BOFA - A Framework for Fairness in Automated Negotiations

Nikolay Blagoev
Supervisor: Sietze Kuilman
Responsible Professor: Luciano Cavalcante Siebert
EEMCS, Delft University of Technology, The Netherlands

June 19, 2022

A Dissertation Submitted to EEMCS faculty Delft University of Technology,
In Partial Fulfilment of the Requirements
For the Bachelor of Computer Science and Engineering

Abstract

As automated negotiating agents become more and more part of our daily life, additional care needs to be taken that the agents can negotiate fairly. Humans each have their own intrinsic view on fairness, which affects the negotiation processes and the degree to which the outcome is viewed as satisfactory. However, most current agents are built around a specific notion of fairness, which could be potentially unwanted in certain scenarios or by some individuals. This paper presents a generic framework through which varying fairness notions can be implemented. The performance of the framework is tested through a selection of proposed agents with different fairness views. The results show that the agents are able to reach an agreement on the pareto frontier that is fair given their notion. Thus they do not sacrifice optimality for fairness.

1 Introduction

Automated negotiation is an iterative search process between two or more parties (be it human, computer, or both), carried out as an exchange of propositions (bids) over some protocol [10]. The negotiation[1] has elements of both cooperation and competition and can be regarded as a repeated general sum game. Application of automated negotiation can be within commerce, but also in the fields of computer networks, energy markets, robot synchronisation [10]. In fact, negotiation can be used for any interaction between at least two parties where some resources have to be distributed or an agreement for future actions has to be reached. A growing body of literature is focusing on achieving fairness within a negotiation, however little to no analysis is put into the applied notion of fairness. Most proposed strategies focus on an egalitarian outcome (for example - Nash Point) and, implicitly, touch upon fairness in the process through tactics such as Tit-for-Tat[13][10].

Within the field of computer science, "fairness" is used as a substitute for "lack of bias" (see the work of Dworak[3]), and thus, by extension, it is equated to equality[2]. As mentioned, most proposed "fair" agents assume as a goal the Nash-point (an egalitarian solution), without any reasoning on the side of the authors as to why this choice was made. In other fields, mainly Economics and Philosophy, other notions of fairness exist, such as equity, in the form of a proportional distribution of the outcome, or reciprocity, in the form of changes in the degree of egoism/altruism in response to another's behaviour[16].

If the end goal of the field of automated negotiation is for agents to represent people[2], then the concept of fairness in all of its facets ought to be considered. Humans naturally have a notion of fairness by which they act and this can explain why their behaviour deviates from the perfectly rational model[1][4]. Additionally, fairness greatly influences whether a particular agreement will be viewed as satisfactory and honoured[6].

The purpose of this paper is to propose a framework through which fairness can be introduced in the field of automated negotiation. To avoid defining fairness in a concrete manner, and thus tying the strategy to specific (and potentially unwanted) fair behaviour - the proposed solution is designed to be as agnostic to a notion of fairness as possible. Instead, the general desired properties of the fairness metrics are outlined and their concrete implementation is left to the user. Three algorithms for incorporating said behaviour are provided as examples in Appendix A - proportional distribution, Nash-equilibrium, and reciprocity to intentions. These can serve as loose guidelines when one wishes to implement a specific fair behaviour within the framework. Chapter 4 shows a concrete negotiation

[1] A negotiation is not a collection of bids, but a relationship between agents, mediated through the bids

[2] A notion of Liberte, Egalite, Automatisation

strategy of an agent built on the proposed framework. Chapter 5 evaluates the performance of the strategy in an agent society.

This paper will constrain its analysis to bilateral negotiations, carried through the Stacked Alternating Offers Protocol[15], where the domains are discrete and the profiles can be expressed through a weighted sum of the different issues. This restriction is imposed merely to keep the paper to a reasonable length, however the proposed ideas can easily be generalised to other domains, as they do not hinge on any of these factors.

Chapter 5 evaluates the performance of the proposed strategy, by answering the research following questions through an experiment:

- Can the proposed strategy achieve high and fair utility for itself in a heterogenous environment (one with many other potentially unknown strategies)?
- Can the proposed strategy perform optimally (i.e. does it violate the Pareto Principle) with different fairness notions?

1.1 Analysing Fairness in Automated Negotiation

Remaining agnostic to the notion of fairness used during the negotiations does not mean fairness cannot be analysed. Specifically, we can investigate where in the negotiation and in what forms can fairness manifest itself. To this end, the work of Cecilia Albin in [7] proves instrumental. It identifies four categories of fairness in negotiations:

- Structural fairness concerns the structure of the negotiation process - the protocol used, any power imbalance between parties, etc[7]. The strategy outlined in this paper assumes that structural fairness has already been achieved, as agents can engage freely with other parties over preferred protocols.
- Process fairness concerns the degree to which the parties use deceptive means to accomplish their goals (for example stating inflated payoffs), to what degree are the agreed upon rules followed by the parties, and how parties' notion of a fair outcome affects the negotiation process[7]. The proposed strategy addresses process fairness in part through the dynamic fairness function (see Chapter 3.3) and the protocol outlined for a fair exchange of preference profiles (see Chapter 2).
- Procedural fairness concerns itself with the techniques employed to arrive at the final result[7]. Although similar to process fairness, it differs by focusing on the techniques used rather than how they are used. Procedural fairness is addressed by the kindness function (see Chapter 3.3) and the (optional) use of the reciprocal concession strategy (see Chapter 4.2).
- Outcome fairness relates to the view of what is a fair distribution of the resources (the outcome)[3][7]. It is what is most commonly associated with fairness in negotiation. Outcome fairness is addressed in the proposed strategy through the static fairness function (see Chapter 3.2)

Fairness in automated negotiations suffers from two major issues. First, agents may have different views on what a fair distribution would constitute. For example, an agent

[3]The actual definition by C. Albin also includes the degree to which the parties view the distribution as fair, manifest in fulfilling this agreement, however I do not think this part to be relevant to automated negotiation

that has contributed twice as much may have the reasonable (and fair) belief that they should receive twice as much as the other party [7]. Agents may be willing to choose a distribution of resources that yields a lower utility for itself (relative to some other point), however such distribution is fairer in its view (for example attempting to approach the Kalai-Sordinsky point[22], instead of proposing its best bid). Thus now the problem becomes twofold. The agents have to agree not only on a strictly material distribution, but also a shared concept of fairness. Second, how can fairness be verified? Since the two agents have no knowledge of each otherâs utility functions, how can they verify that fairness (in outcome, in process) has been achieved [4]. While no clear solution exists for the first issue, as it depends on the agent's strategies and outcome fairness views, the second one is addressed in Chapter 2, in the context of bilateral negotiations.

2 A Fair Protocol to Exchange Profiles

A great many of fair behaviours hinge heavily on an accurate idea of the opponent's motivations and values (a similar conclusion is reached in [4]), which in the context of automated negotiation can be understood in part through their preference profile. Examples would include a proportional distribution[19], Nash-equilibrium distribution[21], reciprocal concessions towards some fair outcome[20]. In the ideal case the two parties can simply exchange their profiles when negotiations begin. However, one party can gain an unfair advantage by holding off on sending their profile (u1) until they have received their opponent's (u2). Then they calculate a preference profile, u1', in such a way that some fair distribution between u1' and u2 is actually unfair under u1 and u2 (which is similar to what Cecilia Albin discusses in regards to process fairness [7]). Although this can be resolved by adding a third party (a mediator) this introduces additional overhead in a bilateral negotiation. A simpler solution is outlined here, in the form a protocol through which agents can exchange profiles, without one gaining an unfair advantage:

- (OPTIONALLY) Each party gives a deposit which is returned at the end of the negotiation process. This prevents premature terminations of the negotiation which can be viewed as fair. This is similar to the work in [9].
- The two parties send each other a commitment of their profile (for example, their profile encrypted under some key)[8].
- Only once they have received the profile of the other, do they reveal the profile (sending the decryption key). Thus, one party cannot make an unfair change to their profile after learning the other one's.

The protocol, however, modifies the SAOP protocol[15] or rather - adds on top of it. Some legacy (or unwilling) parties may not support this extension, thus the proposed strategy does not require that such an exchange has happened (but it can greatly benefit from it). The analysis of the performance of the strategy in Chapter 5 is in two parts. In the first part the agents do not have each other's profiles, thus they must use Opponent Modelling to estimate it. This however introduces a dependency of the strategy on a good opponent model and any malperformances could be the fault of an incorrect estimate. In the second one it is assumed that such an exchange has occurred (thus the agents have each others' profiles), essentially

[4]Without an omniscient mediator everything is permitted

transforming it into a game of perfect information. This is done to focus the evaluation on the framework itself through the specific strategy it is employed within, without additional variables like Opponent Modelling.

3 Introducing BOFA

The proposed framework simply adds an additional component to the BOA framework[12] - mainly Fairness (which is what the F stands for). This way this work can easily be adapted to suit different needs by pairing it with any Bidding, Opponent Modelling, and Acceptance strategy[12].

3.1 The F in BOFA

Instead of strictly focusing on material utility, under BOFA an agent evaluates the value function of each bid, which is:

$$V = c1 * utility + c2 * Fs + c3 * Fd \tag{1}$$

where:

- c1, c2, c3 are simply constants
- Utility is the material utility of the bid
- Fs is the static view of fairness (or the outcome fairness as viewed before the bidding process has started)
- Fd is the dynamic view of fairness (or an outcome fairness dependent on the proceedings of the negotiations)

Varying the three constants results in different behaviour. For example, increasing c1 relative to c2+c3 will result in a more egoistic agent, which prioritises their material utility over some notion of fairness.

3.2 Static Fairness

The Static Fairness function (Fs) accounts for the desired outcome, which is taken as a reference point for fairness. This is where the Nash-point can be used in an attempt to move towards that specific distribution. It is important to note that this function is bound between [-1, 0], as in it measures unfairness rather than fairness. This is because in my opinion once a baseline view is established of fairness, one cannot be more fair than it. For example, let us assume we want to distribute a certain amount of money and I view as fair a 2:1 split (in my favour). Now, if you propose 1:2 in your favour, I will see that as unfair. However, even if you propose a 4:1 I still should see it as unfair, as we have deviated again from the fairness reference point. You might be able to treat me more KINDLY, but not more FAIRLY. Hence, -1 will become the most unfair possible outcome and 0 - the ideal one[5].

A few important properties worth mentioning. First and foremost, the function should not lead to loss of optimality (violation of the Pareto Principle - an important point raised

[5]A similar view is held in [18] regarding inequity aversion

in [11]). This can be avoided by ensuring that, given a perfectly continuous Pareto frontier, at least one point on it has a static fairness evaluation of 0. Secondly, as mentioned before, it should be bound between [-1, 0], with 0 being the ideal fairness case.

3.3 Dynamic Fairness

The Dynamic Fairness function (Fd) accounts for the process and procedural fairness. It essentially modifies the preference for certain bids based on how the opponent has acted throughout the negotiation. One can think of this function as an equivalent to "intention based reciprocity" - treating the other more unkindly when we believe they are mistreating us[6][17]. It has the same properties as the Static Fairness function, with an additional key requirement. As the name would suggests, it can change throughout the negotiation, while Static Fairness is constant. This means that some bid (b1) can be preferred over another (b2) in a round, while in the next one - this relationship can be inverted (b2 preferred over b1). This is achieved through a combination of a kindness and a compensation function.

The kindness function, unlike the fairness ones, is bound on the interval [-1,1]. That is because one can treat another more kindly than some reference point. An example would be two acquaintances seeing each other in passing. The reference point (which would take the value 0 - neutrality) would be a casual pleasantry. If one of the two decided to do a more warm greeting and compliment the other, the kindness evaluation would become positive (tending towards 1 as the kindness becomes more unexpected). In contrast, if instead that person decided to greet the other with a violent onslaught of punches, the latter's kindness evaluation should start tending to -1. Thus the interpretation of the range of the function should become evident. The compensation function accounts for past miscalculations of the kindness function, due to an inaccurate opponent model. That way we can compensate for a mistake in our attitude towards an agent. These two functions, however, do not directly constitute the dynamic fairness, but merely cause it to change (more in favour of us or our opponent) in some manner.

The presence of a kindness function allows the agent to exhibit more complex concessional behaviour. An example would be "intention based" concession. This can be achieved by having a random chance to ignore an opponent's concession if the kindness evaluation tends to -1. Thus the agent acts in a more egoistic manner when it believes its opponent to be treating it unkindly. Additionally, the kindness value can be remembered by agents and thus influence future negotiations between the same parties. Agents can also share their evaluation of other parties between each other, thus the kindness function becoming similar to a trust score in peer-to-peer networks.

4 Proposed strategy for a fair agent

Thus far I have explained the underlying idea of the BOFA framework. This chapter presents an agent which implements the outlined framework. THEMIS[7] is essentially a slight modification of the agent in [10], using their Opponent Modeling strategy verbatim. It also uses their idea of considering rounds as windows of fixed lengths of consecutive offers[10]. The agent's strategy can be presented through the following Finite State Machine:

[6]Of course this relationship can be inverted to suit the user's need. Treat those kind to us unkindly, and vice versa.

[7]Named after the Greek goddess of justice

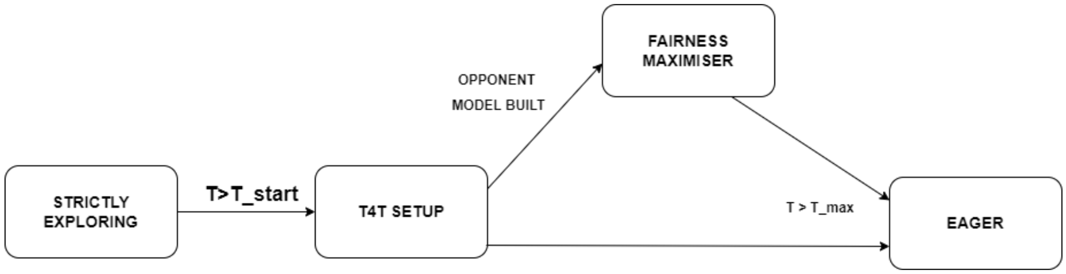

Figure 1: FSM of the agent's strategy

4.1 Strictly Exploring

During this initial phase the agent is exploring the domain and signaling its position to its opponent. This is done by giving its best bids above a certain initial target_min (lowest utility of acceptable bids) in descending order. When each bid is offered there is a p chance for it to be repeated and $1-p$ chance to proceed to the next one, where p is just a constant. While in this state the agent should be reluctant to accept incoming offers (as they could be potentially suboptimal). Thus the acceptance strategy is that an incoming bid is accepted if c1>c2+c3 and its utility is above target_min.

4.2 Tit-for-Tat setup

Once the first two consecutive bid windows are filled (so $2 * window\_size$ bids have been offered by the opponent), the agent moves to the Tit-for-Tat setup state. In this strategy the opponent's behaviour is analysed, hence it is needed that at least 2 windows of bids are present so as to compare them for any significant changes[10]. This state is based on the concession strategy described in [10], with two small differences. First, and where the name setup comes from, is that if a proportional outcome is chosen (in the form of $us/them$), the agent will concede target_min (which is called u_target in [10]) to $0.9 * us/them$. The reasoning is that if a fair outcome is 1/2, it makes no sense to propose bids with a utility above 0.5 for ourselves (as the utility is bound on the interval [0,1]). Second, it is a proportional concessional strategy[16]. That is if in [10] they would lower u_target by some value X, this agent would lower target_min by X/(us/them). For example, under 2/1 proportional distribution a concession of 0.05 would become 0.025.

In this state bids are offered randomly, with the requirement that their utility evaluation is above the target_min and below target_min+0.3 of the current round. Additionally, when a bid is offered there is a 0.9 chance that its utility evaluation is less than the previous one. Lastly, to prevent repetition of bids, once offered, each bid has a p chance to be repeated. The goal of this strategy is to prepare for the next state, learn about the opponent's preferences, and signal our position to them. The acceptance strategy is the same as the one in Chapter 4.1.

4.3 Fairness Optimiser

When the opponent model has been updated 4 times (or 4 rounds of the T4T setup strategy have elapsed when the opponent's profile was already known), the agent moves to a Fairness Optimiser strategy. Here it offers its best bids above target_min based on the Value function described in Chapter 3.1. The functions used are described in Appendix A.1, Appendix A.3, and Appendix A.4. Target_min is continuously updated based on the strategy described in Chapter 4.2, however if the kindness evaluation is above or below certain thresholds there is a random chance to ignore a change and instead move in the opposite direction. Thus, the concession strategy contains also elements of responding to a trend[16]. The acceptance strategy here is to accept if the incoming offer has a utility evaluation above $target\_min$ and has a better value score than any of the next N bids we are about to propose.

4.4 Eager strategy

Once time passes a certain t_max threshold, the agent moves to an eager strategy. Here, if c1>c2+c3 (prioritises more personal utility than notions of fairness), the agent will repeat the best bids it has received, which are above a certain reservation value. The agent will accept any incoming offer above target_min. If no such bids exist, it will simply repeat the last bid it has sent. If c1<=c2+c3, the agent will simply repeat the last bid it has sent and accept anything with a value evaluation greater than it.

5 Experimental Setup and Results

This section assess empirically the performance of the framework through the proposed agents. The aim is to answer the research questions outlined in Chapter 1.

5.1 Metrics

The metrics chosen to evaluate the performance of the agents are as follows:

- Pareto Accuracy (ParAcc) - given as a percentage, it measures the amount of times the agent achieved an agreement on the Pareto Optimal Frontier over all its runs.
- Average Distance to the Pareto Frontier (DistToPar) - this measures the average of the minimal distances to the pareto frontier when the agreement was not on it.
- Average Distance to Kalai-Sordinsky Point (DistToKalai) - this measures the average distance to Kalai-Sordinsky point[22]. This solution attempts to get the ratio of the utilities of the agents as close to each other as possible (ratio of 1:1). In principle, the closer an agent is to this point, the fairer it is given an egalitarian outcome fairness view.
- Nash Product (Nash) - this measures the product of the two agent's utilities. It is a commonly used measurement to assess an agent's social performance[10][24].
- Social Welfare (SW) - this measurement is simply the sum of the two agent's utilities. It measures how well off are everyone in society. Thus the higher - the better.
- Average Agent's utility (Util) - this measures the agent's average utility across its runs. Here, however, it is not necessarily the case that higher result means better performance. For example, a highly fair agent might have a lower utility than a more egoistic one.

5.2 Agents

The following agents were chosen for the empirical analysis:

- Boulware agent - The boulware agent employs a time-based strategy, where it first keeps its utility close to its best and as time progresses - concedes in an exponential manner[16]. It was chosen as a fairly simple, but still well performing agent, which has no notion of fairness. Thus it operates in a purely egoistic manner.

- IAmHaggler2012[24] - this agent was highlighted as the most social agent in the 2012 ANAC competition[10]. It was chosen so that the proposed agents performance can be tested against other real-world social agents.

- Social Agent[8] - this agent is based on the one described in [10]. It can be seen as the parent of the proposed agents as their strategy draws heavily from its own. It was chosen to evaluate whether an improvement has been brought by the modifications.

- Themis Kalai - this agent uses the proposed strategy with a fairness view function as described in Appendix A.1. Its settings are in Appendix C.1. The goal of this agent is the approach an egalitarian split of the resources (the Kalai-Sordinsky point).

- Themis Nash - this agent uses the proposed strategy with a fairness view function as described in Appendix A.2. Its settings are detailed in Appendix C.1. This agent essentially tries to maximise the Nash product in its Fairness Optimisation state.

Two agents which use the strategy outlined in Chapter 4 are listed. That is so this analysis can answer the second of the two research questions - "Can the proposed strategy perform optimally with DIFFERENT fairness notions?" For this experiment, both Themis agents have been "blinded"[9] - they do not have access to the opponent's profile and need to rely on their opponent modelling strategy.

5.3 Domain

The domain over which the negotiations is ran is the England vs Zimbabwe scenario[23]. It was used during the 2011 ANAC competiton and consists of 5 issues with a total of 576 possible resolutions.

5.4 Experimental setup

The experiment consists of a tournament-like setup where each agent is ran against each other for a total of 20 times. After the first 10 runs, each pair of agents switches sides. Thus if an agent was with profile A against profile B, after the 10th run it will be with profile B against profile A. Each negotiation is given a deadline of 200 rounds. The platform chosen to facilitate this experiment was Genius, as it is flexible and used in official automated negotiation tournaments[23][24].

[8] I wanted to call it Uranus, the father of Themis, but I thought the mythology references would get too confusing

[9] Because justice is blind

5.5 Results

The results of the experiment are outlined in the following table:

Agent	ParAcc	DistToPar	DistToKalai	Nash	SW	Util	Disagreements
Boulware	0.6125	0.035	0.0434	0.645	1.608	0.808	0
IAmHaggler2012	0.7125	**0.0189**	0.062	0.649	1.614	0.767	0
Social Agent	0.725	0.129	0.068	0.635	1.575	0.796	0
Themis Nash	0.725	0.0253	0.037	**0.652**	**1.618**	**0.823**	0
Themis Kalai	**0.750**	0.0253	**0.034**	0.647	1.602	0.818	0

Table 1: Tournament Results

As one can see, the two Themis agents performed the best in almost all categories. The only criterion in which they lost to another agent is the average Distance to Pareto Frontier, where the IAmHaggler2012 agent outperformed them by 0.0064. Of interest is the Pareto Accuracy metric, where the proposed agents have the two highest scores. Furthermore, they also had the two highest individual utilities. Thus this corroborates that the proposed strategy is able to obtain fair and pareto optimal results, without sacrificing its own utility. In terms of fairness, each of the Themis agents was also able to perform best in their respective category. Themis Kalai had the lowet distance to the Kalai-Sordinsky point and Themis Nash had the highest Nash product. This proves the proposed agents are able to perform fairer than the other agents in the tournament, which included two social agents - IAmHaggler2012 and Social Agent. It is worth noting that while Themis Kalai did have a lower utility than Themis Nash, its goal was the Kalai-Sordinsky point (0.795, 0.835), which would result in lower utility on average than the Nash-Bargaining point (0.745, 0.91).

5.6 Evaluating more unorthodox fairness views

While the previous section analysed the performance of the Themis agents in a heterogenous society, it utilised egalitarian and commonly used fairness views. This section evaluates the performance of two strategies based on equity. The actors of interest in this chapter are Themis Kind (TK), which is Themis Kalai, but with proportion of 1:2, and Themis Rude (TR) - proportion of 2:1. The setup is similar to the one in the previous chapters, except the metric by which they will be assessed is the Ratio Between Utilities (RBU). This metric is simply the fraction of one agent's utility over the other. That is because, given a certain scenario, depending on which side an agent holds, the desired outcome position is not the same, the solution is not symmetric. In case of a disagreement the RBU evaluation of that round is 0. Also, since the two agents may have different views on a fair distributions, it may be reasonable to assume they will try to "split-the-difference" between their desired outcomes[7].

This point can be exemplified through a situation I have dubbed a "Canadian standpoint". Imagine two agents (Themis Kind), which both want twice as much for their opponent. If they are too hung up on this desired outcome fairness they wouldn't reach an agreement. Which is realistically quite an absurd situation - agreement is not met because the two sides are being too nice to each other (although it does happen sometimes in the real world). Thus, it would be reasonable to expect that the two agents will meet at a 50:50 split of the resources (Kalai-Sordinsky point). This is what the RBU metric evaluates - it will give the average proportion of the outcomes.

Additionally, the agents will be "unblinded" - they will have the opponent's profile before the negotiation begins (exchanged through the protocol outlined in Chapter 2). The results of running the two agents against each other and against the other 2 Themis agents in the previous chapter are outlined in the table below:

Agent1	Agent2	Agent1 Prop.	Agent2 prop.	RBU	Pareto accuracy	Disagreements
TK	TK	0.5	0.5	0.93	0.63	4
TK	TR	0.5	2	0.5	1	0
TK	Themis Kalai	0.5	1	0.9	0.74	1
Themis Kalai	TR	1	2	0.81	1	18
TK	Themis Nash	0.5	NA	0.67	0.83	1
TR	TR	2	2	0	0	20
TR	Themis Nash	2	NA	0	0	20

Table 2: Results of inter-THEMIS runs

Evidently, the success of the negotiation highly depends on the fairness views of the agents and their coefficient ratio (c1:c2:c3 in the Value function). Some combinations of fairness views cannot result in a "split the difference" resolution. For example, the Canadian standoff can be resolved by the agents simply moving towards their opponent's fairness view. However, two Themis Rude agents cannot split the difference so easily in this manner, as it would require them to become more altruistic. This paper largely assumes that individuals can only become egoistic without the presence of external stimuli (such as kindness arising from the conduct of the negotiation). When paired against an agent with a complementary fairness views (the reciprocal of their own), each agent was able to achieve RBU of the desired proportion. Other results of note are that of the two TK agents against each other and a TK agent against a Themis Kalai agent. The former were able to achieve almost the desired RBU of 1, with the slight inaccuracy stemming from the 4 disagreements, which would lower the average. The latter pair achieved an RBU of 0.9, as the settings of the Themis Kalai agent made it to prioritise the Static Fairness evaluation twice as much as the Dynamic Fairness one. Thus Themis Kalai would deviate from its desired outcome much harder.

6 Responsible Research

There are no relevant ethical issues that affect the project. The importance of fairness is extensively discussed in this work and the decision on a specific (and arbitrary) choice of notion of it is left to the end-user. The purpose of this paper is to present a way to incorporate fairness in negotiations rather than a specific implementation of it.

The source code of all the agents proposed by this paper and the domain used are uploaded to GitHub (https://github.com/NikolayBlagoev/FairAgent). To reproduce the results of Chapter 5, simply add the python agents to the GeniusWeb server and run a negotiation with the settings described in the chapter.

With respect to the reproducibility of the results, two important points ought to be raised. First, the agents were developed for the GeniusWeb platform, however most ANAC participants which supported the SAOP protocol were written for the Genius platform. As far as I am aware, the two platforms are not compatible with each other. This meant that the negotiations had to be carried out manually - whenever an agent offered a bid, I had to input it to the other one. This was a very slow and error-prone process. Second, the agents behave in a stochastic manner during their first two strategies, which can affect the outcome of the negotiation. Thus a large enough sample size of runs needs to be considered when analysing the performance.edf

7 Conclusions and Future Work

As automated negotiating agents become more prevalent in our society, more care needs to be put in creating fair negotiations given different fairness notions. Currently, most agents strive for an egalitarian approach for the outcome and the process fairness is touched upon through tactics such as Tit-for-Tat. This paper presents a generic framework that can represent different notions and metrics to evaluate the process and outcome fairness. The benefit of this is that it doesn't tie itself to a specific view of fairness, while also encompassing the different aspects of fairness within negotiation, and thus can readily be adapted to different notions and scenarios.

The performance of the fair agents was evaluated through competition with various agents. The experiment demonstrated that the agent is able to perform both optimally and fairly with different fairness notions. The agents were able to stay close to the desired outcome, while also maintaining high utility for themselves. The framework was also evaluated by having agents with different notions of fairness negotiate with each other. It was demonstrated how through a combination of the different functions a more fair outcome that is satisfactory to both sides can be achieved. Further testing can be performed by running the agents in a society with a lot more participants, so that the different parameters can be better tuned.

Currently the agent has very poor efficiency and tends to stagger a lot on large domains. On the other hand, the opponent modelling strategy tends to suffer on small domains. The proposed strategy ought to be revised to be able to perform within reasonable time in large domains. Furthermore, scenarios, such as the "Nice-or-die"[23] one, that are very small and require rapid concession should be considered as currently the proposed agents fail to reach agreements on them.

A proposed improvement is instead of windows of fixed length to use dynamically created windows of varying length. In most trials this lead to better estimation of the opponent's profile and clear identification of strategy changes. For further information see Appendix B.

A Fairness functions examples

A.1 Static Fairness Example - proportional distribution

This section demonstrates an example of how static fairness (or fairness view for any of the other components) can be calculated given any desired proportional distribution.

$$Fairness\_view = normalise(abs(atan(u1/u2) - atan(prop))) \tag{2}$$

Where $u1$ is our utility of the bid, $u2$ is our opponent's, $prop$ is the desired proportional distribution (expressed as a fraction), and $atan$ is the inverse of the tangent function. This fairness view maintains the desired properties as prescribed in Chapter 3.2. What it does, essentially, is compare the ratio of $u1/u2$ to $prop$. Bids close to the ideal fair outcome have $fairness\_view$ approaching 0. Also, given a continuous Pareto frontier, there will exist a point that has a fairness evaluation of 0 (mainly the intersection of the Pareto frontier with the line $y = x * prop$). To bind the static fairness between -1 and 0, the following formula is used:

$$Fs = -1 + exp(-(fairness\_view/std)^2) \tag{3}$$

Where std is a constant acting as the standard deviation. This also allows to punish bids in an exponential manner, allowing for values of $u1/u2$ close to the desired outcome to be considered with negligible decrease in the Value function. The idea behind this is to avoid suboptimal bids from being preferred, i.e. fairness being the cause for all participants to be worse off, due to a discrete pareto frontier.

A.2 Static Fairness Example - Nash Point

To use the Nash Point as the desired outcome fairness, one can use the solution outlined in the previous chapter, changing only the Fairness view function to:

$$Fairness\_view = 1 - (u1 * u2)/(u1 * u2)* \tag{4}$$

where $(u1*u2)*$ is the Nash Product. The $Fairness\_View$ can then be used directly for the static fairness function or it can be used within the exponential function from the previous chapter.

A.3 Kindness Function Example - Proportional Distribution

The proposed algorithm for calculating the kindness metric at each round is based on distance from desired utility and progress towards a fair outcome. First, an acceptable window of bids values for us is calculated through the formula:

$$window = target\_min - 1 + exp((1 - min(tmax, t)/tmax) * ln(1 - kstart)) \tag{5}$$

where t is the current negotiation time, $tmax$ is the maximum negotiation time during which this strategy is used, $target\_min$ is the minimum value of acceptable bids, $kstart$ is the initial acceptable window size. One can observe that $window$ decreases with time. Thus, at some round with $target\_min$ of 0.7, $window$ could be evaluated to 0.7 - 0.6 = 0.1 (thus 0.1 is the smallest acceptable bid for us to receive). At some future round $window$ could be evaluated to 0.7 - 0.3 = 0.4. This is to prevent from fair offers of too small utility value to be seen as kind. This idea of decreasing window sizes is inspired from the agent strategy of [13] and the constructed function is based on the work in [16]. Then, in the next step of the algorithm, the fairness view, Fv, (see previous chapters) of the bid is calculated. If the utility of our agent is less than the value of $window$, then Fv is adjusted to:

$$Fv = Fv * our\_utility/(window) \tag{6}$$

otherwise Fv is not changed. The current $kindness$ evaluation is calculated as:

$$kindness_{new} = normalised(\Delta Fv - \Delta Fv^h) \tag{7}$$

where ΔFv is the change in Fv from the previous round. ΔFv^h is the expected change in fairness (the reference point for neutral level of kindness):

$$\Delta Fv^h = (-Fv_{prev})/(number\_of\_bids\_left) \tag{8}$$

Thus, the base level of kindness is seen as a uniform progression towards the ideal fair distribution. This method is very similar to the one used in [4], with the difference being that that one is used for one-shot games of perfect information. While in the context of negotiation we should measure their progress to a good outcome, as initially both parties are merely stating their positions, which would probably be unfair. The idea of progression is loosely based on the work of [20]. Finally, if the current kindness level is above a certain threshold, $kindness$ is updated in the following manner:

$$kindness = \alpha * kindess_{prev} + (1 - \alpha) * kindness_{new} \tag{9}$$

Here α can be thought of as the learning parameter. This formula is adapted from [14], where it is used for trust between nodes in a peer to peer network.

A.4 Dynamic Fairness Example - proportional distribution

This method is based on the fairness view used in Appendix A.1. Given a kindness evaluation, the dynamic fairness can be computed as:

$$kindness\_add = \begin{cases} \pi/4, & \text{if } kindness{=}1 \\ atan(1/(1 - kindness)) - \pi/4, & \text{if } kindness{>}0 \\ -atan(-1 - kindness) - \pi/4, & \text{if } kindness{<}0 \end{cases}$$

$$Fairness\_view = normalise(abs(atan(u1/u2) - atan(prop) + kindness\_add)) \tag{10}$$

$$Fd = -1 + exp(-(fairness\_view/std)^2) \tag{11}$$

What this set of formulae essentially does is shift the proportion (or the angle) of the desired distribution in our favour if $kindness$ is less than 0 and in our opponent's if greater than 0. The maximal change that can happen in either direction is $\pi/4$. This decision was largely arbitrary and was a way to prevent the dynamic fairness for swinging too greatly between rounds.

The following two graphics shows the kindness evaluation against a Boulware agent and a HardLiner Agent at each round (of 15 bids). The fairness view used was that of the Themis Kalai (proportional distribution of 1:1). Note that the kindness evaluation started only at the halfway point of the negotiation as before that it is assumed that the agent is merely stating their position.

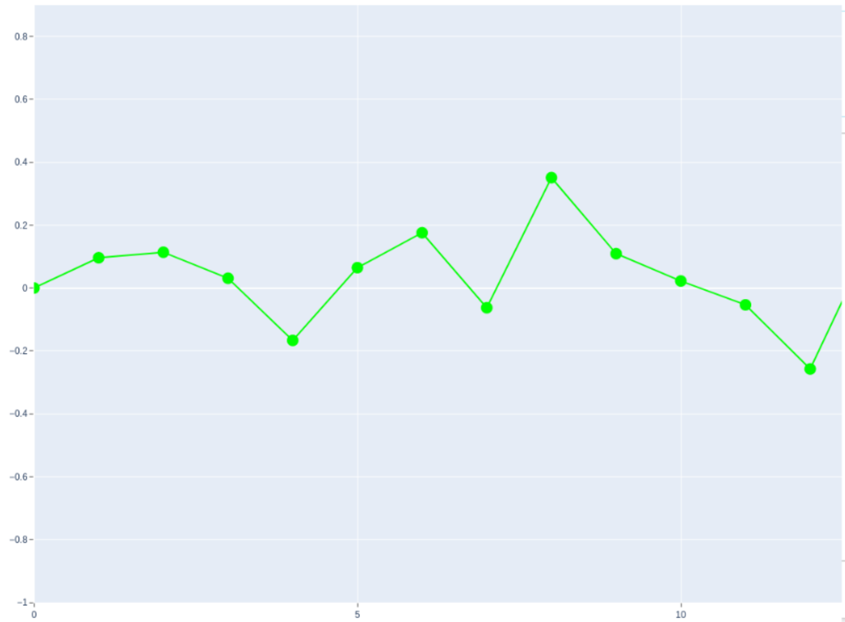

Figure 2: Kindness evaluation against a Boulware agent

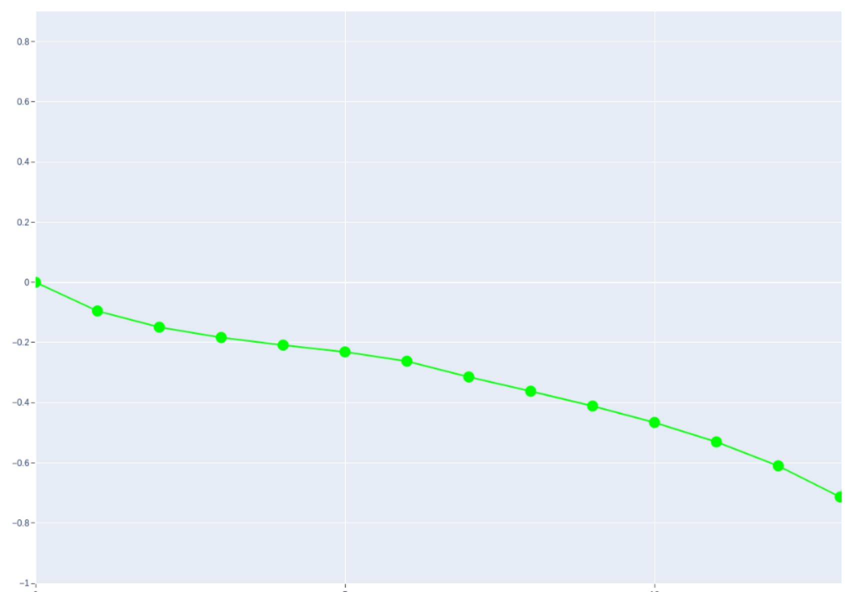

Figure 3: Kindness evaluation against a HardLiner agent

B Dynamic windows

Instead of using windows of fixed length, windows can be created dynamically on the fly based on the bids. In most of my trials this lead to better opponent preference predictions and to better identification when the opponent's strategy has changed.

The algorithm is relatively simple. First the agent's received bid's utilities are listed in chronological order. Then B-spline interpolation is used to "clean up" any sudden changes and thus focus more on the general trend. What are good parameters needs to be tested further. Then agglomerative hierarchical clustering with ward's linkage is used to create clusters of size between some min and max (10 and 20 in my trials).

An example can be seen in Figure 4. There it can be clearly seen that the first three windows consist of Boulware's strategy of holding close to its best bid. Then the boundary between the fourth and third window clearly marks the beginning of the concession phase.

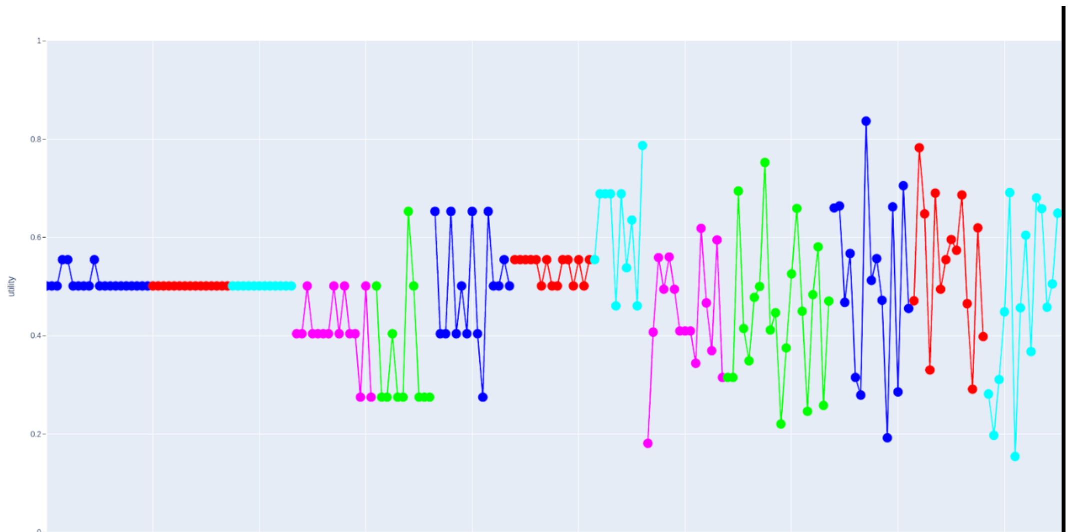

Figure 4: Clustering of the Boulware strategy

C Agent Settings

C.1 Themis Kalai and Themis Nash

$p = 0.9$
$prop = 1$
$c1 = 0.8$
$c2 = 0.7$
$c3 = 0.3$
$k = 15$. This pertains to the window size of bids for a round.
$t\_max = 0.95$

C.2 Themis Kind

$p = 0.7$
$prop = 0.5$
$c1 = 0.8$
$c2 = 0.5$
$c3 = 0.7$
$k = 15$. This pertains to the window size of bids for a round.
$t\_max = 0.95$

C.3 Themis Rude

$p = 0.7$
$prop = 2$
$c1 = 0.8$
$c2 = 0.5$
$c3 = 0.7$
$k = 15$. This pertains to the window size of bids for a round.
$t\_max = 0.95$

References

[1] Welsh, Nancy. 2004. "Fairness: Perceptions of Fairness in Negotiation". Marquette Law Review. 87.

[2] T. Baarslag, M. Kaisers, E. H. Gerding, C. M. Jonker, and J. Gratch, "When will negotiation agents be able to represent us? the challenges and opportunities for autonomous negotiators," Proceedings of the Twenty-Sixth International Joint Conference on Artificial Intelligence, 2017.

[3] Cynthia Dwork, Moritz Hardt, Toniann Pitassi, Omer Reingold, and Richard Zemel. 2012. Fairness through awareness. In Proceedings of the 3rd Innovations in Theoretical Computer Science Conference (ITCS '12). Association for Computing Machinery, New York, NY, USA, 214â226. DOI:https://doi.org/10.1145/2090236.2090255

[4] M. Rabin. 1993. Incorporating Fairness into Game Theory and Economics. The American Economic Review, 83(5), 1281â1302. http://www.jstor.org/stable/2117561

[5] Henry Cerborne. 2021. Providing a Philosophical Critique and Guidance of Fairness Metrics.

[6] Cecilia Albin. 1992. Fairness Issues in Negotiation: Structure, Process, Procedures and Outcome. IIASA Working Paper.

[7] Cecilia Albin. 1993. The Role of Fairness in Negotiation. IIASA Working Paper.

[8] Kiraz, M. S. (2008). Secure and fair two-party computation. Technische Universiteit Eindhoven. https://doi.org/10.6100/IR637269

[9] Andrychowicz, Marcin Dziembowski, Stefan Malinowski, Daniel Mazurek, Åukasz. (2014). Fair Two-Party Computations via Bitcoin Deposits. 105-121. 10.1007/978-3-662-44774-18.

[10] V. Sanchez-Anguix, O. Tunali, R. Aydogan and V. Julian. Can social agents efficiently perform in automated negotiation? Applied Sciences, vol. 11, no. 13, p. 6022, 2021.

[11] L. Kaplow and S. Shavell, "The conflict between notions of fairness and the pareto principle" American Law and Economics Review, vol. 1, no. 1, pp. 63â77, Jan. 1999.

[12] T. Baarslag, K. Hindriks, M. Hendrikx, A. Dirkzwager, and C. Jonker, "Decoupling negotiating agents to explore the space of negotiation strategies," Novel Insights in Agent-based Complex Automated Negotiation, pp. 61â83, Jan. 2014.

[13] S. Mirzayi, F. Taghiyareh, and F. Nassiri-Mofakham, "An opponent-adaptive strategy to increase utility and fairness in agents' negotiation," Applied Intelligence, vol. 52, no. 4, pp. 3587â3603, 2021.

[14] Y. Wang and J. Vassileva, "Bayesian network-based trust model," Proceedings IEEE/WIC International Conference on Web Intelligence (WI 2003), Nov. 2003.

[15] R. Aydogan, D. Festen, K. V. Hindriks, and C. M. Jonker, "Alternating offers protocols for multilateral negotiation," Modern Approaches to Agent-based Complex Automated Negotiation, pp. 153â167, Apr. 2017.

[16] P. Faratin, C. Sierra, and N. R. Jennings, "Negotiation decision functions for autonomous agents," Robotics and Autonomous Systems, vol. 24, no. 3-4, pp. 159â182, 1998.

[17] E. Fehr and K. M. Schmidt, "Chapter 8 the economics of fairness, reciprocity and altruism â experimental evidence and new theories," Handbook of the Economics of Giving, Altruism and Reciprocity, pp. 615â691, 2006.

[18] E. Fehr and K. M. Schmidt, A theory of fairness, competition and Cooperation. Zurich: Institute for Empirical Research in Economics, University of Zurich, 1999.

[19] E. Kalai, "Proportional solutions to bargaining situations: Interpersonal utility comparisons," Econometrica, vol. 45, no. 7, p. 1623, 1977.

[20] T. Baarslag, K. Hindriks, and C. Jonker, "A Tit for tat negotiation strategy for real-time bilateral negotiations" Studies in Computational Intelligence, pp. 229â233, 2013.

[21] J. F. Nash, "The bargaining problem," Econometrica, vol. 18, no. 2, p. 155, 1950.

[22] E. Kalai and M. Smorodinsky, "Other solutions to Nash's bargaining problem," Econometrica, vol. 43, no. 3, p. 513, 1975.

[23] K. Fujita, T. Ito, T. Baarslag, K. Hindriks, C. Jonker, S. Kraus, and R. Lin, "The Second Automated Negotiating Agents Competition (ANAC2011)," Studies in Computational Intelligence, pp. 183â197, 2013.

[24] Williams, C.R., Robu, V., Gerding, E.H. and Jennings, N.R. An overview of the results and insights from the third automated negotiating agents competition (ANAC2012). In Novel Insights in Agent-Based Complex Automated Negotiation; Springer: Berlin/Heidelberg, Germany, 2014; pp. 151â162