themselves, which indicates that once an agent employs a reasonable opponent model, there is not much more to gain after that.

4.1 The Influence of the Opponent

Table 3 also presents the standard deviation of η^2 over the runs, domains and opponents for each of the components. This gives an indication of how much an additional run, domain or opponent would affect the average contribution of each component. Adding more runs will have little effect on the components' contribution, but the same cannot be said for the domains and opponents, which both have high standard deviations, especially for the BS and AC, which indicates that these components highly interact with the domain and opponent.

The opponent is the most important source of variance when considering component contribution. We take a closer look at this phenomenon by focusing on the subset of opponents that employ time-dependent tactics, comparing the η^2 values when negotiating against each of them. The advantage of focusing on the time-dependent opponents is that they are from the same family of tactics, so that only one factor is altered between the different opponents, namely the rate at which they concede. To get a more complete picture, we add two more time-dependent tactics to our opponent pool, namely *Extreme Boulware* ($e = 0.02$) and *Extreme Conceder* ($e = 5$).

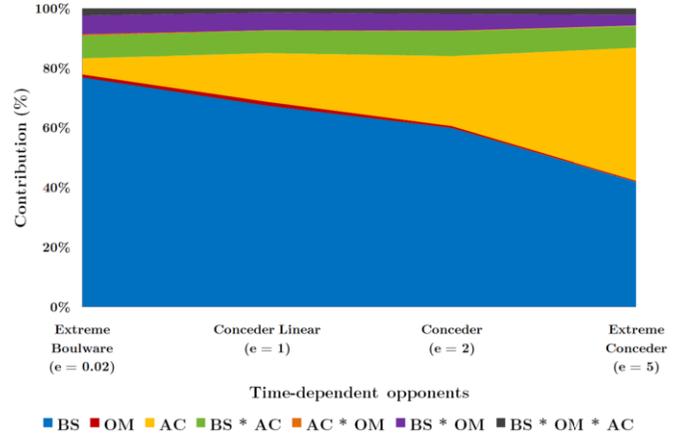


Figure 3. The η^2 values of all components against different time-dependent opponents.

The η^2 values are presented in Figure 3. They show a clear trend for the contribution of the strategy components. As the value of e increases (raising the cooperativeness of the opponent) the contribution of the BS decreases dramatically, from 77% to 42%. Through closer inspection of the negotiation dynamics between the agents, we are able to explain this trend.

When the opponent concedes very little, it is up to the BS of the agent to yield to the opponent to avoid the consequences of no agreement. This places a lot of importance on the BS, because the speed at which the agent concedes dictates the agreement that will be reached, hence its large η^2 value. In these situations, the importance of having an effective OM is at its peak, since it is both more challenging to achieve an acceptable outcome and it is harder to learn the opponent's preferences. Also, when an opponent makes very small concessions, the offers differ very little in utility. This makes the role of the AC less important, because the moment the agreement is reached is not as significant. On the other hand, when the opponent is very cooperative, the choice between accepting and waiting for more concessions can make a big difference in the achieved outcome. Therefore, the AC of the agent and the BS of the opponent will primarily dictate the meeting point between the two agents. This is why, as the opponent concedes more, the contribution of the AC increases from 5% to 47% at the cost of the BS.

4.2 Interaction Effects

Agent components do not perform their function in isolation. First of all, they interact with each other: for example, a learning method may be far less effective when the bidding strategy does not select enough 'exploratory offers' to learn more about the opponent's preferences. There are also interaction effects with the environment, such as the negotiation domain (e.g., learning about the opponent may be harder in big contract spaces), and the opponent (e.g., a good acceptance strategy is more important when the opponent is likely to make attractive offers).

The presence of these interactions requires a more thorough analysis, because the impact of one component can depend on the level of another. We denote interaction effects as $C_1 * C_2$, where C_i represents a component. The term *interaction effect* is used to quantify the effects of the interactions between the components. To determine the η^2 for an interaction we need to calculate the $SS_{b_{c_1 * c_2}}$ as follows:

$$SS_{bC_1 * C_2} = \sum_{i=1}^{|C_1|} \sum_{j=1}^{|C_2|} (x_{ij} - \bar{x}_i - \bar{x}_j + \bar{x})^2 \quad (4)$$

This is then divided by SS_{Total} to determine the η^2 value for the interaction. The component interaction effects are listed at the bottom of Table 3.

There are two important interactions that deserve some attention. One is the interaction between the bidding strategy and the acceptance criteria ($BS * AC$), which is the highest interaction effect of 11%. The BS and AC can be viewed as two simultaneous processes that can each result in a potential agreement. One process consists of offering bids that are appealing to the opponent in the hope that it will be accepted (BS), the other consists of receiving offers from the opponent and deciding whether these should be accepted or not (AC). These two components complement each other, as a good AC can compensate for a bad performing BS by accepting bids from the opponent, while a good BS can compensate for a bad AC by offering enticing counter bids.

Another important interaction is between the bidding strategy and the opponent model ($BS * OM$), which, with 9%, is the second highest interaction effect. The OM directly influences the BS by aiding it in offering bids that are appealing to the opponent, thereby improving the chances of an agreement. Conversely, the effectiveness of an OM depends on the BS. Should the BS not use the OM to its full potential, then its effectiveness will be diminished. For example, *BRAM Agent* [11], a participant of ANAC 2011, presents the OM with only a small selection of bids, combining only the ten most recent offers of the opponent, which reduces the effectiveness of the OM.

Interaction also exists between the acceptance criteria and the opponent model ($AC * OM$); however, this contribution is rather small (1%). This is because the purpose of the OM is to model the opponent's attributes, while the AC typically acts according to the agent's own utility.

There are also three way interactions involving all three strategy components ($BS * OM * AC$), which account for 5% of the variance (calculated analogously to equation (4)). For example, there are agents that not only use the OM to determine the best bid to offer to the opponent, but also use to determine their target utility. More sophisticated AC's make use of this target utility and are thus also affected by the OM, causing a three way interaction.

4.3 Combining the Best Components

Given the interaction effects between the different components, it is not immediately clear whether combining the best components together result in the most effective negotiation strategy. To test this, we used the *Jenks Natural Break Algorithm* [14] to divide each of the three components listed in Table 2 into three different performance groups (*High*, *Medium*, and *Low*). We combined these groups into 27 agent sets, of which we then tested the average utility (see Table 4).

Our results show that the agent set that employed the best performing components in isolation (i.e., the agents that were comprised of components that were all in the *High* category) have the highest average utility, performing significantly better than all other agent groups (*one-tailed t-test*, $p < 0.01$). This demonstrates that independently optimizing the individual components indeed is a feasible approach for developing negotiating agents. The agent group that used the worst performing components (*Low*, *Low*, *Low*) has the lowest average utility. The agent group (*Medium*, *Medium*, *Medium*) is found somewhere in the middle of the rankings, as expected.

Component Combinations				
BS	OM	AC	Avg. Util.	Std Dev
<i>High</i>	<i>High</i>	<i>High</i>	0.790	0.001
<i>Medium</i>	<i>Medium</i>	<i>Medium</i>	0.706	0.003
<i>Low</i>	<i>Low</i>	<i>Low</i>	0.608	0.001

Table 4. Rankings of three agent sets that were created by combining components with *High*, *Medium* and *Low* individual performance

5 Related Work

Combining different negotiation strategies is not a new idea, but as far as the authors are aware, studying the effects of a large group of state of the art negotiation components on the net performance of a negotiating agent is novel. For instance, Faratin et al. [9] analyze the performance of 'pure' negotiation tactics, considering them as components around which full strategies can be built. They discuss the possibility of linearly combining these pure tactics, but they do not investigate this any further.

Sierra et al. obtain promising results in [20] with a negotiation strategy that combines two components: a concession based strategy (either time-based or behavior-based [9]) that decreases a utility threshold to achieve an agreement, and a trade-off strategy [10] that searches for a satisfactory proposal. Our work differs with Sierra et al., as we consider a much wider array of agents of which we are able to not only change the bidding strategy, but also the opponent model and acceptance criteria.

Another negotiation strategy, proposed by Ilany and Gal approach this differently [13]; instead of combining different strategies during one negotiation session, they select the best strategy from a predefined set of agents, based on the characteristics of the domain. Through machine learning this agent is optimized to choose the best strategy for that particular domain. This work, combining existing strategies into one, is similar to our approach. However, we combine different generic *components* of existing strategies, instead of whole negotiating agents.

Some authors use genetic algorithms (GA) to automatically combine certain tactics or strategies. This approach is different to ours, however they do share certain traits, as they view a strategy consisting of different elements and combine them in order to produce a better performing strategy [19, 22]. For example, Matos et al. [19] use a set of baselines negotiation strategies which consist of time dependent, resource dependent and behavior dependent strategies [9] and combine them linearly with the help of GA. Tu et al. [22] on the other hand use *Finite State Machines* to represent strategies and uses GA to modify them as time passes. GA approaches also typically alter the parameters of a negotiation strategy in order to find the best strategy (e.g. Eymann [8]). In our work, we focus on agent components rather than parameters or strategies, and we also investigate their importance.

To be able to find the most important component we employ the statistical measure η^2 . Other multivariate techniques such as *Principle Component Analysis* (PCA) and *Factor Analysis* are also often used to find important or significant variables in the observed data. They do this by reducing the amount of variables, extracting important information and expressing them as a new set of variables known as *principal components* and *factors* respectively [1, 15]. These new variables are a linear combination of the original variables where the first has the largest variance possible, the second the second largest variance, etc. From this it is possible to determine that the first variable is the most important followed by the second, etc. The difference between PCA and factor analysis is that the latter assumes that there

is an underlying causal model and thus attempts to reduce the number of dimensions by using regression modeling techniques, while the former does not use any explicit model [16]. The aforementioned methods are however not applicable to answer our research questions, since we are interested in the importance of our original component set, rather than their linear combinations.

6 Conclusion and Future Work

This paper investigates the performance effects of combining different instantiations of key components of existing negotiating agents, namely the bidding strategy, the opponent model, and the acceptance criteria. For this purpose, we analyzed the key components of a large set of both baseline and state of the art agents. We analyzed each component independently as well as in combination with the other components, using the measure η^2 to quantify their effect.

We found that combining the best agent components indeed results in the strongest agents. This shows that the three-component view of a negotiating architecture not only provides a useful basis for developing negotiating agents, but also enables the analytical approach of optimizing its individual components. By varying the key components of automated negotiators, we are able to demonstrate the contribution of each component to the negotiation result, and thus analyze the significance of each. Moreover, we are able to study the interaction effects between them. With respect to the impact of each of the three key components, we found that the bidding strategy is by far the most important to consider, followed by the acceptance criteria and finally followed by the opponent model.

The low importance of the opponent model is surprising, as the importance of opponent models has been shown on many occasions. Rather, we argue that in our setting, existing learning techniques already do quite well, and no significant effect is to be expected by further improving on the currently existing preference learning techniques for our setting. This is in contrast to the bidding strategy and acceptance criteria, which have a substantial influence on the agent's performance. To put it another way: our results indicate that the majority of the implementation effort of an agent designer should be focused on the bidding and accepting strategy.

This brings us to our goal of shaping the research agenda on negotiating agent design. Based on our results we recommend that research into bidding strategy and acceptance criteria should stay on the agenda. The state-of-the-art in opponent preference modeling is already so good, that we recommend to focus the attention on the research into automated learning of the bidding strategy and acceptance criteria of the opponent.

REFERENCES

- [1] Hervé Abdi and Lynne J. Williams, 'Principal component analysis', *Wiley Interdisciplinary Reviews: Computational Statistics*, **2**(4), 433–459, (2010).
- [2] Tim Baarslag, Katsuhide Fujita, Enrico H. Gerding, Koen Hindriks, Takayuki Ito, Nicholas R. Jennings, Catholijn Jonker, Sarit Kraus, Raz Lin, Valentin Robu, and Colin R. Williams, 'Evaluating practical negotiating agents: Results and analysis of the 2011 international competition', *Artificial Intelligence*, **198**(0), 73 – 103, (May 2013).
- [3] Tim Baarslag, Koen Hindriks, Mark Hendriks, Alex Dirkzwager, and Catholijn Jonker, 'Decoupling negotiating agents to explore the space of negotiation strategies', in *Proceedings of The Fifth International Workshop on Agent-based Complex Automated Negotiations (ACAN 2012)*, (2012).
- [4] Tim Baarslag, Koen Hindriks, Catholijn M. Jonker, Sarit Kraus, and Raz Lin, 'The first automated negotiating agents competition (ANAC 2010)', in *New Trends in Agent-based Complex Automated Negotiations, Series of Studies in Computational Intelligence*, eds., Takayuki Ito, Minjie Zhang, Valentin Robu, Shaheen Fatima, and Tokuro Matsuo, pp. 113–135, Berlin, Heidelberg, (2012). Springer-Verlag.
- [5] Tim Baarslag and Koen V. Hindriks, 'Accepting optimally in automated negotiation with incomplete information', in *Proceedings of the 2013 International Conference on Autonomous Agents and Multi-agent Systems, AAMAS '13*, pp. 715–722, Richland, SC, (2013). International Foundation for Autonomous Agents and Multiagent Systems.
- [6] Réal Carbonneau, Gregory E. Kersten, and Rustam Vahidov, 'Predicting opponent's moves in electronic negotiations using neural networks', *Expert Systems with Applications*, **34**(2), 1266–1273, (February 2008).
- [7] Robert Coe, 'It's the effect size, stupid: What effect size is and why it is important', in *British Educational Research Association Conference*. Education-line, (2002).
- [8] Torsten Eymann, 'Co-evolution of bargaining strategies in a decentralized multi-agent system', in *AAAI fall 2001 symposium on negotiation methods for autonomous cooperative systems*, pp. 126–134, (2001).
- [9] Peyman Faratin, Carles Sierra, and Nick R. Jennings, 'Negotiation decision functions for autonomous agents', *Robotics and Autonomous Systems*, **24**(3-4), 159 – 182, (1998). Multi-Agent Rationality.
- [10] Peyman Faratin, Carles Sierra, and Nick R. Jennings, 'Using similarity criteria to make issue trade-offs in automated negotiations', *Artificial Intelligence*, **142**(2), 205 – 237, (2002). International Conference on MultiAgent Systems 2000.
- [11] Radmila Fishel, Maya Bercovitch, and Ya'akov(Kobi) Gal, 'Bram agent', in *Complex Automated Negotiations: Theories, Models, and Software Competitions*, eds., Takayuki Ito, Minjie Zhang, Valentin Robu, and Tokuro Matsuo, volume 435 of *Studies in Computational Intelligence*, 213–216, Springer Berlin Heidelberg, (2013).
- [12] Jianye Hao and Ho-Fung Leung, 'ABiNeS: An adaptive bilateral negotiating strategy over multiple items', in *Proceedings of the The 2012 IEEE/WIC/ACM International Joint Conferences on Web Intelligence and Intelligent Agent Technology - Volume 02, WI-IAT '12*, pp. 95–102, Washington, DC, USA, (2012). IEEE Computer Society.
- [13] Litan Ilany and Yakov Gal, 'Algorithm selection in bilateral negotiation', in *Proceedings of the Twenty-Seventh AAAI Conference on Artificial Intelligence (AAAI 2013)*, (2013).
- [14] George F. Jenks and Dept. of Geography University of Kansas, *Optimal Data Classification For Choropleth Maps*, Occasional paper, University of Kansas, 1977.
- [15] Ian Jolliffe, *Principal Component Analysis*, John Wiley & Sons, Ltd, 2005.
- [16] John T. Scott Jr., 'Factor analysis and regression', *Econometrica*, **34**(3), pp. 552–562, (1966).
- [17] Raz Lin, Sarit Kraus, Tim Baarslag, Dmytro Tykhonov, Koen Hindriks, and Catholijn M. Jonker, 'Genius: An integrated environment for supporting the design of generic automated negotiators', *Computational Intelligence*, **30**(1), 48–70, (2014).
- [18] Alessio Lomuscio, Michael Wooldridge, and Nicholas Jennings, 'A classification scheme for negotiation in electronic commerce', in *Agent Mediated Electronic Commerce*, eds., Frank Dignum and Carles Sierra, volume 1991 of *Lecture Notes in Computer Science*, 19–33, Springer Berlin Heidelberg, (2001).
- [19] Noyda Matos, Carles Sierra, and Nick R. Jennings, 'Determining successful negotiation strategies: an evolutionary approach', in *Multi Agent Systems, 1998. Proceedings. International Conference on*, pp. 182–189, (1998).
- [20] Raquel Ros and Carles Sierra, 'A negotiation meta strategy combining trade-off and concession moves', *Autonomous Agents and Multi-Agent Systems*, **12**, 163–181, (2006).
- [21] Carles Sierra, Peyman Faratin, and Nick R. Jennings, 'A service-oriented negotiation model between autonomous agents', in *Proceedings of the 8th European Workshop on Modelling Autonomous Agents in Multi-Agent World, MAAMAW-97*, eds., M. Boman and W. van de Velde, volume 1237 of *Lecture Notes in Artificial Intelligence*, pp. 17–35. Springer-Verlag, (1997).
- [22] M. Tuan Tu, Eberhard Wolff, and Winfried Lamersdorf, 'Genetic algorithms for automated negotiations: A fsm-based application approach', in *Proceedings of the 11th International Workshop on Database and Expert Systems Applications, DEXA '00*, pp. 1029–, Washington, DC, USA, (2000). IEEE Computer Society.
- [23] Dajun Zeng and Katia Sycara, 'Bayesian learning in negotiation', *International Journal of Human Computer Systems*, **48**, 125–141, (1998).