

Predicting the Performance of Opponent Models in Automated Negotiation

Tim Baarslag, Mark Hendriks, Koen Hindriks, Catholijn Jonker
Interactive Intelligence Group
Delft University of Technology
Mekelweg 4, Delft, The Netherlands
T.Baarslag@tudelft.nl

Abstract—When two agents settle a mutual concern by negotiating with each other, they usually do not share their preferences so as to avoid exploitation. In such a setting, the agents may need to analyze each other's behavior to make an estimation of the opponent's preferences. This process of *opponent modeling* makes it possible to find a satisfying negotiation outcome for both parties. A large number of such opponent modeling techniques have already been introduced, together with different measures to assess their quality. The quality of an opponent model can be measured in two different ways: one is to use the agent's performance as a benchmark for the model's quality, the other is to directly evaluate its accuracy by using similarity measures. Both methods have been used extensively, and both have their distinct advantages and drawbacks. In this work we investigate the exact relation between the two, and we pinpoint the measures for accuracy that best predict performance gain. This leads us to new insights in how to construct an opponent model, and what we need to measure when optimizing performance.

Keywords—Intelligent agents, Multiagent systems, Machine learning

I. INTRODUCTION

A negotiation between two agents is a dialogue in which both agents aim to reach an agreement better than their status quo. A major challenge in automated negotiation is that agents usually keep their preference information private to avoid exploitation [1], [2]. When the agents have limited knowledge of the other's preferences, the agents may fail to reach an optimal outcome as they cannot take the opponent's desires into account [3].

In order to improve the efficiency of the negotiation and the quality of the outcome, agents may construct a *model* of the opponent's preferences, which aids them in estimating the information that is kept private [1]–[3]. Over time, a large number of such opponent models have been introduced, based on different learning techniques and underlying assumptions, and multiple methods have been used to compare their quality. The different evaluation methods for opponent models make it hard to compare different approaches, as each method has its unique scope of application, together with different advantages and drawbacks. From an engineering perspective, it still remains unclear which opponent model to choose in a particular negotiation setting.

Traditionally, there are two popular ways to measure the quality of an opponent model:

- 1) *Performance measures* evaluate the quality of the outcome, usually measured in utility gain, or distance of the agreement to the Pareto frontier. With this method, the success of an opponent model is expressed in terms of the negotiation *result* (as opposed to the whole negotiation *process*).
- 2) *Accuracy measures* aim to determine the quality of a model in a more fundamental way, by quantifying how well the opponent model *represents* the real preferences of the opponent, using a certain similarity measure. An example is the correlation between the estimated and the real outcome space, or the percentage of correctly inferred Pareto optimal outcomes.

There are various authors that evaluate their opponent model with *performance measures* (e.g. [1]–[3]). Using a performance measure has one very important quality: it measures exactly what needs to be optimized, namely the net effect an opponent model has on the negotiation result. On the other hand, because performance measures are only able to demonstrate improvement of the end result, they may not provide insight into why or how an opponent model works; that is, they measure the result obtained by the negotiation agent as a whole, of which the opponent model is only a single component. This makes the performance measure very sensitive to the specifics of the experimental setup. Moreover, there is usually no clear upper bound in performance gain, so it remains unclear what the highest attainable result is.

Other authors prefer to use *accuracy measures* to evaluate their model (for example [4]–[7]). The main advantage of this approach is that it directly assesses the quality of a model, independent of other factors such as bidding strategy or acceptance strategy. Secondly, it is easier to compare accuracy results between different experimental setups, and to track the accuracy of a model over the course of the negotiation. This, in turn, can reveal valuable information about the reasons for a model's success.

There are also drawbacks of using accuracy measures in negotiation, two of which we will address in this paper. First, it is currently unclear what effect a more accurate opponent model has on the negotiation outcome. It could very well be that from some point, increased accuracy does not translate into better performance. An 80% accurate model for example, could perform just as well as a perfect model.

Second, there are many accuracy measures to choose from, and it is currently unknown which accuracy measure should be selected to ensure a good overall end result; that is, we would like to know what accuracy measure best *predicts* an improvement in performance.

Our work bridges the gap between both approaches by considering opponent models from both a performance and an accuracy perspective. We first test many current opponent modeling techniques in different negotiation settings, measuring both their accuracy through time and their performance. We then analyze how changes in accuracy translate into performance differences. Moreover, we review all ways to measure the accuracy of an opponent model, and we pinpoint the best *predictors* for good performance.

The remainder of this paper is organized as follows. Section II describes the setting and terminology used throughout the paper. In Section III we introduce our research questions and experiments, followed by a discussion of the results in Section IV. Finally, Section V discusses related work and Section VI provides directions for future work.

II. PRELIMINARIES

In this work we focus on bilateral automated negotiations, where agents take turns in exchanging offers using the alternating offers protocol. The agents seek to reach an agreement while at the same time aiming to maximize their own utility. A negotiation scenario consists of the negotiation domain, which specifies the setting and all possible bids (also called the *outcome space*, denoted by Ω), together with a privately-known preference profile for each party. A preference profile is described by a utility function $u(\omega)$, which maps each outcome $\omega \in \Omega$ to a utility in the range $[0, 1]$. We assume a common discrete time line, with a deadline after a specified number of rounds N . Both agents receive utility 0 if they do not succeed in reaching an agreement in time.

We discuss opponent models that attempt to estimate the opponent’s utility function $u_{op}(\omega)$ while relying solely on the information gathered during the negotiation. Many existing models and measures assume the agent’s utility function is linear additive [3], [8]–[10]. Alternatively, agents may have non-linear preferences (e.g., [11]); however, we restrict ourselves to linear additive preference profiles in this work, as to the best of our knowledge, there exists no large set of comparable models for non-linear preferences.

Linear additive functions make explicit that different issues can be of different importance to a negotiating agent and can be used to efficiently calculate the utility of a bid. The utility $u(\omega)$ of an outcome $\omega \in \Omega$ is computed as a weighted sum (as specified by the *issue weights* w_i) of *value weights* $e_i(\omega_i)$:

$$u(\omega) = \sum_{i=1}^n w_i \cdot e_i(\omega_i). \quad (1)$$

III. MEASURING THE QUALITY OF OPPONENT MODELS

As outlined in the introduction, the aim of this work is to answer three research questions:

- 1) How does the accuracy of opponent models depend on negotiation factors, such as domain size, or time?
- 2) What is the relationship between the accuracy of an opponent model and its expected performance gain?
- 3) What accuracy measures are the best predictors for performance gain?

To answer these questions, we first outline our selection of opponent models (Section III-A) and the accuracy measures incorporated in our method (Section III-B). Next, we discuss the experimental setup (Section III-C and III-D).

A. Selection of Opponent Models

We compare a large set of state of the art opponent models, which were applied in the Automated Negotiating Agents Competition (ANAC) [12], [13]. ANAC is a yearly international competition in which negotiation agents compete on a set of scenarios that are unknown beforehand. Our reason for including this set of models is threefold: first, they represent the state of the art; second, to our knowledge, they are the largest set of techniques designed for one common setting consistent with ours; and finally, their code is publicly available. Table I gives an overview of all models evaluated in this work, including three theoretical baselines. We distinguish four types of opponent models:

- 1) *Bayesian models* estimate the opponent’s preferences by first generating a set of candidate preference profiles. Next, Bayesian learning is used to continually update the model, based on certain assumptions about the opponent’s concession function. For these models, we also include *Perfect* variants, which use perfect knowledge about the opponent’s concessions, but are still unaware of the opponent’s exact preferences.
- 2) *Frequency models* estimate both the issue and value weights separately. The issue weights are estimated based on how often their value is changed between sequential bids. The value weights are derived from the frequency they are offered.
- 3) *Value models* are similar to the frequency models, except that the issue weights are assumed to be equal.
- 4) *Theoretical baselines* are used to compare the quality of the models. The *Perfect Model* and *Worst Model* act as an upper and lower bound on quality respectively, while the *Opposite Model* functions as a baseline, since it serves as a good initial guess of the opponent’s preferences.

Each model was isolated from existing negotiation agents (as indicated in Table I), and then generalized to be compatible with any bidding strategy, as in [14]. The advantage of using this approach is that we can interchange the opponent modeling component of each negotiation strategy, so that we

can compare the performance of different opponent models while keeping the bidding and acceptance strategy fixed.

Bayesian Models	
<i>Bayesian Scalable Model</i> [3]	Estimates the issue and value weights separately, using Bayesian learning. The opponent is assumed to concede a constant amount per round.
<i>IAMhaggler Bayesian Model</i> [8]	A Bayesian model in which the opponent is assumed to use a particular time-dependent strategy and only unique bids are used to update the model.
Frequency Models	
<i>HardHeaded Frequency Model</i> [9]	Learns the issue weights based on how often the value of an issue changes. The value weights are estimated based on the frequency they are offered.
<i>Smith Frequency Model</i> [10]	Learns the value weights based on frequency they are offered. The issue weights are estimated based on the distribution of the values.
<i>Agent X Frequency Model</i>	A variant of the <i>HardHeaded Frequency Model</i> that takes the opponent's tendency to repeat bids into account.
<i>N.A.S.H. Frequency Model</i>	Learns the issue weights based on how often the best value for each issue is offered. The value weights are estimated based on their frequency.
Value Models	
<i>Agent LG Value Model</i>	Estimates the value weights based on the frequency they are offered.
<i>CUHKAgent Value Model</i>	Counts how often each value is offered. The utility of a bid is the sum of the score of its values divided by the best possible score. The model only uses the first 100 unique bids for its estimation.
Theoretical Baselines	
<i>Opposite Model</i>	Defines the opponent's utility as one minus the agent's utility.
<i>Perfect Model</i>	Perfect knowledge of the opponent's preferences.
<i>Worst Model</i>	Defines the estimated utility as one minus the real utility.

Table I
OVERVIEW OF OPPONENT MODELS.

B. Selection of Accuracy Measures

We compare the accuracy of opponent models by evaluating how well the models estimate the opponent's preferences when provided with various negotiation traces. In effect, we treat the opponent model as an isolated component that receives offers as input, and yields an estimate of the opponent's preference profile as output, which (hopefully) gets increasingly accurate with every processed bid.

When we assess opponent model accuracy, we require an accuracy measure that quantifies the similarity between the opponent's actual preference profile u_{op} and the estimation u'_{op} . As there is no current standard for accuracy measures, we surveyed all accuracy measures currently in use, as shown in Table II. The first two sets of measures are derived from literature, to which we have added a set of metrics based on the Pareto optimal frontier.

Many of the measures in Table II are formulated in terms of the combined properties of the opponent's utility space and the agent's own utility space – together called the *bid space*. We define the *real bid space* B as

$$B = \{(u_{own}(\omega), u_{op}(\omega)) \mid \omega \in \Omega\}.$$

Outcome Space	
<i>Pearson correlation of bids</i> [6]	Pearson correlation coefficient between real and estimated preferences.
<i>Ranking distance of bids</i> [4], [6]	Ranking distance between real and estimated preference.
<i>Average difference between bids</i>	Average difference between the real and estimated utility of all bids.
Issue Weights	
<i>Pearson correlation of issue weights</i> [6]	Pearson correlation coefficient between real and estimated issue weights.
<i>Ranking distance of issue weights</i> [6]	Ranking distance between real and estimated issue weights.
<i>Average difference between issue weights</i> [7]	Average difference between the real and estimated issue weights.
Pareto Frontier	
<i>Average difference of Pareto frontier</i>	The average difference between the real and estimated utility of the Pareto bids.
<i>Percentage of found Pareto bids</i>	Percentage of real Pareto bids that are also estimated to be a Pareto bid.
<i>Percentage of correct Pareto bids</i>	Percentage of estimated Pareto bids that are also real Pareto bids.
<i>Difference in Pareto frontier surface</i>	Absolute difference in surface under the real and estimated Pareto frontier.

Table II
OVERVIEW OF ACCURACY MEASURES.

The *estimated bid space* B' is defined in terms of the estimated opponent utility function u'_{op} :

$$B' = \{(u_{own}(\omega), u'_{op}(\omega)) \mid \omega \in \Omega\}.$$

To quantify how well u'_{op} approximates the opponent's preferences, we might consider the differences between u'_{op} and u_{op} directly. Alternatively, we can analyze the resulting bid spaces B and B' , or we might concentrate on subsets.

Outcome space accuracy measures quantify the difference between u_{op} and u'_{op} by considering all bids in the outcome space Ω . A straightforward measure that has been used is the *average distance between bids* metric, which calculates the average absolute difference between u_{op} and u'_{op} over Ω . However, as models are usually only concerned with the ranking of outcomes, a more suitable metric is the *Pearson correlation of bids* that measures the correlation between two outcome spaces, which is defined as follows:

$$d_p(u_{op}, u'_{op}) = \frac{\sum_{\omega \in \Omega} (u_{op}(\omega) - \overline{u_{op}})(u'_{op}(\omega) - \overline{u'_{op}})}{\sqrt{\sum_{\omega \in \Omega} (u_{op}(\omega) - \overline{u_{op}})^2 \sum_{\omega \in \Omega} (u'_{op}(\omega) - \overline{u'_{op}})^2}}, \quad (2)$$

where $\overline{u_{op}}$ and $\overline{u'_{op}}$ denote the real and estimated average utility over all outcomes. Alternatively, the *ranking distance of bids* compares all pairwise preference orderings:

$$d_r(u_{op}, u'_{op}) = \frac{1}{|\Omega|^2} \sum_{\omega \in \Omega, \omega' \in \Omega} c_{<u, <u'}(\omega, \omega'), \quad (3)$$

where $c_{<u, <u'}$ is the conflict indicator function, which is equal to one when the ranking of the outcomes ω and ω' differs between the two profiles, and zero otherwise.

Issue weight accuracy measures quantify the difference between the issue weights of u_{op} and u'_{op} . The underlying idea is that these variables are most important to estimate correctly. The metrics used are identical to the metrics above.

Pareto frontier accuracy measures focus only on the similarity between two Pareto frontiers. This is more challenging since their sizes can be different, so two sets of Pareto frontiers cannot be compared in the same way as outcome spaces or weight vectors. The *average difference of Pareto frontier* metric calculates the difference in utility over all Pareto bids of bid space B . The *percentage of found Pareto bids* measure gives the percentage of Pareto optimal bids of space B that are also in B' . Conversely, the *percentage of correct Pareto bids* metric yields the percentage of Pareto optimal bids in B' that are correct (i.e.: in B).

Finally, we introduce the *difference in Pareto frontier surface* measure, which is defined as follows: we take all outcomes in Ω that form the estimated Pareto frontier in B' ; we then map these points onto B . Finally, we compute the absolute difference in surface below these points and the actual Pareto frontier, as shown in Figure 1.




Figure 1. Visualization of the difference in Pareto frontier surface.

C. Quantifying the Estimation Accuracy

The goal of our first experiment is to quantify the accuracy of opponent models, both in different domains and through time. We outline in detail below the factors of the experimental setup that we believe are important to consider.

Influence of the Opponent Model on the Opponent's Actions

When an agent uses an opponent model, it indirectly influences the opponent in two ways: first, a model may *influence the time of agreement*, as a more accurate model may lead to better offers, resulting in earlier agreement; second, a model may cause changes in the *opponent's strategy*.

Both factors influence the bids presented by the opponent, and thus the information available for the model. To ensure every model learns from the same information, and therefore can be compared with the others, we selected non-adaptive opponents that never accept a bid.

Influence of the Opponent's Strategy on the Opponent Model

Opponents differ in how well their behavior corresponds to a model's assumptions. For instance, a model that assumes

that the opponent concedes will likely have problems modeling a very competitive agent. Therefore, we should select a balanced set of opponents to avoid favoring any model. One of the defining factors here is how much information an opponent reveals over time. For example, a conceding opponent reveals more of its preferences than an agent who makes random bids. Furthermore, we should include agents that strongly violate the modeling assumptions as to evaluate the robustness of the models. Taking both factors into account, we selected the following agents:

- 1) *Conceding agents* select a bid depending on the current time $t \in [0, 1]$ according to a target utility of the form $u_t = P_{max} \cdot (1 - t^{1/e})$ [15]. We selected four agents with $P_{max} = 1$ and different concession rates $e \in \{0.1, 0.2, 1.0, 2.0\}$. These agents make up the *predictable* opponents.
- 2) *Random agents* offer a random bid above a target utility m , where we selected $m \in \{0, 0.25, 0.50, 0.75\}$. This type of agent and the others below form the *unpredictable* opponents.
- 3) *Conceding agents with an offset* are time-dependent agents [15] that do not start with their best bid. For this category we use a linear concession rate ($e = 1$) and starting point $P_{max} \in \{0.7, 0.8, 0.9\}$.
- 4) *Non-conceding agents* start with a minimum target utility that increases to the maximum over time. The target utility is calculated as follows: $u_t = P_{min} + (1 - P_{min}) \cdot t$. We use four agents with parameters $P_{min} \in \{0, 0.25, 0.50, 0.75\}$.

Influence of the Scenario on the Opponent Model

We distinguish three features of the negotiation scenario that can significantly influence how well the opponent model is able to estimate the opponent's preferences:

- 1) *Domain size*. The total possible bids directly relates to the amount of parameters of the preference profile.
- 2) *Bid distribution*. The bid distribution is defined as the average distance to the nearest Pareto optimal bid. A scenario with a high bid distribution has a high percentage of outcomes far from the Pareto frontier.
- 3) *Opposition*. The opposition is defined as the distance from the Kalai-Smorodinsky point to the point of perfect satisfaction (maximum utility for both parties). The higher the opposition is, the more competitive the domain.

We made sure to select a balanced set of scenarios that display all characteristics. We chose five domains based on their size: *Irex vs Cypress* [12] (small: 180 bids), *Employer contract* [16] (small: 3125 bids), *ADG* [12] (medium: 15625 bids), *Supermarket* [12] (large: 98784 bids), and *Travel* [13] (large: 188160 bids). For each domain we created a set of scenarios varying in bid distribution and opposition. As we defined three levels of degree for both factors, 45 scenarios are used in total.

For the experiment, we ensured that each model processes exactly the same opponent traces, using a maximum amount of $N = 5000$ rounds. For the three groups of deterministic agents we recorded their (unique) negotiation trace, amounting to a total of 495 unique traces. For the random agents we recorded five different traces per agent, thus 900 traces in total. Combined, this amounts to 1395 traces that were used to train every opponent model.

D. Quantifying the Accuracy/Performance Relationship

The goal of the second experiment is to investigate the relation between accuracy and performance measures, thereby answering our final two research questions. For this experiment, we used a realistic set of opponents whose acceptance strategies are enabled. With realistic opponents, every negotiation is unique, so for this investigation we had to scale down the experimental setup. We selected a set of bidding strategies and scenarios where using a good opponent model would have added value; i.e., tough bidding strategies with limited learning capabilities (i.e., no opponent model), and large, competitive negotiation scenarios. We selected four of the top bidding strategies from ANAC: *Agent K2*, *HardHeaded*, *IAMhaggler2011*, and *The Negotiator*; and four time-dependent agents with concession rate $e \in \{0.1, 0.2, 1.0, 2.0\}$. These eight bidding strategies were combined with all thirteen models (the models in Table I and the two Bayesian models with perfect strategy knowledge) and no model. Each agent competed five times against all opponents (the eight bidding strategies without model) on five scenarios: *Grocery* [12], *Employment contract* [16], *Travel* [13], *Small Energy*, and *Supermarket* (both used in ANAC 2012).

Each agent played both sides of the five scenarios using a round-based protocol of 1000 rounds. Since 112 agents competed 5 times against 8 opponents on 5 scenarios for both preference profiles, 44800 matches were ran in total.

IV. EXPERIMENTAL ANALYSIS

We will now answer our three research questions by analyzing the results of both experiments. Each section corresponds to one of the research questions.

A. Evaluating the Estimation Accuracy of Opponent Models

As outlined in Section III-C, we measured the accuracy of a large set of opponent models to answer our first research question, the results of which are shown in Figure 2–4.

Accuracy over Time

The graphs in Figure 2 show the average accuracy of the opponent models over time. The left graph shows the accuracy over time when playing against predictable opponents; the right graph shows the results against the unpredictable opponents.

First of all, it is surprising to see that many of the state of the art models actually become *less* accurate over time. The

main cause of this phenomenon is that the bids presented later on in the negotiation are incorrectly handled. The value models and frequency models for example, treat every received bid the same way, independent of the time it is received. In effect, this means that when the opponent is conceding, the models increase the estimated utility of less preferred outcomes. This does not hold for the *CUHKAgent Value Model*, which incidentally also performs best, as this model only takes the first 100 unique bids into account. For Bayesian models, the problem is that they assume very particular opponent behavior, which is likely to become increasingly invalid as time progresses, and they perform very poorly as a result. When we disregard this shortcoming by considering the *perfect* Bayesian models, they perform better, and their accuracy then increases monotonically over time. However, even in this case, they come second to the *CUHKAgent Value Model* by a large margin.

Another interesting result is that despite their simplicity, the frequency models and value models perform best against both types of opponents. We believe that this is due to the small number of assumptions they make; i.e., only assuming that values with high utility are offered relatively more often. As the right graph illustrates, these models are rather robust, even though it is clear from the final accuracy that it is harder to model unpredictable agents.

The lesson to take away from this is that to be robust, opponent models need to minimize their assumptions about the opponent’s behavior. Of course, every model needs to make certain educated guesses, but when it does, the model should at least be highly *adaptable*, paying close attention to the opponent’s strategy. The predictions should be revised if, over time, the opponent behavior does not seem to fit the assumptions anymore.

Accuracy per Opponent

We now analyze in more detail the accuracy of the best performing models in every category against different opponents; that is, the best value model (*CUHK Value Model*), the best frequency model (*Smith Frequency Model*), and the best performing Bayesian models (*Perfect Scalable Bayesian Model* and *IAMhaggler Bayesian Model*). The results are shown in Figure 3. While the best value model performs best on average, there is no opponent model that dominates all others.

An interesting result is that the technique of the best value model to only take a limited amount of bids into account does not always pay off. The model performs poorly against the non-conceding agents, who show their most preferred values later in the negotiation. This means that the model can be fooled, which can be a concern in practice.

Accuracy per Scenario

We are also interested in exactly how the specifics of the negotiation scenario influences accuracy. We focus on the




Figure 2. Pearson correlation over time against predictable (conceding) opponents and the unpredictable opponents (other agents). The numbers above a cluster of lines are ordered from high to low accuracy.




Figure 3. Accuracy of four opponent models against different types of opponents, measured using two accuracy measures.

same four opponent models as above, and evaluate their accuracy against predictable opponents. Figure 4 summarizes the results. Note that we consider the average accuracy over all four models here, but we have verified that our conclusions also hold for each model individually.

One of the first observations is that there is a high variance in accuracy over different scenarios, and each factor seems to be equally important to consider. This underlines the importance of using a balanced set of negotiation scenarios. Clearly, *domain size* is a significant factor, as the domain size relates directly to the amount of unknown variables to be learnt. But also for the *bid distribution* and *opposition* we find a strong correlation with learning accuracy. The reasons for both are very similar: when the bid distribution or opposition is low, there are many outcomes of similar utility because the average distance to the Pareto frontier is small. This in turn, entails that the values of an issue are relatively close to each other in utility, which is harder for the models to learn than more extreme preferences.

B. Evaluating the Accuracy vs. Performance Relationship

Our second goal was to investigate the relationship between accuracy and performance of opponent models. Figure 5 visualizes the results for two accuracy measures: *Pearson correlation of bids* and *difference in Pareto frontier surface*.




Figure 4. Accuracy of the best opponent models on varying scenarios.

The performance is expressed in terms of obtained utility by the agents that employ the opponent models, normalized such that the *Worst model's* performance is zero, and the *Perfect model's* performance is 1. Using *no model* falls somewhere in between, since this is still better than using a wrong model.

The first thing to notice is the cluster of the best performers: the value and frequency models. The performance of these models is already quite close to that of the perfect model. To put it differently, we cannot anticipate a significant improvement from any other preference modeling technique over what is already achieved by these rather simple techniques.

The other types of opponent models also form clusters in the diagram. The *(Perfect) Bayesian Model* perform even worse than not using an opponent model; and only slightly better than simply assuming opposite preferences.

The almost linear relationship between accuracy and




Figure 5. Accuracy versus performance for all opponent models. Accuracy is measured using the *difference in Pareto frontier surface* (range $[0, 1]$, where 0 is best) and *Pearson correlation of bids* (range $[-1, 1]$, where 1 is best).

performance is the second thing that stands out. This shows that there is always added value to increasing the accuracy of an opponent model, even when the accuracy is already high. Nevertheless, the added value will necessarily be small, as the performance is already at 90% of its upper limit when the accuracy is at 70%.

An interesting comparison can be made with the results of the previous experiment. Figure 2 clearly shows the decrease in accuracy over time of many of the frequency and value models. How is it that they still manage to perform close to optimally? The reason is that many negotiations end in agreement, and this occurs somewhere *before* the deadline by definition. In these cases, the models are updated with less bids of poor value for the opponent. Therefore, the deciding factor in the success of the value and frequency models lies in their higher initial accuracy.

Finally, it is interesting that the results for the *Pearson correlation of bids* and *difference in Pareto frontier surface* metrics are in fact very similar when we ignore their orientation. Despite that the latter only measures the quality of the Pareto optimal frontier instead of the full outcome space, it seems to be a suitable predictor for performance as well. We explore this idea further in the next paragraph.

C. Evaluating the Usefulness of Accuracy Measures

Our final goal was to find a strong predictor for performance of opponent models, since there are so many different accuracy measures to pick from. Towards this end, we applied all of the accuracy measures shown in Table II and analyzed their correlation with performance; see Figure 6.

The dark line represents the predictive power of each accuracy measure, which is defined by the absolute correlation coefficient $|\rho|$ between the accuracy measure score and the model’s performance. We take the absolute value because some accuracy measures are negatively correlated with performance, while others are positively correlated.

The light grey line indicates what portion of the bid space is learnt by each accuracy measure. For this, we calculate the

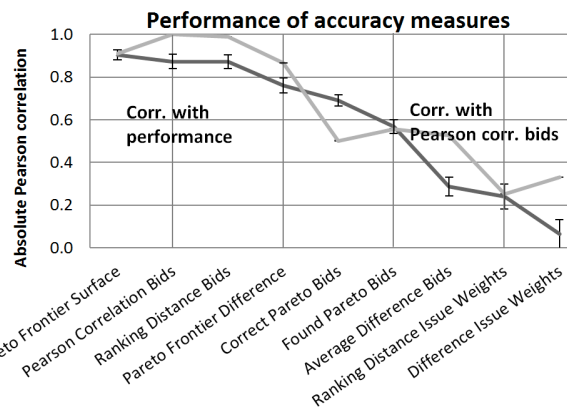


Figure 6. Absolute correlation between accuracy measure scores and two other measures: performance and “Pearson correlation of bids”.

absolute correlation coefficient $|\rho|$ between each accuracy measure and the *Pearson correlation of bids*. At the lower end of the scale we see the accuracy measures that only consider issue weights, which means they are not correlated at all with learning the space as a whole. These measures should not be used to make predictions about performance because they do not convey enough information about the accuracy of a model.

We found three measures that correlate strongly with performance, as indicated by the dark grey line, and therefore are good performance predictors; these are: *difference in Pareto frontier surface*, *Pearson correlation of bids*, and *Ranking distance of bids*. These measures codify sufficient information about the relationship between the real preferences and the learned preferences, and therefore, we can translate these notions to statements about performance. The performance of the top three measures are significantly better than the other measures (*one-tailed t-test*, $p < 0.01$).

Even though it only quantifies the similarity of the Pareto frontier, the *difference in Pareto frontier surface* metric performs best of all (*one-tailed t-test*, $p < 0.02$). This means that for an opponent model, it is sufficient to predict which

bids are Pareto optimal. The reason being that the Pareto frontier is a crucial component of the outcome space, and that many bidding strategies seek Pareto optimal agreements. It is less computationally expensive to calculate than the *ranking distance of bids*, and it has the lowest standard deviation between the runs. Furthermore, it is defined for all inputs, in contrast to the Pearson correlation measures, whose results are undefined when all bids are estimated to have the same utility. Therefore, we recommend the *difference in Pareto frontier surface* as a suitable measure for accuracy.

V. RELATED WORK

We consider two categories of related work. The first category consists of work *evaluating the accuracy* of opponent models. Carbonneau et al. [5] calculate the Pearson correlation between the real and estimated utility of the opponent's next bid. Hindriks and Tykhonov [6] extend this approach by measuring the Pearson correlation of the whole outcome space and discuss analogous definitions for the ranking distance. Our method incorporates both measures.

An alternative approach is to measure the distance between elements of two preference profiles. For example Jazayeriy et al. [7] introduce such measures for the learning error of issue weights. We have incorporated these measures in our method, and we also apply the same measures to quantify the similarity between two full bid spaces.

Finally, there exist accuracy measures tailored to specific learning methods. Buffett and Spencer [4] for example, define a metric for opponent models that use Bayesian learning. The measure is defined as the average likelihood that the correct hypothesis is chosen from the set of candidate hypotheses. Since we employ models that are based on a wide range of learning techniques, we do not incorporate measures specific to a particular learning method.

The second category of related work quantifies the *relationship between accuracy and performance*. In [1], Coehoorn and Jennings introduce a model that estimates the opponent's issue weights and investigate the influence of small prediction errors on performance. Our method takes this a step further, as we analyze the relation between an exhaustive set of accuracy measures – including accuracy of the issue weights – and performance.

VI. CONCLUSION AND FUTURE WORK

In this work we have evaluated a large set of accuracy measures to identify the best method to predict the performance of opponent modeling techniques in negotiation. We introduced a procedure to quantify the accuracy of state of the art opponent models and we identified their strengths and weaknesses. One of our main conclusions is that there is an almost linear correspondence between accuracy and performance of models when we employ the proper accuracy measures. Moreover, the best models are close to being

perfectly accurate, which means there is only limited room for improvement with regard to performance.

Surprisingly, the accuracy of most opponent models decays over time due to the incorrect handling of the opponent's less preferred bids, which are usually offered at a later stage of the negotiation. Especially then, good strategy prediction is needed in order to be effective at preference modeling.

Finally, we analyzed how well accuracy measures can predict the performance of an opponent model. We concluded that three measures in particular are useful predictors of performance, and we found that this can be best achieved by limiting the analysis to difference in Pareto frontier surface between the real and the learned bid space.

REFERENCES

- [1] R. M. Coehoorn and N. R. Jennings, "Learning an opponent's preferences to make effective multi-issue negotiation trade-offs," in *Proceedings of the 6th international conference on Electronic commerce*, 2004.
- [2] D. Zeng and K. Sycara, "Bayesian learning in negotiation," *International Journal of Human Computer Systems*, vol. 48, pp. 125–141, 1998.
- [3] K. V. Hindriks and D. Tykhonov, "Opponent modelling in automated multi-issue negotiation using bayesian learning," in *Proceedings of the 7th international joint conference on Autonomous agents and multiagent systems - Volume 1*, 2008.
- [4] S. Buffett and B. Spencer, "Learning opponents' preferences in multi-object automated negotiation," in *Proceedings of the 7th international conference on Electronic commerce*, ser. ICEC '05. New York, NY, USA: ACM, 2005, pp. 300–305. [Online]. Available: <http://doi.acm.org/10.1145/1089551.1089608>
- [5] R. Carbonneau, G. Kersten, and R. Vahidov, "Predicting opponent's moves in electronic negotiations using neural networks," *Expert Systems with Applications*, vol. 34, no. 2, pp. 1266–1273, 2008.
- [6] K. V. Hindriks and D. Tykhonov, "Towards a quality assessment method for learning preference profiles in negotiation," in *Agent-Mediated Electronic Commerce and Trading Agent Design and Analysis*. Springer Berlin Heidelberg, 2010.
- [7] H. Jazayeriy, M. Azmi-Murad, N. Sulaiman, and N. Izura Udizir, "The learning of an opponent's approximate preferences in bilateral automated negotiation," *Journal of theoretical and applied electronic commerce research*, vol. 6, no. 3, pp. 65–84, 12 2011.
- [8] C. R. Williams, V. Robu, E. H. Gerding, and N. R. Jennings, "Iamhaggler: A negotiation agent for complex environments," in *New Trends in Agent-based Complex Automated Negotiations*. Springer-Verlag, 2012, pp. 151–158.
- [9] T. van Krimpen, D. Looije, and S. Hajizadeh, "Hardheaded," in *Complex Automated Negotiations: Theories, Models, and Software Competitions*. Springer Berlin Heidelberg, 2013.
- [10] N. van Galen Last, "Agent smith: Opponent model estimation in bilateral multi-issue negotiation," in *New Trends in Agent-based Complex Automated Negotiations*. Berlin, Heidelberg: Springer-Verlag, 2012, pp. 167–174.
- [11] T. Ito, M. Klein, and H. Hattori, "A multi-issue negotiation protocol among agents with nonlinear utility functions," *Multiagent Grid Syst.*, vol. 4, no. 1, pp. 67–83, Jan. 2008.
- [12] T. Baarslag, K. Fujita, E. H. Gerding, K. Hindriks, T. Ito, N. R. Jennings, C. Jonker, S. Kraus, R. Lin, V. Robu, and C. R. Williams, "Evaluating practical negotiating agents: Results and analysis of the 2011 international competition," *Artificial Intelligence Journal*, 2012.
- [13] T. Baarslag, K. Hindriks, C. M. Jonker, S. Kraus, and R. Lin, "The first automated negotiating agents competition (ANAC 2010)," in *New Trends in Agent-based Complex Automated Negotiations*. Springer-Verlag, 2012, pp. 113–135.
- [14] T. Baarslag, K. Hindriks, M. Hendriks, A. Dirkzwager, and C. 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.
- [15] P. Faratin, C. Sierra, and N. R. Jennings, "Negotiation decision functions for autonomous agents," *Robotics and Autonomous Systems*, vol. 24, no. 3–4, pp. 159 – 182, 1998, multi-Agent Rationality.
- [16] L. Thompson, *The Mind and heart of the negotiator*, 3rd ed. Upper Saddle River, NJ, USA: Prentice Hall Press, 2000.