

An Autonomous Negotiating Agent Framework with Reinforcement Learning Based Strategies and Adaptive Strategy Switching Mechanism

Ayan Sengupta
 NEC Corporation
 Tokyo, Japan
 a-sengupta@nec.com

Yasser Mohammad
 NEC Corporation
 Tokyo, Japan
 y.mohammad@nec.com

Shinji Nakadai
 NEC Corporation
 Tokyo, Japan
 nakadai@nec.com

ABSTRACT

Despite abundant negotiation strategies in literature, the complexity of automated negotiation forbids a single strategy from being dominant against all others in different negotiation scenarios. To overcome this, one approach is to use mixture of experts, but at the same time one problem of this method is the selection of experts, as this approach is limited by the competency of the experts selected. Another problem with most negotiation strategies is their incapability of adapting to dynamic variation of the opponent's behaviour within a single negotiation session resulting in poor performance. This work focuses on both, solving the problem of expert selection and adapting to the opponent's behaviour with our Autonomous Negotiating Agent Framework. This framework allows real-time classification of opponent's behaviour and provides a mechanism to select, switch or combine strategies within a single negotiation session. Additionally, our framework has a reviewer component which enables self-enhancement capability by deciding to include new strategies or replace old ones with better strategies periodically. We demonstrate an instance of our framework by implementing maximum entropy reinforcement learning based strategies with a deep learning based opponent classifier. Finally, we evaluate the performance of our agent against state-of-the-art negotiators under varied negotiation scenarios.

KEYWORDS

Automated Negotiation; Negotiation strategy; Reinforcement Learning

ACM Reference Format:

Ayan Sengupta, Yasser Mohammad, and Shinji Nakadai. 2021. An Autonomous Negotiating Agent Framework with Reinforcement Learning Based Strategies and Adaptive Strategy Switching Mechanism. In *Proc. of the 20th International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2021)*, Online, May 3-7, 2021, IFAAMAS, 10 pages.

1 INTRODUCTION

Negotiation has been studied for a long time from different perspectives like game theory [42], business [18], psychology [4], neuroeconomics [22] and many more. With the progress of AI-technologies, automated negotiation allows collaboration and negotiation among AI-enabled parties. Automated negotiation aims to achieve win-win

deals for all parties, while simultaneously reducing the time and effort, thus adding significant value to society as a whole [39]. But the complexity of automated negotiation still hinders the deployment of autonomous agents in real-world applications [14].

Though much research already existed in developing negotiation strategies, Automated Negotiating Agents Competition (ANAC) brought significant improvements in strategy development [32]. In spite of such improvements in strategy design, there is no single strategy that is optimal for all possible domains [30]. One natural solution is to choose a pool of strategies and use the approach of mixture of experts during negotiation. At the same time, one needs to choose an appropriate initial set of expert strategies to excel. The questions that can originate while designing such an algorithm are these: What initial set of strategies should we select? On what conditions should we switch strategies? How to improve the initial set of chosen strategies? In this work we give a solution to all three of these questions by introducing our autonomous negotiating agent framework.

The contributions of this work to the existing research in this domain are three-fold. Firstly, we propose an autonomous negotiating agent framework, which facilitates the creation of autonomous negotiating agents capable of classifying opponent's behaviour and adaptively change strategies within a single negotiation session to reach better agreements. Secondly, we propose a mechanism to update the base strategies in an algorithmic manner to improve the overall performance. Finally, we validate this framework and provide insights in general about autonomous negotiating agents by evaluating it extensively against state of the art negotiators.

The rest of the paper is organized as follows: Section 2 gives a sketch of related work in this domain, Section 3 provides the introduction to negotiation settings. Section 4 gives a detailed description of each of the components in our framework and Section 5 describes the experimental setup. Section 6 shows the evaluations of our framework and finally, we conclude with Section 7 by discussing the limitations and provide direction for future research.

2 RELATED WORK

A considerable amount of literature has already been published on autonomous negotiation strategies. However, in recent years, the success of reinforcement learning (RL) algorithms in different fields has drawn significant attention to its application in autonomous negotiation [17, 49]. A part of our work falls under the above mentioned domain. Additionally, the other part of our work is at the intersection of the domains of opponent classification and strategy

selection in autonomous negotiation. In this section we discuss the work done in both of these domains.

2.1 RL in Autonomous Negotiation

Previously many computational methods including Bayesian Learning [29, 52] and Genetic Algorithm [21, 35, 40] have been used in automated negotiation for developing and evaluating negotiation strategies. Then again, in the last couple of decades several studies have looked at the application of reinforcement learning (RL) algorithms like Q-learning [17, 20, 47, 49, 50] and REINFORCE [48] in automated negotiation. Recently, Deep Reinforcement learning (DRL) has been used to learn the target utility values [16], the acceptance strategy [43] or both bidding and acceptance strategies [19]. Moreover, authors of [15] have also shown application of DRL in concurrent bilateral negotiation.

Bakker et al. introduced RLBOA framework [17] based on the BOA architecture [12] for automated negotiating agents, where they trained the bidding strategy of the agent using Q-learning. Their approach involves discretizing utility space and using opponent modelling to choose next offer from a set of offers, where the set of offers at each time step depends on the action taken. A limitation of this method is the loss of information due to discretization of utility space and that leads to further dependence on opponent modelling for the choice of next offer. In contrast to their work, we do not use opponent modelling while training bidding strategies. Moreover, we train the bidding strategy using DRL on continuous state and action spaces.

The authors of [19] have used DRL algorithms for training both bidding and acceptance strategies in continuous state and action spaces. The state space and the action space for their approach includes actual offer from the outcome space and hence limits the scope to a particular negotiation scenario. Furthermore, in their approach one needs to train both acceptance and bidding strategy for every domain. Moreover, experimental setup and evaluations were done against fixed preference profiles, which limits the scope of applicability. In contrast, our approach considers the utility value of the offers projected to self utility axis, thus making our bidding strategies applicable to multiple negotiation scenarios. Additionally, we show the generality of our approach by evaluating in varied negotiation scenarios while training in a single negotiation domain.

Furthermore, evaluations in both [17] and [19] are against primitive agents only, whereas we evaluated our approach against GENIUS [37] based ANAC [32] winning agents.

2.2 Opponent Modelling and Strategy Selection in Automated Negotiation

The Opponent modelling is a fundamental component of BOA architecture proposed in [12]. Commonly, opponent models attempt to learn one or more of the following opponent's attribute: acceptance strategy, deadline, utility function or bidding strategy [11]. However, our approach does not fit any of these usual types of opponent models. Unlike popular approaches of learning the bidding strategy, we classify an opponent depending on the history of bids. In fact our problem of classifying the opponent falls under the domain of continuous opponent strategy classification. Under this domain for instance, authors of [44] used a hierarchical approach with fuzzy

models to perform the opponent strategy classification in a real time strategy game. Preference profile learning by classifying the negotiation trace was done in [29, 38, 39] using Bayesian learning to determine the best match for the opponent's preference profile. However, in this work we classify the opponents bidding behaviour periodically with respect to a set of negotiator's bidding behaviours and select an appropriate strategy for negotiation within a single negotiation session.

Inspired from algorithm selection method [36], authors of [30] developed a meta-agent, which predicts the performance of a set of bilateral negotiation negotiators based on features of domains, and accordingly chooses the negotiator expected to perform best for the given negotiation scenario. Extending this idea to multilateral negotiation settings and using the approach of mixture of experts, authors of [26] solved the problem of how to combine multiple experts. However, in all these approaches a single negotiator is selected throughout a negotiation session. In contrast to that, this paper focuses on selection and switching (or combination) of strategies within a single negotiation session based on the opponent behaviour.

3 NEGOTIATION SETTINGS

A bilateral automated negotiation is a negotiation between two automated entities. We will denote these entities as *negotiators*. A negotiation setting consists of a negotiation protocol, the concerned negotiators and negotiation scenario [9]. A negotiation protocol defines the rules of the encounter, specifying which actions each negotiator can perform at any given moment. A negotiation scenario consists of the preference profiles of each negotiator and the negotiation domain. In this work, a *strategy* of a negotiator is the combination of an acceptance strategy and a bidding strategy. Additionally we denote the opponent negotiators drawn from GENIUS platform [37] as *agents*.

The negotiation protocol used throughout this paper is the stacked alternating offers protocol. Under this protocol, a negotiation session consists of rounds of consecutive turns where each negotiator can either make an offer, accept offer, or walk away from the negotiation [7]. The negotiation session ends if both negotiators find a joint agreement or a deadline is reached or one of the negotiators decides to walk away from the negotiation resulting in no agreement. The deadline can be measured in number of rounds or actual wall-time. Negotiations are non-repeated, that is one negotiation session cannot impact actions of any negotiator in subsequent sessions.

A negotiation domain consists of one or more issues. To reach an agreement, the negotiators must settle on a specific value for each negotiated issue. The outcome space of a negotiation domain denoted by Ω is the set of all possible negotiation outcomes. The outcome space can be defined as the Cartesian product of negotiation issues and is formally denoted as $\Omega = \{\omega_1, \dots, \omega_n\}$ where ω_i is a possible outcome and n is the cardinality of outcome space. A preference profile or utility profile defines a preference order \leq that ranks the outcomes in the outcome space. Usually, a preference profile of a negotiator is specified by a utility function, which assigns a utility value to an outcome ω_i denoted by $U(\omega_i)$. Utility functions are private information and the negotiators only know

their own utility functions. The preference profile of an agent also specifies a reservation value. The reservation value u_r is the utility that the negotiator receives in case of no agreement.

4 PROPOSED NEGOTIATOR FRAMEWORK

In this section we provide the structure and explain the details of our proposed Autonomous Negotiating Agent Framework, a framework that facilitates the creation of autonomous negotiating agents which are capable of classifying opponents in real time and switching strategies accordingly within a single negotiation session. First, we introduce the components of our framework and then describe an approach for designing each of the components. The proposed framework is comprised of four main components: negotiator-strategy pair, opponent classifier, strategy switching mechanism and reviewer. Figure 1 outlines all the components of the framework, while each of them are discussed in the remainder of the section.

The first component of our framework is a set of negotiator-strategy pairs, where negotiators can be any autonomous negotiating agents and strategies are bidding strategy trained against the negotiators in addition to fixed acceptance strategy for each chosen negotiator. The framework gives the user the flexibility to choose any set of negotiators which we will call the base negotiators for the rest of this paper. To illustrate, the base negotiators can include from simple time or behaviour dependent strategies [24] to state-of-the-art negotiators like ANAC winning agents. In this work, as described in Section 4.1, we have trained deep reinforcement learning (DRL) based bidding strategies against each negotiator to form *negotiator-strategy* pairs and subsequently showed the superiority and generality for such class of strategies.

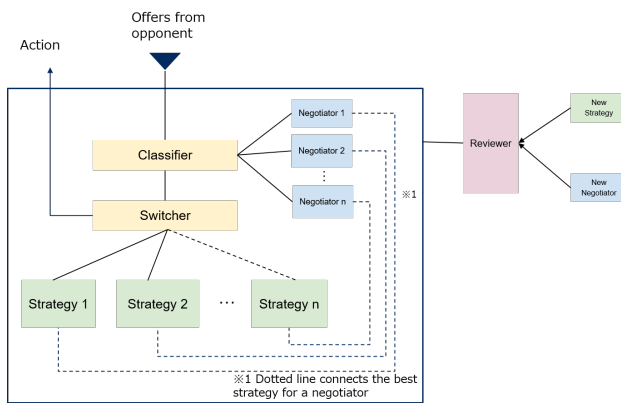


Figure 1: Block diagram of proposed framework showing the following blocks: n base negotiators blocks (blue), n trained strategies block (green), the classifier block (yellow), strategy switching block (yellow) and the reviewer block (purple). The dashed lines connecting a single negotiator to a single strategy represents negotiator-strategy pair. The components inside the solid box are utilised within a negotiation session whereas the blocks outside the solid box are used outside a negotiation session.

The second component is an opponent classifier that classifies the opponent’s bidding behaviour with respect to the bidding behaviour of base negotiators. After every negotiation round during a negotiation session, the classifier takes a sequence of opponent’s bid as input and accordingly assigns an estimated probability to each base negotiator. In this work, we have used only the sequence of opponent’s offers projected on the self utility axis as input since our implementation showed similar accuracy with additional information of self offers. Deep learning based approach, described in Section 4.2 is used for training the opponent classifier and the results in Section 6 show the versatility of such classifier in different negotiation scenarios.

The third component is a switching mechanism, that switches or combines the strategies learned against base negotiators depending on the output of the opponent classifier. While the opponent classifier classifies the behaviour of the opponent negotiator at every time step, the switching mechanism has added flexibility of changing decisions after certain intervals. Note that the strategy switching as narrated in Section 4.3 is performed within a single negotiation session, and hence makes our negotiator framework adaptive.

All the three aforementioned components are active components that function when a negotiation is underway. Whereas, the Reviewer is a passive component that does not actively take part in the negotiation process. Outside the negotiation session, the Reviewer provides a mechanism that decides if a new negotiator or a new strategy should be included in the framework. To the best of our knowledge, all the meta-agent strategies proposed in literature do not have a mechanism that can enhance their capability by evaluating and adding new strategies. This component is crucial to the design of our framework as it insures the framework against depreciation in the future. In Section 4.4, we provide the algorithm of the Reviewer and in the remaining section, we discuss the approaches for building each component of our framework.

4.1 Deep Reinforcement Learning Based Strategies

The prime components of our framework are the negotiation strategies trained against the base negotiators. The whole framework is based on the presumption that one can successfully learn an effective strategy against each of the base negotiators. Additional requirement is that the approach should facilitate the framework to perform well in a domain-independent manner. Due to recent success of RL algorithms in training strategies of automated negotiators, we used Soft Actor-Critic [27, 28], a DRL algorithm to train a bidding strategy against each base negotiator. For acceptance strategy we adopted the approach of combined acceptance condition as proposed in [13]. In contrast to RLBOA framework in [17], no opponent model component is used while training against an opponent negotiator.

4.1.1 State Space, Action Space and Reward. A major problem for developing a domain independent negotiator framework is the fact that the outcome space Ω varies significantly across different negotiation scenarios. Moreover the offers $\omega_i \in \Omega$ are usually non numerical in nature, which again demands an approach to convert the offers to numerical value in a meaningful way. To overcome

both of these problems, we took a similar approach to [17] and represented every outcome ω_i by $U_s(\omega_i)$ where U_s is the self utility function. To avoid any information loss we considered the continuous outcome space rather than discretizing it. For our DRL approach everything in the negotiation scenario including the opponent is considered as the environment. Let us denote the state and action in an environment as s_t and a_t respectively. The state consists of only the information about the offers and the action determines what utility value to bid next. For a negotiation session with time limit T , we defined our state space and action space as

$$\begin{aligned} s_t &= \{t_r, U_s(\omega_s^{t-2}), U_s(\omega_o^{t-2}), U_s(\omega_s^{t-1}), \\ &\quad U_s(\omega_o^{t-1}), U_s(\omega_s^t), U_s(\omega_o^t)\} \\ a_t &= u_s^{t+1} \text{ such that } u_r < u_s \leq 1 \end{aligned}$$

where t_r denotes the relative time, ω_s^t and ω_o^t denotes offers by self and opponent at time step $t < T$ respectively. The self reservation value is denoted by u_r and u_s^{t+1} denotes the utility value of the next offer. To get the actual offer from the utility value we need an inverse map $U_s^{-1} : u_s \rightarrow \omega_s$ of the self utility function U_s which can be a one-to-one or one-to-many mapping. One simple way of defining the inverse utility function is given in Equation (1).

$$\begin{aligned} U_s^{-1}(u_s) &= \underset{\omega}{\operatorname{argmin}} f(\omega), \quad \text{where} \\ f(\omega) &= (U_s(\omega) - u_s)^2 \quad \forall \omega \in \Omega. \end{aligned} \quad (1)$$

The goal of the strategies trained is to maximise the average utility against the corresponding base negotiator. So the reward function R is defined as

$$R(s_t, a_t, s_{t+1}) = \begin{cases} U_s(\omega_a), & \text{if there is an agreement } \omega_a \\ -1, & \text{for no agreement and } s_{t+1} \\ & \text{is terminal state,} \\ 0, & \text{otherwise.} \end{cases}$$

There is an immediate reward of 0 after every step in a negotiation session when negotiation has not ended.

4.1.2 Soft Actor-Critic Algorithm. Soft actor-critic (SAC) [27, 28] is an off-policy algorithm based on maximum entropy reinforcement learning that aims to maximize both the expected reward and the policy's entropy. Policies with higher entropy have more randomness, which means that maximum entropy reinforcement learning learns a policy that has maximum randomness yet achieves a high reward. Normal reinforcement learning algorithms try to maximize the expected reward only. On the contrast the reason for maximizing the the entropy of the policy is to improve both the algorithm's robustness to hyperparameters and its sample efficiency [27]. In automated negotiation, this randomness is desirable to reduce the opponent's ability to predict the behaviour of an agent and exploit this information.

An optimal policy π^* in entropy-regularized reinforcement learning can be expressed as

$$\pi^* = \operatorname{arg max}_{\pi} \mathbb{E}_{\pi} \left[\sum_{t=0}^{\infty} \gamma^t \left(R(s_t, a_t, s_{t+1}) + \alpha H(\pi(\cdot|s_t)) \right) \right],$$

where R is the reward function, γ is the discount factor, H denotes the entropy of policy π and $\alpha > 0$ is the entropy regularization coefficient. s_t and a_t denotes the state and action at time-step t

respectively. Now, the corresponding action-value function $Q^{\pi}(s, a)$ for state s and action a can be expressed as

$$Q^{\pi}(s, a) = \mathbb{E}_{\pi} \left[\sum_{t=0}^{\infty} \gamma^t R(s_t, a_t, s_{t+1}) + \alpha \sum_{t=1}^{\infty} \gamma^t H(\pi(\cdot|s_t)) \mid s, a \right]$$

SAC concurrently learns a policy π and two Q-functions. In our implementation we used the approach proposed by Haarnoja et al. [28] where entropy regularization parameter α is also a trainable parameter. The details of the hyperparameters for our implemented model are provided in the supplementary material of [45].

4.2 Deep Learning Based Opponent Classifier

Opponent modelling is a fundamental block of BOA architecture proposed in [12]. Although we do not have an opponent modelling block while training the DRL based strategies, our framework contains a classifier that classifies an unknown opponent's bidding behaviour with respect to the base negotiator's behaviour in the framework. Our approach uses 1D-Convolutional Neural Networks (1D-CNN) based classifier to classify an unknown opponent at every time step of a negotiation. In the following sections we will describe our classifier's input/output and the model architecture.

4.2.1 Input and Output. The input to the classifier is a sequence of offers by the opponent projected to self utility axis. Similar to Section 4.1, we will denote each offer as ω_i and the self utility value for the offer as $U_s(\omega_i)$. The choice of using self utility value ensures the input sequence to be numerical all the time and can be directly provided to the classifier without any requirement of pre-processing. Another significant benefit is, it allows the framework to work across different negotiation scenarios without retraining the classifier for each domain. The output of the classifier are the estimated probabilities for each base negotiator. Let us denote the input to the classifier at the current negotiation time step t as I_c^t and output as O_c^t . Then,

$$\begin{aligned} I_c^t &= \{U_s(\omega_i)\}_{i=t-k}^{i=t-1} \quad \text{where } k \in \mathbb{Z}_+ \text{ and } k > 1. \\ O_c^t &= [p_1, \dots, p_n] \end{aligned}$$

where n is the number of base negotiators in the framework. \mathbb{Z}_+ denotes the set of positive integers. p_i denotes the estimated probability of the opponent behaving as the i^{th} negotiator. Values of input array I_c^t are zero before the first opponent offer, that is, $U_s(\omega_i) = 0$ if $i < 0$. k is the window length and is fixed before training. Moreover, greater the value of k , greater is the information provided to the classifier.

4.2.2 Classifier Model. The input to the classifier at every time step I_c^t is a time-series. For such time series data, Long short-term memory (LSTM) or recurrent neural network (RNN) architectures perform incredibly well. But, a major problem with LSTM and RNN is that they require datasets of massive sizes and large computational resources for training. To overcome such difficulties 1D-CNNs have shown great promise [33]. 1D-CNN based classifiers have been successfully used in structural damage detection [2, 5], fault detection in modular multilevel converters [34], condition monitoring in rotating mechanical machine parts [23, 31].

In our classifier model, we have consecutive 1D-CNN layers followed by consecutive Dense layers. The depth of the model will

increase or decrease with the increase or decrease of the window length k respectively. Moreover, a greater value of k will result in a larger part of history of opponent offers to be considered by the classifier. This will reduce the model accuracy. On the other hand a very small value of k will make the model myopic and will be error prone in classification. Some hyperparameter tuning is required for adjusting the value of the window length. The model architecture and the hyperparameters are provided in the supplementary material of [45].

4.3 Strategy Switching Mechanism

Depending on the output of the classifier O_c^t , this component switches or combines the strategies to take the next action. The algorithm for the switching mechanism is provided in Algorithm 1. Although the opponent classification is done at every time step, the approach of choosing next offer need not change after every time step. In Algorithm 1 the parameter β_i where $i \in [1, n]$ tunes the algorithm from a hard switcher to a combination mechanism. The algorithm becomes a pure switching algorithm if $\beta_i = 1$ and $\beta_k = 0 \forall k \neq i$.

Algorithm 1: Algorithm for strategy switching

Input: $O_c^t = \{p_1, \dots, p_n\}$ from the opponent classifier
Data: $S_b = \{s_1, \dots, s_n\}$, the set of base strategies corresponding to the set of base negotiators
 $N_b = \{N_1, \dots, N_n\}$
Output: Action: next offer or Accept
 Choose initial strategy $s_{init} \in S_b$, $S_{action} = s_{init}$
 $i = \text{argmax}(O_c^t)$ where N_i denotes the base negotiator with highest classification probability
if action by strategy s_i is Accept **then**
 | Accept opponent's offer
else
 | $w^{t+1} = U_s^{-1}(\sum_{k=0}^{k=n} \{\beta_k * u_{s_k}\})$ where u_{s_k} is the utility value by strategy s_k for next time step and β_k is the weight parameter
end

4.4 Reviewer Mechanism

This component enables the addition of new negotiators or strategies or both to the RL-agent instantiated by our framework. To show the basic operation of this component, we implemented an evaluation based approach for the Reviewer. The algorithm is provided in Algorithm 2 where the parameters α and β are threshold parameters. When a new negotiator N_{new} is introduced to the Reviewer, first a new strategy S_{train} is trained against it. Subsequently, S_{train} and the RL-agent are evaluated against N_{new} . Finally, depending of the parameter α if the evaluation with the S_{train} is better in comparison with RL-agent, then the reviewer will provide confirmation and S_{train} will added to the pool of strategies and the classifier will be retrained with a new class N_{new} .

Moreover, The new trained strategy S_{train} is cross-evaluated with base negotiators and compared with the base strategies. Depending on the evaluation and parameter β , base strategies may

Algorithm 2: Algorithm for Reviewer component

Input: New strategy S_{new} or new negotiator N_{new}
Data: $S_b = \{s_1, \dots, s_n\}$, the set of base strategies corresponding to the set of base negotiators
 $N_b = \{N_1, \dots, N_n\}$, $Eval$ function that provides an evaluation score of strategy for a negotiator.
Function StrategyEvaluation($S_{test}, N_i, s_i, \beta$):
 $\hat{e}_k = Eval(N_i, S_{test})$
 $e_k = Eval(N_i, s_i)$
if $\hat{e}_k \geq \beta * e_k$ **then**
 | **return** Accept and replace s_i with S_{test}
else
 | **return** Reject
end
if Input is N_{new} **then**
 | Train new strategy S_{train} against N_{new}
 $e_f = Eval(N_{new}, \text{RL-agent})$
 $e_s = Eval(N_{new}, S_{train})$
if $e_s \geq \alpha * e_f$ **then**
 | Accept N_{new} and S_{train}
else
 | Reject
end
for ($k \in [1, n]$)
 | StrategyEvaluation($S_{train}, n_k, s_k, \beta$)
end
if Input is S_{new} **then**
 | **for** ($k \in [1, n]$)
 | StrategyEvaluation(S_{new}, n_k, s_k, β)
end

be updated with S_{train} . In this manner the Reviewer provides a mechanism for gradual improvement of the agent.

5 EXPERIMENTAL SETUP

The goals of our experiments are two fold. First, we introduce our RL-agent based on the proposed framework with small number of base negotiators and show that the RL-agent, while generalizing over negotiation scenarios performs on average better than state of the art ANAC winning agents. Secondly, we show the value of the reviewer mechanism by evaluating and adding new negotiators and subsequently show improvement of our RL-agent.

We analyzed our proposed system in 18 domains of ANAC 2013 with cardinality of outcome space ranging from 3 to 56700 and opposition [10] ranging from 0.002 to 0.606 as shown in Table 1. All the negotiation experiments are conducted using the NEGotiation MultiAgent System (NegMAS) platform [41, 51]. For the purpose of calculating benchmarks and evaluating performance we used the given preference profiles of ANAC 2013 for each domain. Among the negotiation settings, the reserved value is kept zero and the discount factor is ignored for all negotiations. Moreover, we used min-max normalisation for normalising the utility values between

Table 1: Overview of the ANAC 2013 domains. We considered 18 domains with a pair of utility functions.

Domain	Opposition	Outcome Space
Acquisition	0.104	384
Animal	0.15	1152
Camera	0.076	3600
Coffee	0.279	112
Defensive Charms	0.193	36
Dog Choosing	0.002	270
Fifty Fifty	0.498	11
House Keeping	0.13	384
Ice-cream	0.01	720
Kitchen	0.219	15625
Laptop	0.076	27
Lunch	0.246	3840
Nice Or Die	0.177	3
Outfit planes	0.049	128
Smart Phone	0.606	27
Smart Phone	0.022	12000
Ultimatum	0.319	9
Wholesaler	0.128	56700

0 and 1. For performance comparisons average utility values are calculated on negotiation data obtained from 50 to 100 negotiations between a pair of agents for each negotiation scenario.

For training a bidding strategy against a given negotiator we generate a random utility function for the opponent and used fixed self utility function. This ensures the that maximum entropy reinforcement learning algorithm can learn a stochastic policy that performs well in varied negotiation scenarios. Additionally, the training of strategies are done in a single domain and evaluations are done in all 18 domains. For our experiments all strategies are trained on the *camera* domain and tested across other domains. Moreover, for simplicity all hyperparameters of SAC algorithm were kept fixed while training against different negotiators. Training of RL-agent was done using the TF-Agents [46] library.

For the opponent classifier we used a window length of 20 and the model consists of three consecutive 1D-CNN layers followed by two Dense layers. The training data for the classifier was generated by multiple simulations of the base negotiators while training the RL strategies against them. Additionally, the classifier is trained only on the data generated in the camera domain while the same classifier has been used in the evaluation for all other domains. The training of the classifier was done using TensorFlow library [1].

Evaluations of our proposed framework are done by first instantiating an RL-agent with the number of base negotiators $n = 1$ and then increasing it to $n = 3$ with the help of Reviewer mechanism. We used parameters $\alpha = \beta = 1.1$ for the Reviewer and for simplicity, opted for a pure switching mechanism instead of combining strategies. Furthermore, we conducted t-tests with Bonferroni’s conservative multiple-comparisons correction [3] for analysing if the difference in mean utilities are statistical significant.

6 RESULTS

In this section we present the results according to the experimental setup of Section 5. First we will present the detailed results of our experiments with 7 ANAC winning agents. Next we present the evaluations of the Reviewer mechanism. Finally, we present a summarised result against 47 GENIUS agents. For comparison, we created three different benchmarks and then compared the performance of RL-agent against each of the benchmarks. Now, $U_a^{a \times b : d}$ denotes the average utility achieved by agent a against b in domain d over 100 runs with two different utility functions. A and D denotes the set of agents and domains respectively over which the benchmark is calculated and $|\cdot|$ denotes the cardinality of a set.

- (1) *Self utility benchmark*: In this benchmark, score of an agent $S_a = \frac{1}{|A| \times |D|} \sum_{d \in D} \sum_{b \in A} U_a^{a \times b : d}$ is the mean utility acquired by the agent a when negotiating with every agent $b \in A$ in all negotiation scenarios.
- (2) *Utility against opponent benchmark*: In this benchmark, score against an agent a , $O_a = \frac{1}{(|A|-1) \times |D|} \sum_{d \in D} \sum_{b \in A/a} U_b^{a \times b : d}$ denotes the mean utility acquired by agents $b \in A/a$ while negotiating with agent a in all negotiation scenarios.
- (3) *Domain utility benchmark*: In this benchmark, score of a domain d , $D_d = \frac{1}{|A|} \sum_{b \in A} U_a^{a \times b : d}$ denotes the mean utility obtained by all agents $a \in A$ in domain d , while negotiating with every agent $b \in A$.

6.1 Performance against ANAC Winning Agents

Before evaluating our agent, we first calculate the benchmark scores with 7 ANAC winning agents and then compare the score of our RL agent with each benchmarks as shown in Figure 2. The agents selected are Atlas3, ParsAgent, RandomDance, ParsCat, AgentYX, Caduceus and PonpokoAgent [6, 8, 25]¹. It is clearly visible from Figure 2, that our RL-agent outperformed all 7 agents in all the benchmarks. The error bars in Figure 2a and Figure 2b denote the standard deviation of average utilities obtained over the domains. In comparison with self utility benchmark, RL-agent performed 25% better than the agent which acquired highest average utility in that benchmark as shown in Figure 2a. Overall, our agent’s score was 37% higher than the average scores of all other agents. Glancing at the error bars, one can understand that the RL-agent has the minimum standard deviation among all other agents and hence shows the robustness of the agent in varied negotiation domains. In comparison with utility against opponent benchmark, the average utility obtained by the RL-agent outperformed the benchmark scores in a range of 11% to 50% as shown in Figure 2b. This shows that the RL-agent outperforms the average score of the opponents against each agent. Proceeding to the comparison with domain benchmark illustrated in Figure 2c, one can clearly visualize that the scores of RL-agent is better than the highest score by any agent in 13 out of 18 domains. In fact the average score of RL-agent is more than the utility benchmark by a range of 4% to 450%.

¹Atlas3 (2015 winner), ParsAgent (2015 2nd position), RandomDance (2015 3rd position), Caduceus (2016 winner), ParsCat (2016 2nd position), AgentYX (2016 2nd position) and PonpokoAgent (2017 winner)

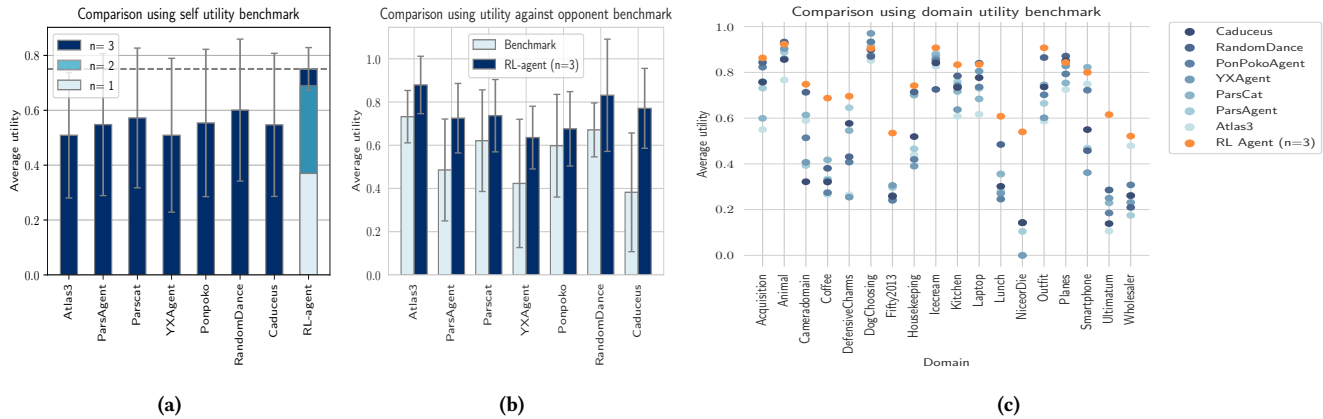


Figure 2: (a) Comparison of RL-agent with self utility benchmark consisting of 7 ANAC winning agents. It also shows the performance of RL-agent with $n = 1$, $n = 2$ and $n = 3$. (b) Comparison of RL-agent with utility against opponent benchmark consisting of 7 ANAC winning agents. (c) Comparison of RL-agent with domain benchmark consisting of 18 domains.

6.2 Performance of Reviewer

Although the final comparison shown in Figure 2 is the instance of RL-agent with three negotiator-strategy pairs, the framework was first initialized with only one negotiator, that is, with the Random negotiator. One RL based strategy was trained against it and included in the negotiator-strategy pair. Opponent classifier training was not needed as number of base negotiator $n = 1$. As expected the performance of the initial RL-agent was poor as shown in Figure 2a. Next we introduced a simple time dependent agent called Boulware agent [24], to the Reviewer. After receiving acceptance from reviewer, new RL based strategy was added to the strategy pool and the opponent classifier was trained for $n = 2$. In a similar way, a new behavioural strategy based agent, Naive tit-for-tat and corresponding RL based strategy was included in the RL-agent and classifier was trained for $n = 3$. Furthermore, the strategy trained against random negotiator was replaced by the RL strategy trained against Naive tit-for-tat as per the evaluation of the Reviewer. Moving ahead, we introduced other agents from the pool of 7 ANAC agents to the Reviewer, but the Reviewer rejected the inclusion of any additional agents. It is to be noted that all evaluations by the Reviewer were also restricted to a single domain only (camera domain). The performance of the RL-agent with different number of base negotiators is shown in Figure 2a. It can be clearly seen that the performance of the RL-agent has increased with the addition of strategies and hence shows the significance of the Reviewer mechanism in our framework. To keep the visualisations simple, further comparisons with other benchmarks as shown in Figure 2b and Figure 2c are only shown with number of base negotiators $n = 3$. The results obtained have two fold significance, firstly it shows that the bidding behaviour of all 7 opponent agents can be approximated by a piece-wise function of Boulware type, behavioural type or random type negotiator. So the strategy switching technique within a negotiation session, with strategies specifically trained for three basic negotiators works undoubtedly well against these 7 opponents. Secondly, it shows the Reviewer mechanism plays a

key role in addition of base strategies and thereby improving the RL-agent’s overall performance.

6.3 Performance against GENIUS Agents

To show the versatility of the proposed framework, we chose the already created RL-agent with only 3 base negotiators and evaluated it against the benchmarks scores of 47 GENIUS agents which includes ANAC competitors from year 2015 to 2017. The results are illustrated in Figure 4 and Figure 3 where the error bars denote the standard deviation over the domains. In comparison with self utility benchmark, our RL-agent performed better than other agents in the range of 11% to 139%. The average improvement is 37.4% with an improvement of more than 50% against 12 agents as shown in Figure 4a. The relatively low standard deviation marks the robustness of our agent in different domains. Next, comparison with the utility against opponent benchmark as shown in Figure 3, reveals that the RL-agent outperformed in 40 out of 47 comparisons. Individual performance improvement ranges from 1% to 428% with an improvement of over 50% against 19 agents and 100% against 5 agents. Finally, comparison with domain benchmark is shown in Figure 4b, where it is visible that the scores of RL-agent is better than the highest score by any agent in 13 out of 18 domains. Additionally, the RL-agent outperformed the average benchmark scores in all the domains. The performance improvement over the benchmark scores in each domain ranges from 10% to 488% with an improvement of at least 50% in 7 domains. Finally, we calculate the statistical significance of the differences. For the utility differences to be statistically significant, following Bonferroni’s conservative multiple-comparisons correction the p-values of each t-tests in self utility benchmark, utility against opponent benchmark and domain benchmark should be less than 0.0011, 0.0011 and 0.0028 respectively. It turns out that for self utility benchmark, differences of utility in 30 out of 47 comparisons were statistically significant whereas in utility against opponent benchmark 32 out of 47 were statistically significant. In case of domain benchmark, all differences of utilities were statistically significant.




Figure 3: Comparison of RL-agent with utility against opponent benchmark consisting of 47 GENIUS Agents in 18 domains.




Figure 4: Comparison of the performance of RL-agent (a) with self utility benchmark consisting of 47 GENIUS Agents and (b) with domain benchmark consisting of 18 domains.

7 CONCLUSION AND FUTURE WORK

In this work we proposed an autonomous negotiating agent framework with four components: negotiators paired with trained strategies, an opponent classifier, a switching mechanism and a reviewer mechanism. Strategies included are RL based strategies whereas strategy switching depends on the classification probabilities of the opponent classifier. The proposed opponent classifier classifies the opponent's bidding behaviour with respect to the base negotiator's behaviour at every time step thus allows the agent to switch or combine strategies within a single negotiation session. These functionalities together gives our RL-agent versatility even with a small pool of base negotiators as illustrated in our evaluations. Furthermore, the reviewer mechanism helps in the decision making of adding more negotiators and strategies to the existing pool of base negotiators and strategies. This helps in incremental improvement of the RL-agent and also restricts unnecessary addition of base entities.

In our experimental setup, all training and evaluations were done by removing the discount factor and keeping the reserved utility value as zero. It would be interesting to obtain and compare the results with varying discount factors and reserved utility values. Moreover, while training bidding strategies, we have used the same acceptance strategies against all negotiators. At the same time, it has been noticed that the performance of trained strategy depends on the choice of acceptance strategy. So our future work involves training the acceptance strategy together with the bidding strategy. Also, to show the concept of reviewer mechanism, we have implemented an evaluation based reviewer. Another interesting direction could be using unsupervised clustering algorithms on the negotiators bidding behaviour to differentiate a new negotiator from the pool of negotiators. That will remove the additional step of training a new strategy each time for evaluation by the reviewer and at the same time will give us a concrete picture about the type of base negotiators that make the RL-agent perform better across various negotiation scenarios and against varied opponents.

REFERENCES

- [1] Martín Abadi, Ashish Agarwal, Paul Barham, and ET al. 2015. TensorFlow: Large-Scale Machine Learning on Heterogeneous Systems. <https://www.tensorflow.org/Software> available from tensorflow.org.
- [2] Osama Abdeljaber, Onur Avci, Serkan Kiranyaz, Moncef Gabbouj, and Daniel J Inman. 2017. Real-time vibration-based structural damage detection using one-dimensional convolutional neural networks. *Journal of Sound and Vibration* 388 (2017), 154–170.
- [3] Hervé Abdi. 2007. Bonferroni and Šidák corrections for multiple comparisons. *Encyclopedia of measurement and statistics* 3 (2007), 103–107.
- [4] J Stacy Adams. 1965. Inequity in social exchange. In *Advances in experimental social psychology*. Vol. 2. Elsevier, Amsterdam, Netherlands, 267–299.
- [5] Onur Avci, Osama Abdeljaber, Serkan Kiranyaz, and Daniel Inman. 2020. Convolutional neural networks for real-time and wireless damage detection. In *Dynamics of Civil Structures, Volume 2*. Springer, Berlin, Germany, 129–136.
- [6] Reyhan Aydoğan. 2016. ANAC2016 - Automated Negotiating Agents Competition 2016. TUDelft. Retrieved October 8, 2020 from <http://web.tuat.ac.jp/~katfui/ANAC2016/>
- [7] Reyhan Aydoğan, David Festen, Koen V Hindriks, and Catholijn M Jonker. 2017. Alternating offers protocols for multilateral negotiation. In *Modern Approaches to Agent-based Complex Automated Negotiation*. Springer, Berlin, Germany, 153–167.
- [8] Reyhan Aydoğan, Katsuhide Fujita, Tim Baarslag, Catholijn M Jonker, and Takayuki Ito. 2018. ANAC 2017: Repeated multilateral negotiation league. In *International Workshop on Agent-Based Complex Automated Negotiation*. Springer, Stockholm, Sweden, 101–115.
- [9] Tim Baarslag. 2014. *What to bid and when to stop*. Ph.D. Dissertation. Delft University of Technology.
- [10] Tim Baarslag, Katsuhide Fujita, Enrico H Gerding, Koen Hindriks, Takayuki Ito, Nicholas R Jennings, Catholijn Jonker, Sarit Kraus, Raz Lin, Valentin Robu, et al. 2013. Evaluating practical negotiating agents: Results and analysis of the 2011 international competition. *Artificial Intelligence* 198 (2013), 73–103.
- [11] Tim Baarslag, Mark JC Hendriks, Koen V Hindriks, and Catholijn M Jonker. 2016. Learning about the opponent in automated bilateral negotiation: a comprehensive survey of opponent modeling techniques. *Autonomous Agents and Multi-Agent Systems* 30, 5 (2016), 849–898.
- [12] Tim Baarslag, Koen Hindriks, Mark Hendriks, Alexander Dirkzwager, and Catholijn Jonker. 2014. Decoupling negotiating agents to explore the space of negotiation strategies. In *Novel Insights in Agent-based Complex Automated Negotiation*. Springer, Berlin, Germany, 61–83.
- [13] Tim Baarslag, Koen Hindriks, and Catholijn Jonker. 2014. Effective acceptance conditions in real-time automated negotiation. *Decision Support Systems* 60 (2014), 68–77.
- [14] Tim Baarslag, Michael Kaisers, Enrico H. Gerding, Catholijn M. Jonker, and Jonathan Gratch. 2017. When Will Negotiation Agents Be Able to Represent Us? The Challenges and Opportunities for Autonomous Negotiators. In *Proceedings of the Twenty-Sixth International Joint Conference on Artificial Intelligence, IJCAI-17*. IJCAI, Melbourne, Australia, 4684–4690. <https://doi.org/10.24963/ijcai.2017/653>
- [15] Pallavi Bagga, Nicola Paoletti, Bedour Alrayes, and Kostas Stathis. 2020. A Deep Reinforcement Learning Approach to Concurrent Bilateral Negotiation. In *Proceedings of the Twenty-Ninth International Joint Conference on Artificial Intelligence, IJCAI-20*, Christian Bessiere (Ed.). International Joint Conferences on Artificial Intelligence Organization, Yokohama, Japan, 297–303. <https://doi.org/10.24963/ijcai.2020/42> Main track.
- [16] Pallavi Bagga, Nicola Paoletti, and Kostas Stathis. 2020. Learnable Strategies for Bilateral Agent Negotiation over Multiple Issues. arXiv:2009.08302 [cs.MA]
- [17] Jasper Bakker, Aron Hammond, Daan Bloembergen, and Tim Baarslag. 2019. RLBOA: A Modular Reinforcement Learning Framework for Autonomous Negotiating Agents.. In *AAMAS IFAAMAS*, Montreal, Canada, 260–268.
- [18] Max H Bazerman, George F Loewenstein, and Sally Blount White. 1992. Reversals of preference in allocation decisions: Judging an alternative versus choosing among alternatives. *Administrative science quarterly* 37, 2 (1992), 220–240.
- [19] Ho-Chun Herbert Chang. 2020. Multi-Issue Bargaining With Deep Reinforcement Learning. arXiv:2002.07788 [cs.MA]
- [20] Lihong Chen, Hongbin Dong, Qilong Han, and Guangzhe Cui. 2013. Bilateral multi-issue parallel negotiation model based on reinforcement learning. In *International Conference on Intelligent Data Engineering and Automated Learning*. Springer, Hefei, China, 40–48.
- [21] Nirmal Choudhary and Kamal Kant Bharadwaj. 2019. Evolutionary learning approach to multi-agent negotiation for group recommender systems. *Multimedia Tools and Applications* 78, 12 (2019), 16221–16243.
- [22] Giorgio Coricelli and Rosemarie Nagel. 2009. Neural correlates of depth of strategic reasoning in medial prefrontal cortex. *Proceedings of the National Academy of Sciences* 106, 23 (2009), 9163–9168.
- [23] Levent Eren, Turker Ince, and Serkan Kiranyaz. 2019. A generic intelligent bearing fault diagnosis system using compact adaptive 1D CNN classifier. *Journal of Signal Processing Systems* 91, 2 (2019), 179–189.
- [24] Peyman Faratin, Carles Sierra, and Nick R Jennings. 1998. Negotiation decision functions for autonomous agents. *Robotics and Autonomous Systems* 24, 3-4 (1998), 159–182.
- [25] Katsuhide Fujita, Reyhan Aydoğan, Tim Baarslag, Koen Hindriks, Takayuki Ito, and Catholijn Jonker. 2017. The sixth automated negotiating agents competition (ANAC 2015). In *Modern Approaches to Agent-based Complex Automated Negotiation*. Springer, Berlin, Germany, 139–151.
- [26] Taha D Güneş, Emir Arditi, and Reyhan Aydoğan. 2017. Collective voice of experts in multilateral negotiation. In *International Conference on Principles and Practice of Multi-Agent Systems*. Springer, Nice, France, 450–458.
- [27] Tuomas Haarnoja, Aurick Zhou, Pieter Abbeel, and Sergey Levine. 2018. Soft Actor-Critic: Off-Policy Maximum Entropy Deep Reinforcement Learning with a Stochastic Actor. arXiv:1801.01290 [cs.LG]
- [28] Tuomas Haarnoja, Aurick Zhou, Kristian Hartikainen, George Tucker, Sehoon Ha, Jie Tan, Vikash Kumar, Henry Zhu, Abhishek Gupta, Pieter Abbeel, and Sergey Levine. 2019. Soft Actor-Critic Algorithms and Applications. arXiv:1812.05905 [cs.LG]
- [29] Koen Hindriks and Dmytro Tykhonov. 2008. Opponent modeling in automated multi-issue negotiation using bayesian learning. In *Proceedings of the 7th international joint conference on Autonomous agents and multiagent systems-Volume 1*. IFAAMAS, Estoril, Portugal, 331–338.
- [30] Litan Ilany and Ya'akov Gal. 2016. Algorithm selection in bilateral negotiation. *Autonomous Agents and Multi-Agent Systems* 30, 4 (2016), 697–723.
- [31] Turker Ince, Serkan Kiranyaz, Levent Eren, Murat Askar, and Moncef Gabbouj. 2016. Real-time motor fault detection by 1-D convolutional neural networks. *IEEE Transactions on Industrial Electronics* 63, 11 (2016), 7067–7075.
- [32] Catholijn M Jonker, Reyhan Aydoğan, Tim Baarslag, Katsuhide Fujita, Takayuki Ito, and Koen Hindriks. 2017. Automated negotiating agents competition (ANAC). In *Thirty-first AAAI conference on artificial intelligence*. AAAI, San Francisco, California, USA, 5070–5072.
- [33] Serkan Kiranyaz, Onur Avci, Osama Abdeljaber, Turker Ince, Moncef Gabbouj, and Daniel J. Inman. 2019. 1D Convolutional Neural Networks and Applications: A Survey. arXiv:1905.03554 [eess.SP]
- [34] Serkan Kiranyaz, Adel Gastli, Lazhar Ben-Brahim, Nasser Al-Emadi, and Moncef Gabbouj. 2018. Real-time fault detection and identification for MMC using 1-D convolutional neural networks. *IEEE Transactions on Industrial Electronics* 66, 11 (2018), 8760–8771.
- [35] Raymond YK Lau, Maolin Tang, On Wong, Stephen W Milliner, and Yi-Ping Phoebe Chen. 2006. An evolutionary learning approach for adaptive negotiation agents. *International journal of intelligent systems* 21, 1 (2006), 41–72.
- [36] Kevin Leyton-Brown, Eugene Nudelman, Galen Andrew, Jim McFadden, and Yoav Shoham. 2003. A portfolio approach to algorithm selection. In *IJCAI*, Vol. 3. Morgan Kaufmann Publishers, Acapulco, Mexico, 1542–1543.
- [37] Raz Lin, Sarit Kraus, Tim Baarslag, Dmytro Tykhonov, Koen Hindriks, and Catholijn M. Jonker. 2014. Genius: An Integrated Environment for Supporting the Design of Generic Automated Negotiators. *Computational Intelligence* 30, 1 (2014), 48–70. <https://doi.org/10.1111/j.1467-8640.2012.00463.x>
- [38] Raz Lin, Sarit Kraus, Jonathan Wilkenfeld, and James Barry. 2006. Bounded Rational Agents with Incomplete Information'. In *ECAI 2006: 17th European Conference on Artificial Intelligence, August 29-September 1, 2006, Riva Del Garda, Italy; Including: Prestigious Applications of Intelligent Systems (PAIS 2006); Proceedings*, Vol. 141. IOS Press, Riva Del Garda, Italy, 270.
- [39] Raz Lin, Sarit Kraus, Jonathan Wilkenfeld, and James Barry. 2008. Negotiating with bounded rational agents in environments with incomplete information using an automated agent. *Artificial Intelligence* 172, 6-7 (2008), 823–851.
- [40] Noyda Matos, Carles Sierra, and Nicholas R Jennings. 1998. Determining successful negotiation strategies: An evolutionary approach. In *Proceedings International Conference on Multi Agent Systems (Cat. No. 98EX160)*. IEEE, Paris, France, 182–189.
- [41] Yasser Mohammad. 2020. NEGotiation MultiAgent System (NegMAS). <https://github.com/yasserfarouk/negmas> [Online; accessed 2020-09-21].
- [42] Howard Raiffa, John Richardson, David Metcalfe, et al. 2002. *Negotiation analysis: The science and art of collaborative decision making*. Harvard University Press, Harvard, US.
- [43] Yousef Razeghi, Celal Ozan Berk Yavuz, and REYHAN Aydoğan. 2020. Deep reinforcement learning for acceptance strategy in bilateral negotiations. *Turkish Journal of Electrical Engineering & Computer Sciences* 28, 4 (2020), 1824–1840.
- [44] Frederik Schadd, Sander Bakkes, and Pieter Spronck. 2007. Opponent Modeling in Real-Time Strategy Games.. In *GAMEON*. GAMEON, Bologna, Italy, 61–70.
- [45] Ayan Sengupta, Yasser Mohammad, and Shinji Nakadai. 2021. An Autonomous Negotiating Agent Framework with Reinforcement Learning Based Strategies and Adaptive Strategy Switching Mechanism. arXiv:2102.03588 [cs.AI]
- [46] Sergio Guadarrama and Anoop Korattikara and Et al. 2018. TF-Agents: A library for Reinforcement Learning in TensorFlow. <https://github.com/tensorflow/agents>. <https://github.com/tensorflow/agents> [Online; accessed 25-June-2019].
- [47] M. Sridharan and G. Tesauro. 2000. Multi-agent Q-learning and regression trees for automated pricing decisions. In *Proceedings Fourth International Conference on Multi-Agent Systems*. IEEE, Boston, MA, USA, 447–448.

- [48] Vishal Sunder, Lovekesh Vig, Arnab Chatterjee, and Gautam Shroff. 2018. Prosocial or Selfish? Agents with different behaviors for Contract Negotiation using Reinforcement Learning. arXiv:1809.07066 [cs.LG]
- [49] Gerald Tesauro. 2000. Pricing in agent economies using neural networks and multi-agent Q-learning. In *Sequence learning*. Springer, Berlin, Germany, 288–307.
- [50] Gerald Tesauro and Jeffrey O Kephart. 2002. Pricing in agent economies using multi-agent Q-learning. *Autonomous agents and multi-agent systems* 5, 3 (2002), 289–304.
- [51] Yasser Mohammad, Shinji Nakadai and Amy Greenwald. 2020. NegMAS: A platform for Situated Negotiations. *International Conference on Principles and Practice of Multi-Agent Systems* (2020).
- [52] Dajun Zeng and Katia Sycara. 1998. Bayesian learning in negotiation. *International Journal of Human-Computer Studies* 48, 1 (1998), 125–141.