Predicting Concession Curves of Negotiating Agents Using Machine Learning

Khalid Mansour
Faculty of Computer Studies, Arab Open University, Bahrain

***Abstract*—Accurate opponent modeling is critical for effective automated negotiation, enabling agents to adapt their strategies based on the type of opponent. This study investigates machine learning approaches for classifying negotiation agent strategies from offer sequences across three scenarios: time-dependent agents following predetermined concession functions, strategic agents adapting to opponent behavior with deadline-only termination, and strategic agents with realistic termination through mutual agreement or deadline expiration. We systematically evaluate four algorithms—Naive Bayes, Random Forest, Support Vector Machines, and Neural Networks— on a number of simulated negotiations, comparing classification performance with and without temporal feature augmentation. A key contribution of this work is the introduction of temporal feature augmentation, where quarterly concession patterns and variance metrics are used to capture adaptive negotiation behavior that raw offer sequences alone cannot reveal. The augmented features encode temporal adaptation characteristics that distinguish Boulware, Linear, Conceder, and strategic negotiation behaviors. Feature augmentation produced statistically significant improvements in 7 of 12 model–scenario combinations, with the most notable gains observed in strategic agent identification.**

***Keywords*—Automated negotiation; strategy classification; machine learning; feature engineering; strategic agents**

I. Introduction

Automated negotiation has emerged as a main component of modern multi-agent systems, enabling autonomous agents to reach agreements without human intervention [1]. As software agents increasingly mediate transactions in e-commerce, resource allocation, supply chain management, and collaborative decision-making systems, the ability to negotiate effectively becomes essential for optimizing outcomes and improving efficiency [2] [3]. The growing complexity of these systems, coupled with the need for real-time decision-making in dynamic environments, has driven substantial research interest in developing sophisticated automated negotiation mechanisms that can adapt to diverse scenarios and opponent behaviors [4] [5]. Automated negotiation is widely used in various domains such as online marketplaces, grid resource allocation, web service composition, and intelligent transportation systems, where autonomous agents act on behalf of their users to reach agreements that benefit all parties [6].

The success of automated negotiation heavily depends on an agent's ability to model and predict opponent behavior accurately [7]. Opponent modeling involves understanding the preferences, strategies, and future actions of negotiating parties, which enables agents to make informed decisions [8]. Traditional approaches to opponent modeling have relied on game-theoretic frameworks and heuristic methods that assume rational behavior and complete information [9]. However, these approaches often struggle to capture the complexity and variability of real-world negotiation scenarios, particularly when dealing with heterogeneous agents employing diverse strategies, incomplete information, and bounded rationality [10]. The limitations of classical game-theoretic approaches have motivated researchers to explore alternative methodologies that can learn from experience and adapt to unknown opponent characteristics.

A particularly important aspect of opponent modeling is predicting the concession behavior of negotiating agents [11]. Concession curves, which describe how agents modify their offers over time, provide crucial insights into an agent's negotiation strategy, urgency, and reservation values. Understanding concession patterns enables agents to anticipate opponent moves, identify optimal timing for offers, and avoid unnecessary concessions while maximizing agreement utility [12]. Faratin et al. [12] introduced the concept of time-dependent and behavior-dependent negotiation strategies, which have become foundational frameworks for automated negotiation research. Time-dependent strategies make concessions based on negotiation deadlines and remaining time, with agents becoming more conciliatory as deadlines approach. Behavior-dependent strategies, in contrast, adjust concessions based on opponent actions, such as matching opponent concessions or responding to patterns in offer sequences [12]. These two strategy types represent fundamental approaches to automated negotiation and exhibit distinct concession curve characteristics that can be exploited for prediction purposes.

The prediction of concession curves has important implications for negotiation strategy selection and adaptation. If an agent can accurately predict how an opponent will concede over time, it can optimize its own strategy to achieve better outcomes, such as delaying concessions when the opponent is expected to concede more rapidly or accelerating agreement when both parties are approaching their reservation values [13]. Furthermore, concession curve prediction can help identify negotiation deadlocks early, allowing agents to explore alternative proposals or adjust their aspiration levels to facilitate agreement [14]. Despite these potential benefits, concession curve prediction remains challenging due to the diversity of negotiation strategies, the presence of noise in observed behaviors, and the strategic considerations that may lead agents to disguise their true preferences.

Despite the growing body of research on opponent modeling, there remains a need for systematic evaluation of machine learning techniques specifically for predicting concession curves across different negotiation strategies [15]. While previous work has explored various aspects of opponent modeling, limited research has focused on comparing the effectiveness of different machine learning algorithms in capturing the temporal dynamics of concession behavior. The work in [7] provided a comprehensive survey of opponent modeling techniques, but noted that comparative evaluations across multiple learning methods remain scarce. Furthermore, most studies have concentrated on specific learning methods without providing comprehensive comparative analysis across multiple algorithms and strategy types. The lack of systematic comparison makes it difficult for negotiation agent designers to select appropriate learning methods for their specific applications and undermines our understanding of which algorithmic characteristics contribute to effective concession curve prediction.

While existing opponent modeling techniques primarily focus on time-dependent agents with fixed concession functions, the classification of adaptive strategic agents — whose behavior is driven by opponent response rather than a predetermined schedule — has received considerably less attention, and the role of temporal feature engineering in improving such classification remains largely unexplored.

This study addresses these gaps by evaluating four machine learning algorithms—Neural Networks, SVM, Random Forest, and Naïve Bayes—representing diverse learning paradigms, for classifying negotiation agent strategies under both time-dependent and adaptive behavior-dependent scenarios. A key contribution of this work is the introduction of temporal feature augmentation, where quarterly concession patterns and variance metrics are used to explicitly capture adaptive negotiation behavior that raw offer sequences alone cannot reveal. Results demonstrate that augmentation yields statistically significant improvements across most scenarios. These findings offer practical insights for designing more effective opponent models in automated negotiation systems

The remainder of this study is organized as follows: Section II reviews related work on automated negotiation and opponent modeling techniques. Section III describes the methodology, including the negotiation framework, agent strategies, dataset generation, feature engineering. Section IV presents the experimental settings, model configurations, and preprocessing and evaluation. Section V discusses the results and provides analysis of algorithm performance across different datasets and strategy types. Finally, Section VI concludes the study and outlines directions for future research.

II. Related Work

The task of predicting an opponent's concession curve sits at the intersection of automated negotiation strategy and machine learning-based opponent modeling. This section reviews the foundational work in these areas, beginning with classical negotiation tactics and opponent modeling techniques, then discussing data-driven and machine learning approaches, and finally identifying the specific gap that this research aims to address.

Recent advances in machine learning have opened new avenues for opponent modeling in automated negotiation [16]. Machine learning techniques offer the capability to learn patterns from historical negotiation data and make predictions about opponent behavior without strong assumptions about opponent rationality [17]. Several researchers have explored the application of various learning algorithms to predict opponent preferences, utilities, and negotiation strategies. Bayesian learning approaches have been employed to update beliefs about opponent preferences during negotiations, allowing agents to refine their models incrementally as new information becomes available [18]. The authors in [19] proposed a Bayesian learning method to estimate opponent preferences in multi-issue negotiations, demonstrating that learning-based approaches can significantly improve negotiation outcomes. Reinforcement learning methods have been used to develop adaptive negotiation strategies that improve through experience, enabling agents to learn optimal policies without prior knowledge of opponent behavior [20].

In addition to Bayesian and reinforcement learning approaches, other machine learning techniques have been applied to opponent modeling. The work in [21] investigated the use of kernel density estimation and regression techniques to learn opponent preferences from limited negotiation data. Their work highlighted the importance of sample efficiency in opponent modeling, as negotiations are typically conducted under time constraints with limited opportunities for information gathering. More recently, deep learning approaches have been explored for modeling complex negotiation dynamics, with neural networks demonstrating the ability to capture nonlinear relationships between negotiation features and opponent behavior [22]. The diversity of machine learning methods applied to opponent modeling reflects the multi-faceted nature of the challenge and the need for techniques that can handle different types of uncertainty and strategic complexity.

A. Classical Negotiation Tactics and Concession Strategies

The formalization of concession strategies is a cornerstone of automated negotiation. The seminal work of [12] introduced a taxonomy of negotiation tactics. Among these, time-dependent tactics are particularly relevant for modeling concession curves. These tactics define an agent's willingness to concede based on the remaining negotiation time, typically parameterized by a function (e.g., polynomial or exponential) and a β parameter that dictates the convexity of the concession curve. A tough negotiating agent would have $\beta < 1$ (Boulware strategy) offers small concessions until the deadline approaches then it concedes to its reservation value, with a $\beta > 1$, the agent is considered using a Conceder strategy which makes rapid early concessions. If $\beta = 1$, the agent offers linearly increasing amount of concessions and the strategy is called Linear. Infinite number of concession curves can be produced by changing the value of β. Behavior-dependent tactics, such as Tit-for-Tat, which mimic the opponent's previous concession, are also common.

B. Data-Driven and Machine Learning Approaches

The proliferation of large-scale negotiation tournaments, such as the Annual Automated Negotiating Agents Competition (ANAC) [23], has provided a rich source of data, catalyzing the application of machine learning (ML) to opponent modeling.

C. Machine Learning in Automated Negotiation

Machine learning techniques have become increasingly important in automated negotiation, enabling agents to learn effective strategies and adapt to diverse opponents. Recent research has focused primarily on reinforcement learning approaches, which allow negotiation agents to learn optimal policies through interaction with opponents or simulated environments.

Deep reinforcement learning (DRL) has emerged as a particularly promising approach for automated negotiation. Bagga et al. [24] developed ANEGMA, a negotiation model using an actor-critic architecture with model-free reinforcement learning for concurrent bilateral negotiations in e-markets. Their approach pre-trains strategies using supervised learning on synthetic market data, reducing exploration time during actual negotiations. Similarly, Chen et al. [25] proposed an effective negotiating agent framework based on deep offline reinforcement learning, addressing the challenge of expensive online data collection by learning from fixed datasets without active opponent interaction.

Deep reinforcement learning to multi-issue negotiation problems is applied showing that neural network-based approaches can handle the complexity of negotiations involving multiple interdependent issues [26].

Transfer learning and adaptive frameworks have also been investigated to enable agents to generalize across different negotiation scenarios. The authors in [27] developed an autonomous negotiating agent framework with reinforcement learning-based strategies and adaptive strategy switching mechanisms, allowing agents to select appropriate strategies based on opponent classification. This work is extended by [28] by incorporating emergent communication in coalitional negotiation games, demonstrating that deep reinforcement learning agents can learn to exchange information effectively to facilitate cooperation.

Despite these advances, most reinforcement learning approaches focus on learning strategies for specific negotiation domains or opponent types, with limited generalization to unseen scenarios. Additionally, the computational cost of training deep neural networks and the sample complexity of reinforcement learning remain challenges for practical deployment. Our work complements these approaches by investigating supervised learning methods for strategy classification, which can provide rapid identification of opponent behaviors without requiring extensive online training. In addition, we employed data augmentation to enhance performance.

III. METHODOLOGY

A. Agent-Based Negotiation Framework

The methodology follows an agent-based simulation framework in which negotiation data are generated and analyzed using supervised machine learning techniques. The experimental work is conducted in two main phases. In Phase 1, negotiation data are generated from different types of agents. In Phase 2, the generated data are used to train supervised machine learning models. The trained models are then evaluated based on their ability to accurately model and classify different agent types using the observed negotiation data. Regarding the complexity and scalability of machine learning algorithms used in this study, see Table I.

Table I summarizes the computational complexity and scalability of the four machine learning algorithms employed in this study, where n denotes the number of training sessions, r denotes the number of negotiation rounds per session, a denotes the number of augmented features, T denotes the number of trees in the ensemble, E denotes the number of training epochs, L denotes the number of layers, and h denotes the number of hidden units per layer. Without augmentation, the feature vector reduces to $2r$, comprising solely the raw offer and counter-offer values at each round, whereas augmentation expands the feature representation to $2r + a$, providing additional discriminative information that improves classification performance across most models.

TABLE I. COMPLEXITY AND SCALABILITY OF ML MODELS

Model	Complexity	Scalability
Naive Bayes	$\mathcal{O}(n(2r+a))$	Excellent
Random Forest	$\mathcal{O}(Tn(2r+a)\log n)$	Good
SVM	$\mathcal{O}(n^2)$–$\mathcal{O}(n^3)$	Poor
Neural Network	$\mathcal{O}(EnLh^2)$	Good

B. Agent Types and Behavioral Models

We implement six distinct negotiation agent types, divided into two categories: time-dependent agents and strategic agents. Time-dependent agents employ concession strategies based on temporal factors, while strategic agents adapt their behavior in response to opponent actions.

1) Time-dependent agents: The first category includes three agent types—Boulware, Linear, and Conceder—which generate offers according to the time-dependent tactical function, see Eq. (1) and Eq. (2):

$$\alpha_l(t) = \left(\frac{t}{T}\right)^{1/\beta_l} \tag{1}$$

In Eq. (1), $\alpha_l(t)$ refers to the time-dependent scaling factor of agent l at time t, T is the deadline of agent l and β_l is the convexity of the concession curve of agent l.

$$x_l(t) = IV_l + \alpha_l(t)\big(FV_l - IV_l\big) \tag{2}$$

$$x_l(t) = IV_l + (1 - \alpha_l(t))\big(FV_l - IV_l\big) \tag{3}$$

Eq. (2) and Eq. (3) define the offer calculation, where IV_l and FV_l are the initial and final (reservation) values of agent l, and $x_l(t)$ denotes the offer at time t. Eq. (2) is applied when the offer increases over time, while Eq. (3) is used when the offer decreases.

In this study, we consider a bilateral negotiation scenario between a buyer agent and a seller agent over a single price issue. The buyer agent increases its offer at each time step according to Eq. (2) until reaching its reservation value, while the seller agent decreases its offer at each time step according to Eq. (3) until reaching its reservation value. For example, given a reservation interval of $[10, 20]$, the buyer agent offers 10 at $t = 0$, while the seller agent offers 20. As time progresses, the buyer's offer increases above 10, and the seller's offer decreases below 20, potentially leading to an agreement. For the remainder of this study, $x_i(t)$ and $x_j(t)$ denote the buyer's offer and the seller's offer at time t, respectively.

$$x_j(t) = x_j(t-1) - \rho\,(x_i(t-1) - x_i(t-2)) \tag{4}$$

Eq. (4) (with $\rho \in [0, 1]$) is employed by strategic seller agents who respond to various time-dependent buyer strategies, including Boulware, linear, and conceder. In this equation, $x_i(t-1)$ and $x_i(t-2)$ denote the two most recent offers made by the buyer agent. Seller agents using Eq. (4) are categorized as Strategic Boulware ($Stra\_Bou$), Strategic Linear ($Stra\_Lin$), or Strategic Conceder ($Stra\_Con$), depending on the type of buyer agent they are negotiating with. For example, when a seller agent negotiates with a Boulware buyer agent, the seller agent is referred to as $Stra\_Bou$. In our experiments, the data used to train the machine learning models consisted of offers generated by seller agents employing various negotiation strategies.

C. Dataset Generation

We generated three distinct negotiation datasets to train and evaluate machine learning models for automated negotiation strategy classification. Each dataset simulates bilateral negotiations between buyer and seller agents with varying behavioral patterns and termination conditions.

1) Case 1: Time-dependent sellers: This dataset contains 3,000 seller offer sequences (1,000 per strategy type) generated by time-dependent seller agents negotiating in isolation. Each agent is assigned a random deadline $T \in [10, 15]$ rounds and a β value sampled uniformly from its strategy's range. Agents generate offers from $t = 0$ to $t = T$, producing sequences that capture pure time-dependent negotiation patterns without opponent interaction. The interval of β determines the type of agent: *Boulware*: (0.2, 0.9), *Linear*:(0.9, 1.2), *Conceder*: (1.2, 4.0). This case serves as a baseline for evaluating how well machine learning models can distinguish fundamental negotiation strategies based solely on offer progression patterns.

2) Case 2: Strategic sellers with deadline-only termination: This dataset simulates 3,000 bilateral negotiations (1,000 per seller agent type) between strategic sellers and time-dependent buyer agents. Both agents are assigned randomly sampled deadlines in the range [10,15] and negotiate in an alternating-offer protocol until one of the agents reaches its deadline, whichever occurs first. In this setting, negotiations proceed until a deadline is reached even if a mutually acceptable agreement emerges earlier, thereby capturing the full negotiation trajectory up to the deadline.

Each negotiation proceeds as follows:

- At $t = 0$, the seller offers $IV_j = 20$.
- For $t \geq 1$, the buyer generates a time-dependent offer $x_i(t)$ and the seller responds strategically based on the buyer's concession magnitude using equation 4.
- The process continues until $t = \min(T_i, T_j)$.

This case evaluates whether machine learning models can identify agents' strategic adaptation patterns as they dynamically respond to opponent behavior throughout the entire negotiation horizon.

3) Case 3: Strategic sellers with realistic termination: This dataset extends Case 2 by incorporating realistic negotiation termination conditions. Negotiations terminate when either: 1) an agent reaches its deadline, or 2) one agent accepts the opponent's offer. An agent accepts an offer when the opponent's current offer is at least as favorable as the agent's next planned offer.

This acceptance mechanism produces variable-length sequences that reflect real-world negotiations, where agreements are often reached before deadlines.

This case represents the most realistic scenario and tests whether models can classify strategic behaviors even when negotiations terminate at different stages, producing sequences of varying lengths.

D. Data Characteristics

Table II presents summary statistics for all three cases. Each dataset maintains class balance with exactly 1,000 sequences per strategy type. Sequence lengths vary by case: Case 1 exhibits uniform length distributions within the deadline range [10, 15], while Case 2 sequences reach the full deadline length. Case 3 shows the greatest length variability due to early termination through acceptance, with sequences ranging from as few as 5 rounds to 15-rounds.

TABLE II. DATASET SUMMARY STATISTICS

Characteristic	Case 1	Case 2	Case 3
Total sequences	3,000	3,000	3,000
Sequences per class	1,000	1,000	1,000
Number of classes	3	3	3
Sequence length (min)	10	10	5
Sequence length (max)	15	15	15
Sequence length (mean)	12.0	12.0	9.2
Strategy labels	Boulware, Linear, Conceder	Stra_Boul, Stra_Lin, Stra_Con	Stra_Boul, Stra_Lin, Stra_Con

All datasets employ stratified train-test splits (80/20) to ensure balanced representation across classes during model evaluation.

To maintain dataset balance and ensure consistency across different sequence lengths, all offer sequences are truncated to the shortest observed length.

E. Feature Augmentation

To enhance classification performance, we augment the raw offer sequences with six derived features that capture temporal concession patterns and underlying negotiation dynamics. Let $x_{s_j} = \{x_j(0), x_j(1), \ldots, x_j(n)\}$ denote the offer sequence of a seller agent s_j with length n, where $x_j(t)$ represents the offer made by agent s_j at time t. The augmented features are defined as follows:

1) Quarterly concession features: We partitioned each negotiation into four temporal quarters and computed the concession magnitude within each period [see Eq. (5) to Eq. (8)]:

$$Q_1 = x_j(0) - x_j(\lfloor n/4 \rfloor) \tag{5}$$
$$Q_2 = x_j(\lfloor n/4 \rfloor) - x_j(\lfloor n/2 \rfloor) \tag{6}$$
$$Q_3 = x_j(\lfloor n/2 \rfloor) - x_j(\lfloor 3n/4 \rfloor) \tag{7}$$
$$Q_4 = x_j(\lfloor 3n/4 \rfloor) - x_j(n) \tag{8}$$

These features capture the temporal distribution of concessions, enabling discrimination between Boulware agents (small Q_1, large Q_4), Linear agents (uniform Q_i), and Conceder agents (large Q_1, small Q_4).

2) Variance features: To measure concession consistency and smoothness, we computed variance statistics on the first-order difference of the offer sequence [see Eq. (9)]:

$$\begin{aligned} \mathrm{Var}_{\mathrm{1st}} &= \mathrm{Var}\Big(\{\Delta x_j(t)\}_{t=0}^{T_{s_j}}\Big), \\ \mathrm{Var}_{\mathrm{2nd}} &= \mathrm{Var}\Big(\{\Delta^2 x_j(t)\}_{t=0}^{T_{s_j}-1}\Big), \end{aligned} \tag{9}$$

where,

$$\begin{aligned} \Delta x_j(t) &= x_j(t) - x_j(t+1), \\ \Delta^2 x_j(t) &= \Delta x_j(t) - \Delta x_j(t+1). \end{aligned}$$

The first-order difference $[\Delta x_j(t)]$ shows the magnitude of concessions between consecutive offers, while the second-order difference $[\Delta^2 x_j(t)]$ reflects variations in concession behavior over time, indicating whether an agent accelerates or decelerates its concessions. In contrast, the variance of first-order differences ($\mathrm{Var}_{\mathrm{1st}}$) captures the consistency of concession magnitudes across negotiation rounds, while the variance of second-order differences ($\mathrm{Var}_{\mathrm{2nd}}$) reflects the smoothness of concession dynamics over time.

3) Feature vector construction: The final augmented feature vector for each sequence is constructed as in Eq. (10):

$$\mathbf{x}_{\text{aug}} = [x_j(0), x_j(1), \ldots, x_j(n), Q_1, Q_2, Q_3, Q_4, \mathrm{Var}_{\mathrm{1st}}, \mathrm{Var}_{\mathrm{2nd}}] \tag{10}$$

where, $\mathbf{x}_{\text{aug}} \in \mathbb{R}^{n+6}$. This representation preserves the complete offer sequence while adding explicit temporal and statistical features that facilitate classification by machine learning models.

IV. Experimental Settings

The experimental dataset comprises 6,000 negotiation agents evenly distributed across six strategy types: Boulware, Linear, Conceder, strategic Boulware (Stra_Bou), strategic Linear (Stra_Lin), and strategic Conceder (Stra_Con), with exactly 1,000 instances per strategy. Each agent participates in a bilateral negotiation scenario with an initial offer value of 20 and a reservation price of 10 for the seller agent, with the corresponding values reversed for the buyer agent. Negotiation deadlines are randomly assigned from a discrete uniform distribution over $[10, 15]$, resulting in variable-length offer sequences. In this study, we classify seller agent types; therefore, only the offers generated by the seller agent are collected to train the machine learning models.

The concession parameter (β) for the time-dependent agents used in our experiments is defined as follows:

- $\beta \in [0.2, 0.9]$ for the *Boulware* strategy, which exhibits stubborn behavior with minimal early concessions;
- $\beta \in [0.9, 1.2]$ for the *linear* strategy, characterized by steady concessions;
- $\beta \in [1.2, 4.0]$ for the *conceder* strategy, which makes rapid concessions early in the negotiation.

Following sequence generation, all offers are truncated to the minimum observed length to maintain uniform dimensionality across data points.

A. Model Configurations

1) Neural network architecture: The MLP classifier consists of an input layer matching the input dimensionality, followed by two densely connected hidden layers with 100 and 75 neurons respectively, and a 3-class softmax output layer. The network employs ReLU activation functions in hidden layers and is trained for a maximum of 500 iterations with mini-batch gradient descent using the Adam optimizer. Training incorporates early stopping with a validation split of 10%, monitoring validation loss over 10 consecutive iterations with no improvement before stopping. The Adam optimizer is configured with a learning rate $\eta = 0.001$, momentum coefficients $\beta_1 = 0.9$ and $\beta_2 = 0.999$, an L_2 penalty $\alpha = 0.0001$, and a convergence tolerance $\varepsilon = 10^{-8}$.

2) Random forest configuration: The ensemble comprises 100 decision trees constructed using bootstrap sampling with no maximum depth constraint, allowing trees to expand until all leaves are pure or contain fewer than the minimum samples required for splitting. Each split requires a minimum of 2 samples, while leaf nodes require at least 1 sample. The Gini impurity criterion guides feature selection at each node, and bootstrap aggregation is enabled with out-of-bag error estimation for internal validation.

3) SVM configuration: The support vector classifier employs an RBF kernel with regularization parameter C = 1.0 and automatic gamma computation based on feature variance ($gamma =' scale'$, which equals $1/(n\_features * variance)$). Class weights are balanced inversely proportional to class frequencies to manage potential class imbalance risks.

4) Naive Bayes configuration: The Gaussian variant assumes feature independence with individual feature variance estimations. A variance smoothing parameter of 10^{-9} is applied to all features to prevent numerical underflow and ensure computational stability.

5) Training configuration: All models are trained using an 80/20 train-test split with stratified sampling to maintain class distribution. Features are standardized using z-score normalization (StandardScaler) with zero mean and unit variance.

Random state is fixed at 42 for reproducibility across all experiments.

B. Preprocessing and Evaluation

All features undergo standardization via z-score normalization, transforming each feature to zero mean and unit variance based on training set statistics. Label encoding converts categorical strategy names to integer indices $\{0, 1, 2\}$ corresponding to alphabetically ordered class names (e.g., for Case 1: Boulware=0, Conceder=1, Linear=2; for Cases 2 and 3: Stra_Boul=0, Stra_Con=1, Stra_Lin=2).

The dataset is partitioned using stratified random sampling with an 80/20 train-test split (random seed: 42), ensuring proportional class representation in both subsets. This yields 2,400 training samples and 600 test samples, with 800 and 200 instances per class, respectively.

Model performance is assessed using accuracy, $accuracy = \frac{TP+TN}{TP+TN+FP+FN}$.

V. Results and Discussion

This section presents the experimental results and discussion, followed by a presentation of the research implications and limitations.

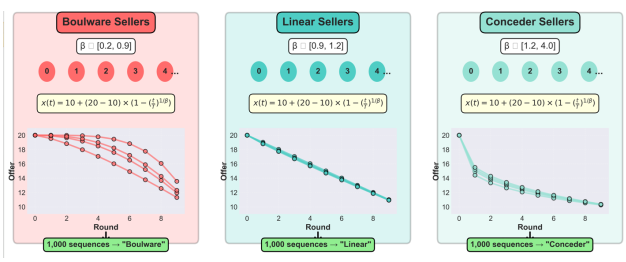

Fig. 1. Sample offer sequences by seller agents: Time-dependent agents.

Concession curves for time-dependent seller agents in Case 1 of the data generation are depicted in Fig. 1.

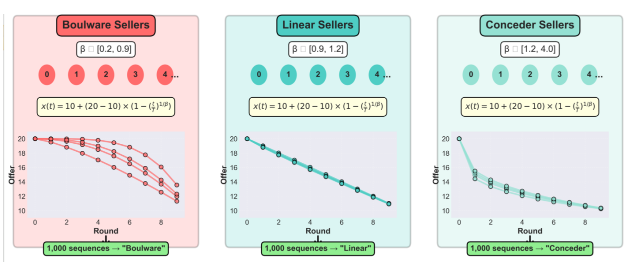

Fig. 2. Sample offer sequences of strategic seller agents, showing negotiations that stop at deadlines.

Fig. 2 illustrates the concession curves of strategic seller agents for Case 2 of the data generation process where negotiation stops at either the buyer agent's deadline or at the seller agent's deadline.

Fig. 3 shows the concession curves of strategic seller agents in Case 3, where negotiations conclude at an agent's deadline or upon reaching agreement.

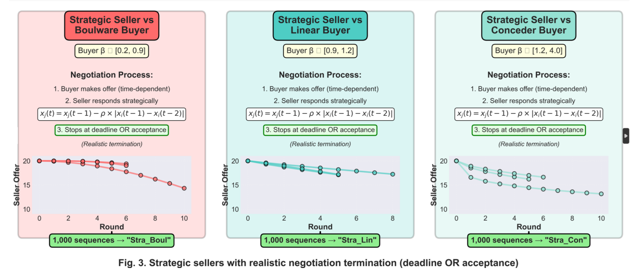

Fig. 3. Sample offer sequences of strategic seller agents, with negotiations ending at deadlines or upon reaching agreement.

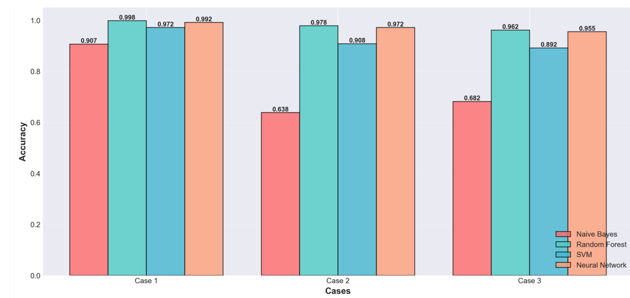

Fig. 4. ML algorithm accuracy without feature augmentation.

Fig. 4 compares the accuracy of four machine learning algorithms without feature augmentation. Random Forest and Neural Network achieve the highest accuracy, whereas Naive Bayes performs the lowest. SVM performs significantly better in Case 1 compared to Cases 2 and 3.

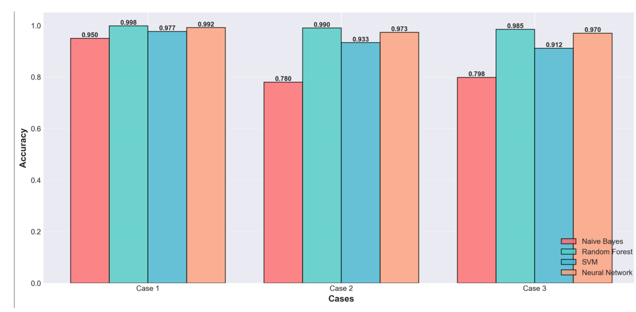

Fig. 5. ML algorithm accuracy with feature augmentation.

Fig. 5 compares the accuracy of four machine learning algorithms with feature augmentation. All ML algorithms show improvement in both Case 2 and Case 2. Random Forest and Neural Network achieve the highest accuracy, whereas Naive Bayes performs the lowest. SVM performs significantly better in Case 1 compared to Cases 2 and 3.

The performance improvements observed in Cases 2 and 3 (strategic agents) compared to Case 1 (time-dependent agents) stem from the fundamental complexity of strategic negotiation behavior. While time-dependent agents follow deterministic mathematical functions that machine learning models can readily extract from raw offer sequences, strategic seller agents adapt dynamically to opponent behavior producing variable, opponent-dependent trajectories with complex temporal patterns. Table III presents the accuracies across all cases.

TABLE III. Classification Accuracy Comparison: Without vs. With Feature Augmentation.

Model	Case 1		Case 2		Case 3	
	W/O	W/	W/O	W	W/O	W
Naive Bayes	0.907	0.950	0.638	0.780	0.682	0.798
Random Forest	0.998	0.998	0.978	0.990	0.962	0.985
SVM	0.972	0.977	0.908	0.933	0.892	0.912
Neural Network	0.992	0.992	0.972	0.973	0.955	0.970

To further highlight the impact of feature augmentation, Fig. 6 illustrates the performance differences between models trained with and without feature augmentation across all cases. Naive Bayes consistently demonstrates the largest improvements (+4.33%, +14.17%, +11.67%), confirming that explicit feature engineering is essential for models with limited capacity to learn complex temporal patterns. The dramatic improvements in Cases 2 and 3 (strategic agents) compared to Case 1 (time-dependent agents) validate that the augmented features—quarterly concessions and variance measures—effectively capture the temporal adaptation patterns inherent in strategic negotiation behavior. Random Forest and SVM show moderate, consistent improvements across strategic cases (+1.17%–2.50%), while Neural Networks exhibit minimal gains (+0.00%–1.50%), suggesting these models already learn similar feature representations internally. Notably, all models either improve or maintain performance with augmentation, with zero cases of degradation, demonstrating the robustness of the feature engineering approach. The dramatic improvements in Cases 2 and 3 compared to Case 1 validate that explicit temporal features become increasingly valuable as negotiation complexity increases from deterministic time-dependent behaviors to adaptive strategic interactions.

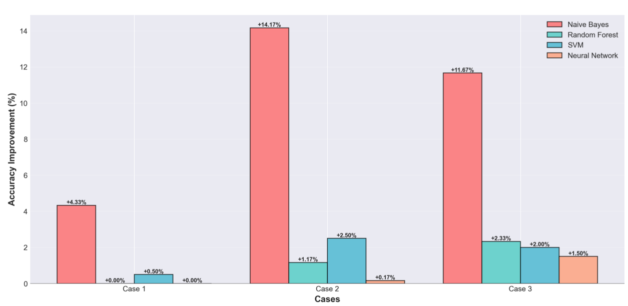

Fig. 6. Accuracy improvement from feature augmentation.

Feature augmentation provides an overall +3.3 percentage point improvement across all models and cases, with strategic agents (Cases 2 and 3) benefiting nearly 4× more than time-dependent agents (Case 1). Across all 12 model-case combinations, 7 showed statistically significant improvements ($p < 0.05$), with no cases of performance degradation. Naive Bayes exhibited highly significant gains in all three cases ($p < 0.001$). Random Forest and SVM showed significant improvements primarily in strategic agent scenarios (Cases 2 and 3), with Random Forest gaining +1.2% (p = 0.022) and +2.3% ($p < 0.001$), and SVM improving +2.5% (p = 0.023) in Case 2. Neural Networks showed minimal but statistically significant improvement only in Case 3 (+1.5%, p = 0.041), suggesting that deep learning models already approximate the augmented features through their learned representations.

Feature importance was assessed using Random Forest's built-in Gini importance metric, which measures the average reduction in impurity across all trees. This approach was chosen due to its computational efficiency and interpretability.

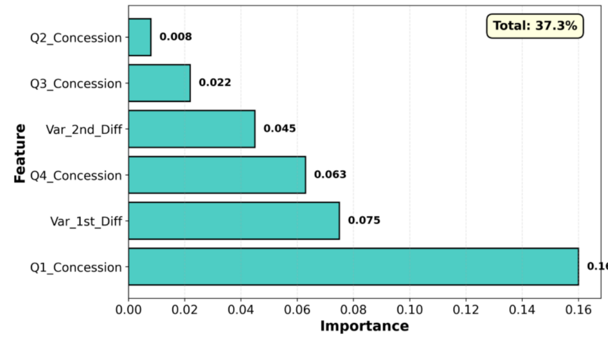

Fig. 7. Augmented feature importance for Case 1.

Fig. 7 shows the feature importance for Case 1. Q1 concession shows the most important feature, while Q2 shows the least important feature.

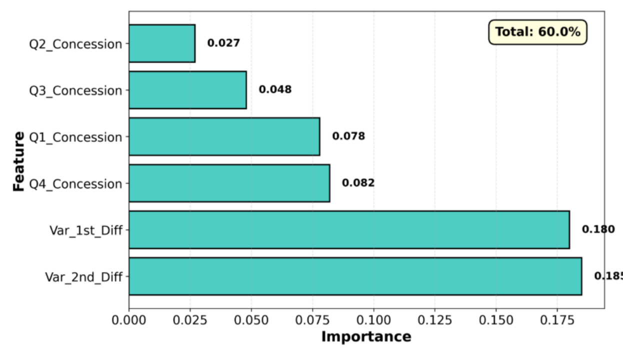

Fig. 8. Augmented feature importance for Case 2.

Fig. 8 shows the feature importance for Case 2. The variance of the second order difference was the best performer while Q2 shows the least important feature.

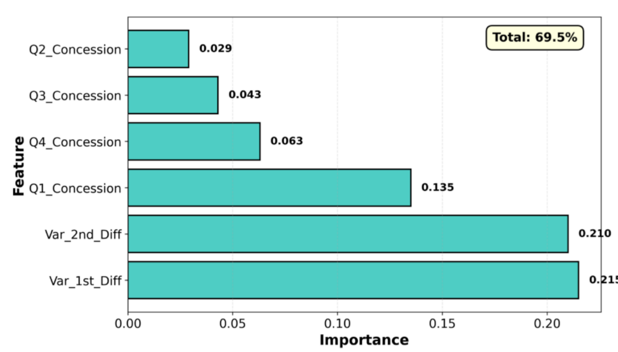

Fig. 9. Augmented feature importance for Case 3.

Fig. 9 shows the feature importance for Case 3. Again, the variance of the second order difference was the best performer while Q2 shows the least important feature.

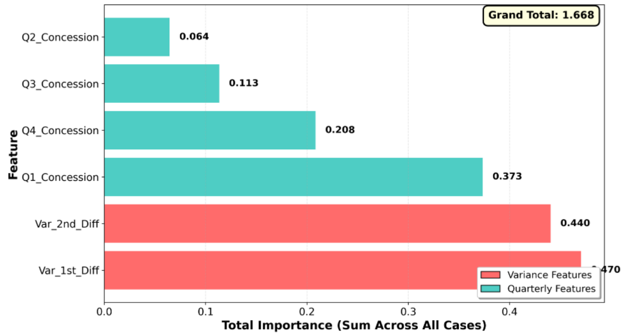

Fig. 10. Augmented feature importance: Across all cases.

Feature importance analysis using Random Forest's Gini importance metric reveals that the augmented features contribute substantially to classification performance, with their relative importance varying systematically across negotiation scenarios. Aggregating importance values across all three cases, variance features (Var_1st_Diff: 0.470, Var_2nd_Diff: 0.440) collectively account for 54.6% of total augmented feature importance, while quarterly concession features comprise 45.4%, see Fig. 10. The total contribution of augmented features increases progressively with scenario complexity: 37.3% for time-dependent agents (Case 1), 60.0% for strategic agents with deadline termination (Case 2), and 69.5% for strategic agents with realistic termination (Case 3). Notably, feature importance patterns shift dramatically between time-dependent and strategic scenarios. In Case 1, Q1_Concession dominates (16.0% importance), reflecting that early-round behavior strongly predicts time-dependent strategy types. Conversely, in strategic cases (2 and 3), variance features become paramount (combined 60.8% and 61.2% of augmented importance respectively), indicating that consistency and volatility patterns in concession rates are critical discriminators for strategic agent classification. This pattern confirms that strategic agents, which adapt dynamically to opponent behavior, require fundamentally different feature representations focused on behavioral consistency rather than temporal position, validating our feature engineering approach for opponent modeling in automated negotiation.

A. Implications for Automated Negotiation

The high prediction accuracy achieved by our approach has important implications for automated negotiation systems. Agents equipped with accurate strategy identification capabilities can dynamically adapt their counter-strategies rather than relying on static negotiation protocols. For instance, recognizing a Boulware opponent who maintains high offers until near deadline allows an agent to adopt patient negotiation tactics, while identifying a Conceder opponent who makes early concessions enables more aggressive value claiming. The effectiveness of feature-augmented Naive Bayes and SVM (achieving 78-93% accuracy) demonstrates that even resource-constrained environments can deploy effective opponent modeling without requiring computationally expensive deep learning infrastructure.

B. Limitations

While our results demonstrate strong classification performance, several limitations warrant acknowledgment. First, our study examines six specific strategy types derived from classical time-dependent and strategic negotiation tactics. Real-world negotiations often involve hybrid strategies that blend multiple tactics, time-varying behavior where agents switch strategies mid-negotiation, or deceptive practices where agents deliberately misrepresent their intentions. Second, our experimental design focuses on single-issue bilateral negotiations (price negotiation between buyer and seller). Multi-issue negotiations with interdependent preferences, trade-offs across issues require fundamentally different feature representations and classification approaches. Third, our models classify complete offer sequences post-negotiation rather than partial histories during ongoing negotiations. For real-time opponent modeling, incremental classification methods that can identify strategies from incomplete sequences would be necessary, potentially using sliding window approaches or recurrent architectures. Fourth, our datasets assume perfect information about opponent offers; realistic scenarios may involve noisy observations, missing data, or communication delays.

VI. Conclusion and Future Work

This study investigated machine learning approaches for classifying negotiation agent strategies from offer sequences, addressing a fundamental challenge in opponent modeling for automated negotiation systems. We systematically evaluated four algorithms —Naive Bayes, Random Forest, SVM, and Neural Networks— across three negotiation scenarios of increasing complexity: 1) time-dependent sellers following predetermined concession functions, 2) strategic sellers adapting to opponent behavior with deadline-only termination, and 3) strategic sellers with realistic termination through mutual agreement or reaching the deadline. Our experimental evaluation using 9,000 simulated negotiations (3,000 per case) examined both raw offer sequences and augmented feature representations incorporating quarterly concession patterns and variance metrics.

Our experimental results yield several significant insights. First, feature augmentation consistently improves classification accuracy across all tested model types, with 7 out of 12 model-case combinations showing statistically significant gains. Naive Bayes demonstrated the most substantial improvements, while Neural Networks showed minimal gains, confirming that explicit temporal feature engineering benefits models with limited capacity to learn complex patterns while providing diminishing returns for deep learning approaches capable of automatic feature discovery.

Second, strategic agent classification (Cases 2 and 3) benefited substantially more from feature augmentation than time-dependent agents. This validates that the augmented features—quarterly concessions and variance measures—effectively capture the temporal adaptation patterns inherent in strategic negotiation behavior, which are difficult to learn from raw sequences alone.

Third, achieving high classification accuracy proved feasible across all scenarios, with Random Forest reaching 99.0%

accuracy in Case 2 and 98.5% in Case 3 with feature augmentation. Even simple probabilistic models like Naive Bayes achieved 78.0% and 79.8% accuracy in strategic scenarios after augmentation, demonstrating that appropriate feature engineering enables competitive performance without requiring complex architectures.

Finally, the augmented features contribute substantially to classification performance, with their importance increasing from time-dependent to strategic negotiation scenarios. While early concession features are most predictive for time-dependent agents, variance-based features dominate in strategic cases, indicating that behavioral stability and volatility are key discriminators for adaptive agents.

A. Future Work

Building on this foundation, we propose several specific research directions:

- Online learning for partial sequences: Develop methods that classify strategies from incomplete negotiation histories, enabling real-time adaptation during active negotiations.
- Multi-issue negotiation: Extend the feature engineering framework to capture interdependent preferences and issue trade-offs.
- Strategy switching and deception detection: Investigate techniques for identifying when opponents change strategies mid-negotiation or deliberately misrepresent their intentions to gain strategic advantage.

Acknowledgment

This research was supported by the Arab Open University, Bahrain.

References

[1] N. R. Jennings, P. Faratin, A. R. Lomuscio, S. Parsons, M. J. Wooldridge, and C. Sierra, "Automated negotiation: prospects, methods and challenges," *Group Decision and Negotiation*, vol. 10, no. 2, pp. 199–215, 2001.

[2] S. Fatima, S. Kraus, and M. Wooldridge, "Principles of automated negotiation," *Cambridge University Press*, 2014.

[3] K. Mansour, Y. Al-Lahham, S. Tawil, R. Kowalczyk, and A. Al-Qerem, "An effective negotiation strategy for quantitative and qualitative issues in multi-agent systems," *Electronics*, vol. 11, no. 17, 2022. [Online]. Available: https://www.mdpi.com/2079-9292/11/17/2754

[4] R. Lin, S. Kraus, T. Baarslag, D. Tykhonov, K. Hindriks, and C. M. Jonker, "Designing and evaluating novel trading agent competition mechanisms," *Artificial Intelligence*, vol. 217, pp. 1–44, 2014.

[5] K. Mansour, "A hybrid concession mechanism for negotiating software agents in competitive environments," *International Journal on Artificial Intelligence Tools*, vol. 29, no. 06, p. 2050016, 2020. [Online]. Available: https://doi.org/10.1142/S0218213020500165

[6] C. M. Jonker, R. Aydogan, T. Baarslag, K. Fujita, T. Ito, and K. Hindriks, "Automated negotiating agents: From theory to practice," in *Proceedings of the 16th Conference on Autonomous Agents and MultiAgent Systems*, 2017, pp. 1233–1237.

[7] T. Baarslag, M. J. Hendrikx, K. V. Hindriks, and C. M. Jonker, "Learning about the opponent in automated bilateral negotiation: a comprehensive survey of opponent modeling techniques," *Autonomous Agents and Multi-Agent Systems*, vol. 30, no. 5, pp. 849–898, 2017.

[8] C. Hou, "Predicting the performance of opponent models in automated negotiation," in *2014 IEEE/WIC/ACM International Joint Conferences on Web Intelligence (WI) and Intelligent Agent Technologies (IAT)*, vol. 3. IEEE, 2014, pp. 355–361.

[9] A. Muthoo, "Bargaining theory with applications," *Cambridge University Press*, 1999.

[10] A. Rubinstein, "Perfect equilibrium in a bargaining model," *Econometrica: Journal of the Econometric Society*, pp. 97–109, 1982.

[11] M. A. Lopez-Carmona, I. Marsa-Maestre, E. de la Hoz, and J. R. Velasco, "Learning in negotiation: An evolutionary perspective," *Applied Soft Computing*, vol. 13, no. 1, pp. 318–331, 2013.

[12] P. Faratin, C. Sierra, and N. R. Jennings, "Negotiation decision functions for autonomous agents," *Robotics and Autonomous Systems*, vol. 24, no. 3, pp. 159–182, 1998, multi-Agent Rationality. [Online]. Available: https://www.sciencedirect.com/science/article/pii/S0921889098000293

[13] T. Baarslag, K. Hindriks, C. Jonker, S. Kraus, and R. Lin, "Predicting the performance of opponent models in automated negotiation," in *Proceedings of the 2013 international conference on Autonomous agents and multi-agent systems*, 2013, pp. 59–66.

[14] S. Chen and G. Weiss, "Strategy selection in automated negotiation," *Multiagent and Grid Systems*, vol. 9, no. 2, pp. 121–142, 2013.

[15] R. Aydogan, D. Festen, K. V. Hindriks, and C. M. Jonker, "A heuristic negotiation framework for preference elicitation in multi-criteria negotiations," *Applied Intelligence*, vol. 47, no. 2, pp. 502–524, 2017.

[16] S. Kraus, "Strategic negotiation in multiagent environments," *MIT press*, 2001.

[17] J. Bakker, T. Baarslag, K. Hindricks, and C. M. Jonker, "Preference elicitation and modelling in social choice," *Trends in Computational Social Choice*, pp. 137–158, 2017.

[18] K. Hindriks and D. Tykhonov, "Opponent modelling in automated multi-issue negotiation using bayesian learning," in *Proceedings of the 7th international joint conference on Autonomous agents and multiagent systems*, vol. 1, 2008, pp. 331–338.

[19] K. V. Hindriks and D. Tykhonov, "Learning opponent's preferences in multi-attribute negotiation using neural networks," *Electronic Commerce Research and Applications*, vol. 9, no. 4, pp. 274–284, 2010.

[20] K. Fujita, T. Ito, and M. Klein, "Automated strategy adaptation for multi-times bilateral closed negotiations," in *Proceedings of the 2011 IEEE/WIC/ACM International Conferences on Web Intelligence and Intelligent Agent Technology*, vol. 2. IEEE, 2011, pp. 3–7.

[21] R. M. Coehoorn and N. R. Jennings, "Learning opponent's preferences in multi-issue negotiation," in *Proceedings of the fourth international joint conference on Autonomous agents and multiagent systems*, 2005, pp. 729–736.

[22] M. He, N. R. Jennings, and A. Prugel-Bennett, "Learning and optimizing through multi-agent bargaining," *IEEE Transactions on Cybernetics*, vol. 48, no. 5, pp. 1479–1493, 2016.

[23] T. Baarslag, K. Hindriks, C. Jonker, S. Kraus, and R. Lin, "The first automated negotiating agents competition (anac 2010)," in *New Trends in Agent-Based Complex Automated Negotiations*, T. Ito, M. Zhang, V. Robu, S. Fatima, and T. Matsuo, Eds. Berlin, Heidelberg: Springer Berlin Heidelberg, 2012, pp. 113–135.

[24] P. Bagga, N. Paoletti, B. Alrayes, and K. Stathis, "Anegma: An automated negotiation model for e-markets," *Autonomous Agents and Multi-Agent Systems*, vol. 35, no. 2, p. 27, 2021.

[25] S. Chen, J. Zhao, G. Weiss, R. Su, and K. Lei, "An effective negotiating agent framework based on deep offline reinforcement learning," in *Proceedings of the Thirty-Ninth Conference on Uncertainty in Artificial Intelligence*. PMLR, 2023, pp. 324–335.

[26] H.-C. H. Chang, "Multi-issue negotiation with deep reinforcement learning," *Knowledge-Based Systems*, vol. 211, p. 106544, 2021.

[27] A. Sengupta, Y. Mohammad, and S. Nakadai, "An autonomous negotiating agent framework with reinforcement learning based strategies and adaptive strategy switching mechanism," in *Proceedings of the 20th International Conference on Autonomous Agents and MultiAgent Systems*, 2021, pp. 1163–1172.

[28] S. Chen, Y. Yang, and R. Su, "Deep reinforcement learning with emergent communication for coalitional negotiation games," *Mathematical Biosciences and Engineering*, vol. 19, no. 5, pp. 4592–4609, 2022.