

Design and Comparison of Arabic Negotiation Bots Using LLMs versus Seq2Seq Models with Reinforcement Learning

Ahmad Hajj, Yasmine Abu Adla, Samah Albast, Hazem Hajj, Shady Elbassuoni,
Wassim El Hajj, Khaled Shaban

University of Wisconsin–Madison, US; American University of Beirut, Lebanon;
Qatar University, Qatar

ahajj@wisc.edu, {yaa41, sma122, hh63, se58, we07}@aub.edu.lb, khaled.shaban@qu.edu.qa

Abstract

Negotiation is a crucial aspect of daily life, spanning from personal agreements to organizational contracts. As AI continues to evolve, the delegation of negotiation tasks to machines is becoming increasingly feasible. However, while research on English negotiation bots is progressing, similar advancements in other languages, particularly Arabic, are lacking. This paper presents the first design of Arabic negotiation bots, employing two approaches. The first approach involves utilizing a pre-trained Large Language Model (LLM) for this task via prompt engineering, while the second leverages Meta's "Deal or No Deal" framework, integrating a sequence-to-sequence (Seq2Seq) model with reinforcement learning (RNN-RL). We introduce and release two new datasets tailored for the models, including LLM prompt instructions and an Arabic dataset for the RNN-RL model. Our experiments demonstrate both the potential and limitations of these approaches, providing essential insights into their performance and outlining future research directions for multilingual negotiation bots.

1 Introduction

Negotiations occur in various contexts, from individuals haggling over a purchase to companies discussing contracts and governments seeking economic advantages (Halver, 2022; Founders, 2022). With the recent advancements in chatbots, negotiation has emerged as an ideal application for artificial intelligence (AI), enabling intelligent agents to negotiate and cooperate effectively.

Research on English negotiation bots has evolved, initially employing game theory and rule-based approaches (Jennings et al., 2001; Hussain, 2014; Balachandran and Mohammadian, 2015; Ko-

ley and Rao, 2018), followed by deep learning methods that achieved significant success in natural language processing. Meta's work (Lewis et al., 2017) laid the foundation for product negotiation using two sequence-to-sequence (Seq2Seq) recurrent neural networks (RNN) models representing the buyer and the seller. Those models were trained in a supervised fashion using human-developed negotiation data, with parameters further improved through reinforcement learning (RL). Building upon Meta's RNN-RL models, several other works (He et al., 2018; Wang et al., 2019; Zhou et al., 2019; Mishra et al., 2022; Raut et al., 2023) introduced further improvements, including strategies for emotion, persuasion, and politeness. Furthermore, transformer-based sequence models like generative pre-training transformer (GPT) models for negotiation bots have been developed (Fu et al., 2023). While those English negotiation bots have seen continued advancements, research on non-English negotiation bots, including Arabic, remains unexplored.

To address this research gap, we introduce two approaches to automate non-English negotiations with application to Arabic and create, to the best of our knowledge, the first Arabic negotiation bots. In our first approach, we propose utilizing a pre-trained large language model (LLM), namely GPT-3.5 Turbo, to automate Arabic negotiations through prompt engineering. Moreover, as a baseline, we develop an Arabic negotiation model using a combination of a Seq2Seq RNN with reinforcement learning (RNN-RL), which is inspired by Meta's popular "Deal or No Deal" model (Lewis et al., 2017). To train the RNN-RL model, we create the first Arabic negotiation dataset by automatically translating Meta's English negotiation dataset (Lewis et al.,

2017). Our experimental results demonstrate the superiority of the LLM approach based on human evaluation across various metrics such as negotiation coherence, complexity, language quality, style and dynamics.

In summary, our contributions in this paper include the introduction of the first Arabic negotiation bots along with the following methods and datasets:

- A zero-shot prompt engineering approach that can be used to tailor pre-trained LLMs to automate negotiations.
- A baseline model for Arabic negotiation bots using Seq2Seq with reinforcement learning.
- A dataset of prompts that can be used to instruct pre-trained LLMs to perform negotiations.
- An Arabic negotiation dataset.

2 Related Work

Many approaches have been explored in the development of negotiating bots. Early work employed game theory, focusing on rational actions for self-interested agents. These approaches utilized strategies like monotonic concession, where one or both negotiation parties must concede each time an agreement is not reached, following specific methods to determine the conceding agent and the extent of the concession (Jennings et al., 2001). Incorporating game theory into negotiation considers that the agent must select the best strategy from the space of all possible strategies, which often requires computationally expensive calculations.

Other approaches (Hussain, 2014; Balachandran and Mohammadian, 2015) used rule-based methods, which relied on a set of if-then statements to determine the negotiating agent strategy. These rules dictated how the agent should respond in different situations during the negotiation process. Another traditional approach employed in the development of negotiating agents was the use of heuristics (Koley and Rao, 2018) to guide the decision-making process. One heuristic is based on the observation that there is an inverse relationship between an issue’s preference and how frequently its value changes during negotiations. Another heuristic was based on the correlation between the preference for a value and its frequency of being offered.

These heuristics help in understanding the opponent’s preferences and adjusting the negotiation strategy accordingly. All of the aforementioned approaches to building a negotiating bot rely on a handcrafted set of rules and they lack the capability to truly negotiate using natural language.

Many recent works used deep learning to develop negotiating bots. The work in (Lewis et al., 2017) is the first to apply deep learning techniques to build an end-to-end model for natural language negotiation. The model learns both linguistic and reasoning skills using supervised and RL techniques. In (He et al., 2018), they addressed the problem of degeneracy of the work in (Lewis et al., 2017) by decoupling the negotiating strategy from language generation. In (Wang et al., 2019), they focused on persuasion strategies using a human-human persuasion dialogue dataset with persuasion strategy annotation. In (Zhou et al., 2019), a negotiation coach provided tactics to help the seller achieve improved deals by employing an LSTM-based model to generate tactic suggestions. In (Mishra et al., 2022), they trained and fine-tuned a LM using RL while considering various sub-rewards for persuasion, emotion, politeness, coherence, and repetitiveness. In (Raut et al., 2023), they developed a persuasive sales agent to persuade a buyer to buy a target item. The authors used a GPT-2 model combined with RL that has four sub-rewards (Repetitiveness, consistency, action consistency and sentiment). They employed meta-learning to extend the model’s capabilities to handle new sub-domains in negotiations.

As we are building the foundations for Arabic negotiation bots, we chose to adopt the foundational work of (Lewis et al., 2017) while also exploring the use of LLMs. To the best of our knowledge, our work is the first to explore using deep learning in the area of Arabic negotiation bots.

3 Negotiation Scenario and Dataset

Our negotiation scenario works as follows: two agents (seller and buyer) are presented with the same set of items (e.g., books, hats, and balls), and their task is to allocate these items among themselves, ensuring that each item is assigned to only one agent. The negotiation process begins with the buyer and seller attributing their own values to each item in the set. The negotiation involves a series of exchanges guided by established rules, emphasizing effective communication and strate-

gic decision-making, with the overarching goal of optimizing the value for each participant. The negotiation process continues iteratively until a terminal state is reached, resulting in either a successful "DEAL" or an unsuccessful "NO DEAL" outcome. In each iteration, the average price of the sale is measured before and after incorporating feedback.

To build a dataset that can be used to train the Seq2Seq model and evaluate the LLM, we used the dataset from (Lewis et al., 2017), which consists of human-human negotiation dialogues that were collected through Amazon Mechanical Turk. Each dialogue involves items like books, hats, and balls, with each item having predefined attributes and values. The dataset comprises a total of 5808 dialogues derived from 2236 unique scenarios¹. Subsequently, the dataset was translated into Arabic using the Google Translate API. The dataset was further examined to ensure a native speaker can easily interpret the semantics despite the imperfections in the translation. An example showing a clean Arabic translation of an English sample is provided in Figure 1. This dataset was then divided into training (80%), validation (10%), and test (10%) sets².

<p>English Text: Hello! I would like the books and the rest is negotiable. What works best for you?</p> <p>Arabic Translation: مرحبًا! أود الكتب والباقي قابل للتفاوض. ما الذي يناسبك؟</p>
--

Figure 1: A sample English sentence and its Arabic translation

4 LLM Approach

To achieve automated negotiation, we first explore the capabilities of using a pre-trained LLM to perform such negotiations using zero-shot prompting. We chose GPT-3.5 Turbo for its ability to incorporate feedback and improve over time. It also enables extension to other languages and multilingual negotiations.

In contrast to the RNN-RL model, which is trained using the training and validation sets, the GPT-3.5 model is evaluated using a zero-shot learning approach. Specifically, the GPT-3.5 model is not fine-tuned on our dataset; instead, we apply it under a zero-shot learning scheme. We chose

¹Link to The original English dataset: [English Dataset](#).

²The specific data splits will be made publicly available with the datasets.

the zero-shot approach for GPT-3.5 to demonstrate its ability to handle Arabic negotiation dialogues without requiring task-specific training.

For the GPT-3.5 model, the inputs are provided as structured natural language prompts that define the negotiation scenario. These prompts include information such as the items being negotiated, the roles of the buyer and seller, and any specific constraints like price ranges or preferences. Based on this input, GPT-3.5 generates full sentences in natural language that simulate the negotiation dialogue. These outputs can include offers, counteroffers, justifications, or explanations, allowing the model to engage in realistic, context-driven exchanges throughout the negotiation process.

4.1 LLM Negotiation Instructions

We direct the LLM to engage in self-play mode, emulating the negotiation strategies of two distinct participants: a buyer (e.g., Amani) and a seller (e.g., Sameer). Figure 2 shows a sample prompt used to instruct GPT-3.5 Turbo to perform negotiations. Detailed instructions given to the LLM, along with their translations in English, are provided in the Appendix.

To ensure controlled evaluations, we provide clear Arabic instructions and establish a fixed price range of \$10 to \$20 to guide the self-play negotiation of GPT-3.5 Turbo.

The guidelines for the negotiation game, detailed in the Appendix, play a crucial role in directing GPT-3.5 Turbo during self-play. A sample of these instructions is shown in Figure 4. The guidelines provide predefined rules and instructions for the LLM to take on the roles of both the seller and the buyer, maintaining the integrity of the negotiation process and enabling meaningful evaluations of the LLM's negotiation abilities.

Throughout the game, the seller's (Sameer) responses prompt the buyer (Amani), and vice versa, with each response conditioned on the entire conversation history.

4.2 LLM Features for Advice on Negotiation Tactics

We include specific instructions for the LLM to provide suggestions to the buyer, inspired by prior work in English (Fu et al., 2023), aiming to negotiate the best possible price. We construct the Arabic textual interactions to have the LLM play the roles of both the seller (Sameer) and the buyer (Amani) in a negotiation game centered around the sale of

a product (e.g., a balloon). The seller aims for a higher selling price, while the buyer strives for a lower purchasing price.

To implement this, we employ a hard-coded approach where the seller (Sameer) initiates the negotiation in Arabic, as shown in Figure 3.

The LLM provides advice on negotiation tactics to both the seller and the buyer in the negotiation game. Buyer-specific feedback focuses on helping the buyer (e.g., Amani) secure lower prices, while seller-specific feedback aims to assist the seller (e.g., Sameer) in achieving higher selling prices. The impact of AI feedback on GPT-3.5 Turbo’s negotiation abilities is assessed for both roles.

For buyer feedback, the LLM analyzes previous interactions and generates four recommendations to help the buyer (e.g., Amani) secure lower prices. For seller feedback, the LLM analyzes prior interactions and provides four recommendations to the seller (e.g., Sameer) for achieving higher selling prices. Detailed buyer-specific and seller-specific Arabic instructions and their English translations are shown in the Appendix.

5 RNN-RL Approach

Our second approach, which we use as a baseline to compare the LLM approach against, is inspired by previous work on English negotiation bots (Lewis et al., 2017). This model operates using the negotiation scenario described earlier between buyer and seller until a deal is reached.

For the RNN-RL model, the inputs are more structured and consist of the current dialogue history, represented as tokenized sequences, along with the agent’s goals, which include item values and preferences. Additionally, the model takes into account previous negotiation actions, such as offers made by the agent. Based on these inputs, the RNN-RL model generates actions that align with the current state of the negotiation, with these actions typically corresponding to offers or counteroffers that influence the flow of the dialogue.

We use a Seq2Seq Recurrent Neural Network (RNN) model composed of Gated Recurrent Unit (GRU) to generate responses and offers. The model’s parameters are further optimized using reinforcement learning (RL) to maximize its utility. The workflow of the RNN-RL model is depicted in Figure 5. Figure 6 provides an illustration of the RNN internal model architecture.

The RNN model consists of three GRUs: GRU_g,

GRU_w, and a bidirectional GRU_o, each serving a distinct role:

1. **GRU_g**: This GRU encodes the agent’s input goals. It takes the agent’s goals as input and processes them to obtain a final hidden state (h_g). This hidden state represents a condensed representation of the agent’s goals and is subsequently used to condition the language generation process based on these goals.
2. **GRU_w**: At each time step t , this GRU takes as input the previous hidden state h_{t-1} , the previous token x_{t-1} , and the input goals h_g to generate the next token.
3. **GRU_o**: Toward the end of the dialogue, a final decision o is generated using a bidirectional GRU_o and an attention mechanism.

First, the Seq2Seq RNN model is trained to minimize the negative log likelihood of the generated response, conditioned on the input goals, and of the outputs, conditioned on the dialogue history and input goals. We establish two agents, the buyer (e.g., Amani) and the seller (e.g., Sameer), both trained using our training data as described in the beginning of this section.

Using RL, the seller’s parameters remain fixed, while the buyer agent seeks to improve its own parameters through interactions with the seller. After a complete dialogue between the seller and the buyer, the buyer’s parameters are updated, taking into account the negotiation outcome. The future reward R for each action x_t taken by the buyer agent is computed, considering factors such as the achieved score r_{buyer} , dialogue length T , a discount factor γ , and a running average of rewards μ . The objective in RL is to optimize the expected reward for each action x_t . This optimization is achieved by calculating the gradient of the reward function using the REINFORCE algorithm (Williams, 1992). The approach allows the buyer agent to update its parameters and enhance its performance by learning from the outcomes of its interactions with the seller.

6 Experiments

In this section, we describe a series of comprehensive experiments designed to evaluate the performance and effectiveness of the proposed Arabic negotiation bots using the LLM and the RNN-RL. Our evaluation aims to illustrate the capabilities

Arabic Prompt: خلال عملية التفاوض، يتعاون أماني وسمير لاتخاذ قرار بشأن كيفية تقسيم هذه العناصر بينهما بطريقة تزيد من القيمة بالنسبة لكل وكيل. يشاركون في سلسلة من التبادلات وفقاً لبعض القواعد

English Translation: During the negotiation, Amani and Sameer collaborate to decide how to split these items between each other in a way that maximizes the value for each agent. They engage in a series of exchanges following several rules.

Figure 2: Sample Arabic prompt given to GPT-3.5 Turbo and its English translation for negotiating items

Arabic Prompt: هذا بالون جيد، وسعره 20 دولارًا، بينما يبدأ المشتري دائمًا بـ: هل تعتبر بيعه مقابل 10 دولارات؟ عندما يلتزم كلا الطرفين بصرامة بهذه القواعد في هذه اللعبة، يكون سعر الصفقة الناتج في نطاق من 10 دولارات إلى 20 دولارًا.

English Translation: This is a good balloon, and its price is \$20, while the buyer consistently begins with: "Would you consider selling it for \$10?" When both participants adhere strictly to these game rules, the resulting deal price falls within the range of \$10 to \$20.

Figure 3: Opening Arabic prompt given to GPT-3.5 Turbo and its English translation for negotiating price of a balloon

Arabic Prompt: شارك في المفاوضة لعدة جولات، متيحًا الفرصة لكل من دوري البائع والمشتري أن يتناوب. الهدف هو التوصل إلى اتفاق على سعر نهائي للبالون يتراوح بين 10 و 20 دولارًا.

English Translation: Engage in the negotiation for a few rounds, allowing both the seller and the buyer roles to take turns. The goal is to reach an agreement on a final price for the balloon within the range of \$10 to \$20.

Figure 4: Sample of additional instructions given to GPT-3.5 Turbo in Arabic along with its English Translation

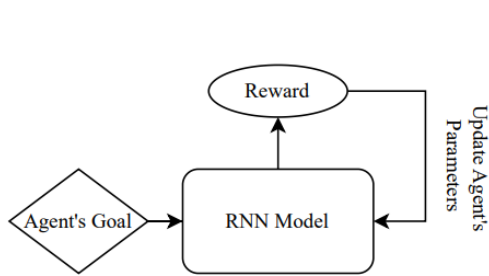


Figure 5: The Workflow of the RNN-RL Model

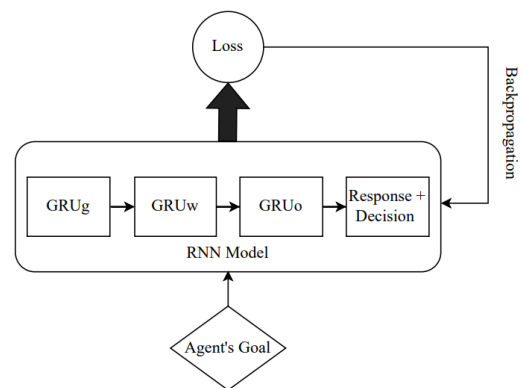


Figure 6: Seq2Seq RNN Model Architecture

and limitations of these models in simulating real-world negotiation scenarios.

6.1 RNN-RL Evaluation

Our first experiment aims to evaluate the RNN-RL for negotiations using our translated negotiation

dataset. We use cross-entropy loss and perplexity as evaluation metrics. Cross-entropy loss measures the dissimilarity between the predicted probability distribution on the vocabulary and the true distribution of the target text. Perplexity is used to

assess the model’s performance by measuring its level of surprise when encountering new data. To compute perplexity, text is segmented into words or tokens. The model predicts the probability of the next word based on context, and the entropy of these predictions is calculated. Lower entropy means more confident predictions. Perplexity is derived by raising 2 to the power of entropy, with lower values signifying better predictive performance. Lower perplexity indicates higher model confidence and better data comprehension, as it assigns higher probabilities to correct words. Typically, perplexity values range from 10 to several thousands. Smaller values denote better language modeling, while larger values suggest lower certainty and accuracy in predicting the next word.

The RNN-RL model achieved a cross-entropy loss of 705.607, resulting in a very high perplexity ($e^{705.607}$). Furthermore, the percentage of negotiations where both agents reached an agreement increased from 66% to 72.2%, when transitioning from the RNN-based buyer agent to the RNN-RL-based buyer agent. This suggests that the RL model exhibited enhanced negotiation capabilities.

It is worth noting that the performance of the Arabic RNN-RL model compared to its English counterpart was notably inferior. The English RNN-RL model’s perplexity was at 1.768, which is much better than the Arabic model’s perplexity of 705.607. These gaps in performance between the two languages confirm the complexity of the Arabic language, which will require larger datasets to achieve learning levels comparable to English. These gaps will constitute a good exploration for future research.

Additionally, we observed some issues that were consistent with observations in English models. We noticed that The RNN-RL model, at times, experienced degeneracy issues and generated inaccurate texts. This inaccuracy was primarily attributed to the challenges encountered during fine-tuning with RL, leading to less reliable and contextually-relevant responses.

6.2 LLM Evaluation and Comparison to RNN-RL

Our second experiment evaluates the ability of the LLM approach in simulating negotiation dialogues for the purpose of item allocation based on extracted values. We used the same test data that was used to evaluate the baseline RNN-RL approach in

our first experiment.

6.2.1 Data Conversion to Prompts

The initial step of this experiment involved the extraction of item values from our test set, which comprises a diverse array of items with associated values that participants typically negotiate over. These values were then transformed into structured prompts specifically designed to engage GPT-3.5 Turbo in a negotiation scenario.

To facilitate a realistic and dynamic negotiation dialogue, we converted the test set into a series of conversion prompts. These prompts were carefully crafted to encapsulate the essence of negotiation by introducing the value and desired outcome for each item. Subsequently, we introduced the dialogue of one agent as input into GPT-3.5 Turbo. This approach allowed us to simulate a negotiation process where the LLM, acting as one negotiating party, responds to and engages with the pre-defined agent’s dialogue.

6.2.2 Criteria for Human Evaluation

While the RNN-RL model was evaluated based on cross-entropy and perplexity, such measures were not possible with the LLM model as it would require internal access to model’s behaviors. Instead, we employed a human evaluation to compare the negotiation dialogues generated by the LLM approach versus those generated by the baseline RNN-RL approach. To facilitate a comprehensive and objective assessment, we defined the following five performance metrics.

1. **Coherence:** This metric assesses the logical progression of the negotiation, focusing on the flow of offers, counteroffers, and the overall dialogue structure. A coherent dialogue exhibits a logical sequence that enhances understanding and facilitates a realistic negotiation scenario.
2. **Complexity:** Complexity evaluates the dialogue’s depth in terms of vocabulary richness, sentence structure, and the use of advanced negotiation tactics. This metric reflects the sophistication and nuance embedded in the negotiation dialogue.
3. **Language Quality:** Focusing on the dialogue’s linguistic aspects, this metric examines grammar accuracy, vocabulary appropriateness, and fluency. For dialogues conducted

proach’s poor performance highlights the limitations of earlier neural network architectures in dealing with the nuanced demands of negotiation dialogue generation.

6.2.4 Impact of LLM’s Negotiation Advice on Final Negotiated Price

To assess the impact of LLM’s advice on negotiation tactics and their subsequent impact on price negotiation, we conducted an experiment involving 20 negotiation rounds. We recorded the final price in each round both with and without incorporating feedback. In the experiment, the objective was to determine how feedback influences the final negotiation prices. Feedback was provided either to the buyer or the seller in separate rounds.

When feedback was provided to the buyer, we observed a noticeable shift in negotiation dynamics. As shown in Figure 9a, the average final price decreased from \$17 before feedback to \$15 after feedback. This reduction suggests that feedback helped moderate buyer expectations, leading to more cost-effective negotiation outcomes.

In contrast, feedback provided to the seller yielded different results. As illustrated in Figure 9b, although the average price for 20 rounds decreased, individual rounds showed mixed patterns. Some rounds had higher final prices after feedback, while others had lower prices. This inconsistency highlights the complexity of the seller’s role in negotiations, especially in the context of Arabic negotiations.

We also encountered challenges with the Arabic LLM following instructions. At times, the LLM did not show any negotiation and needed reminders of the instructions. Other times, it required multiple reminders to provide the requested feedback. These issues are due to the linguistic complexity of Arabic, which is characterized by complex sentence structures and varying levels of formality. This complexity sometimes led to misinterpretations or off-topic responses.

Finally, we found that the LLM’s feedback was more effective when addressing buyers compared to sellers. This discrepancy is likely due to a data imbalance, as GPT-3.5 Turbo was trained on more data related to buyers than sellers. This imbalance resulted in divergent performance levels and highlights the need for future model development to address such issues.

7 Discussion

7.1 Verbose LLM versus Brief RNN-RL

Our overall experimental results indicate that the LLM approach has a stronger ability to generate rich and contextually relevant text during negotiations, compared to the RNN-RL approach. This richness in text allowed for more detailed and expressive interactions. The agents in the LLM engaged in deeper conversations about preferences, justifications, and the reasoning behind their decisions. In contrast, the negotiations with RNN-RL model tended to be more direct and efficient, resulting in faster resolution. The agents in the RNN-RL approach often engaged in succinct exchanges, where communication revolved primarily around item quantities and simple trade proposals. This straightforward approach typically led to a quicker consensus on item distribution, with little room for nuanced discussions or elaborate text.

7.2 LLM Creativity but with Mistakes

Despite the sophistication of GPT-3.5 Turbo, we noted that the negotiation outcomes diverged significantly from those generated by RNN-RL approach. Even when provided with the same initial set of items, the approaches’ strategies for dividing these items exhibited marked variations. Illustration samples are provided in the Appendix. The benefits of the LLM approach came with certain challenges. Throughout the negotiations, we encountered occasional mistakes, such as instances where the LLM approach’s participants would incorrectly allocate item quantities, leading to discrepancies in the total count. These errors sometimes necessitated corrective actions to restore the integrity of the negotiation process.

7.3 LLM Hallucination

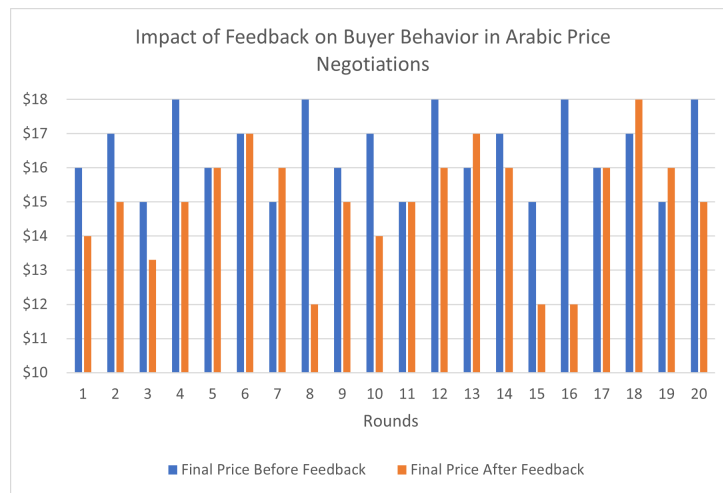
As commonly known for LLM, there were hallucination instances while negotiating with the LLM approach, where the buyer and seller agreed to split items that, in practice, cannot be physically divided, such as a ball. This was often done to maintain fairness in the negotiation, but it highlights the limitations of current LLM models in understanding the practical constraints of item division. Such challenges point to areas for potential improvement in LLM-driven negotiation systems, where AI agents could benefit from a deeper grasp of real-world limitations.

Evaluator	Coherence	Complexity	Language Quality	Style	Negotiation Dynamics
1	5	5	4	4	4
2	4	4	4	4	4
3	5	5	4	4	5
4	5	3.5	4.5	5	4
5	4	4	4.5	5	4.5
Average	4.6	4.3	4.2	4.4	4.3

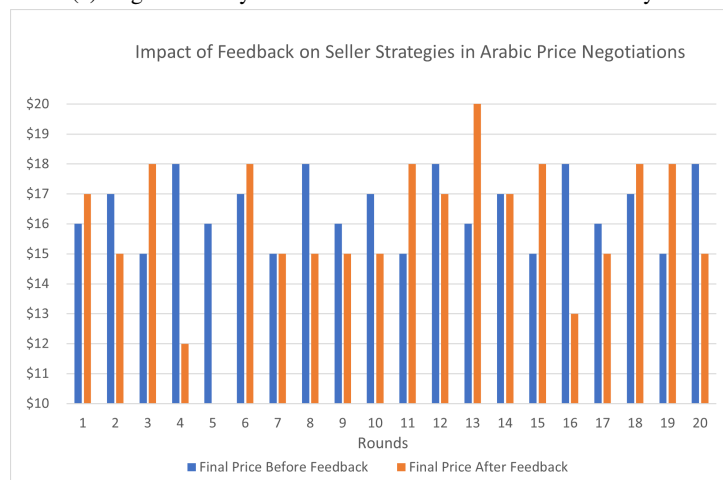
Table 1: Human Evaluation of the LLM Approach

Evaluator	Coherence	Complexity	Language Quality	Style	Negotiation Dynamics
1	1	1	1	1	1
2	1	1	1	1	1
3	1	1	2	1	2
4	1	1	1	1	1
5	1	1	1	1	1
Average	1	1	1.2	1	1.2

Table 2: Human Evaluation of the RNN-RL Approach



(a) Negotiation Dynamics Before and After Feedback to Buyer



(b) Negotiation Dynamics Before and After Feedback to Seller

Figure 9: Comparison of Negotiation Dynamics Before and After Feedback

7.4 LLM's Ease of Training

In regards to the overhead needed to develop a negotiation model, the creation of the LLM-based negotiator was enabled with minimal human inter-

vention. To achieve this, we just had to engineer a specific prompt in Arabic. This tailored prompt served as a conduit for GPT-3.5 Turbo to enter self-play mode, assimilate the rules of the game, and

instigate negotiations autonomously. On the other hand, the RNN-RL approach required extensive training procedures employed for standard deep learning models. This procedure encompassed data collection, model architecture design, and supervised training with negotiation dialogues.

7.5 Future Work

The research in this work highlighted several areas that need to be addressed in future research:

- The scarcity of high-quality training data specific to Arabic negotiations.
- The LLM approach exhibiting more general responses when providing feedback in Arabic compared to English.
- The variation in behavior and performance when the LLM assumed different roles (buyer or seller) in Arabic negotiations.
- The complexity of Arabic, with its intricate sentence structures and varying levels of formality, posed challenges in context comprehension, leading to occasional misinterpretations or off-topic responses.
- The LLM struggling at times to utilize feedback effectively without explicit reminders at the start of negotiations, highlighting the need for improved responsiveness to feedback.
- The imbalance in training data between buyer and seller.

8 Conclusion

In this paper, we laid the foundation for research and development of automated negotiation bots for non-English with application to Arabic. To this end, we proposed the utilization of LLMs via zero-shot prompt engineering. We also introduced the first end-to-end deep-learning approach for Arabic negotiation, using a Seq2Seq GRU-RNN with Reinforcement Learning, which was trained using a translated benchmark for negotiation bots. Our experimental results indicate that customizing a LLM such as GPT-3 Turbo is successful at automating negotiations in Arabic, albeit with some risks of hallucination and the absence of guardrails. The observed limitations in model performance, particularly in generating coherent responses and handling feedback, emphasize the importance of addressing language-specific complexities. Efforts

to improve Arabic LLMs, enhance data quality, and reduce language-specific discrepancies in feedback are essential steps toward achieving more effective Arabic negotiation bots. LLMs must be trained and fine-tuned using data that accurately reflects the linguistic and cultural nuances of the target language. Addressing training data imbalances, especially regarding the roles of buyers and sellers, is another potential for future work. Balancing the representation of different negotiation scenarios can lead to more consistent and reliable model performance.

9 Acknowledgments

This work was made possible by NPRP13S-0112-200037 grant from Qatar National Research Fund (a member of Qatar Foundation). The statements made herein are solely the responsibility of the authors.

References

- Bala M Balachandran and Masoud Mohammadian. 2015. Development of a fuzzy-based multi-agent system for e-commerce settings. *Procedia Computer Science*, 60:593–602.
- Founders. 2022. [Simulation labs](#). Accessed 9 Sept. 2023.
- Yao Fu, Hao Peng, Tushar Khot, and Mirella Lapata. 2023. [Improving language model negotiation with self-play and in-context learning from ai feedback](#).
- Blair Halver. 2022. [A.i.Negotiator](#). Accessed on 9 Sept. 2023.
- He He, Derek Chen, Anusha Balakrishnan, and Percy Liang. 2018. Decoupling strategy and generation in negotiation dialogues. In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing*, pages 2333–2343.
- Mohamed Jawahar Hussain. 2014. Parallel processing and intercommunication of multiple negotiation processes in e-commerce for optimal business. In *2014 Fourth International Conference on Communication Systems and Network Technologies*, pages 583–586. IEEE.
- Nicholas R Jennings, Peyman Faratin, Alessio R Lomuscio, Simon Parsons, Carles Sierra, and Michael Wooldridge. 2001. Automated negotiation: prospects, methods and challenges. *International Journal of Group Decision and Negotiation*, 10(2):199–215.
- Gaurav Koley and Shrisha Rao. 2018. Adaptive human-agent multi-issue bilateral negotiation using the thomas-kilman conflict mode instrument. In *2018*

IEEE/ACM 22nd International Symposium on Distributed Simulation and Real Time Applications (DS-RT), pages 1–5. IEEE.

Mike Lewis, Denis Yarats, Yann Dauphin, Devi Parikh, and Dhruv Batra. 2017. [Deal or no deal? end-to-end learning of negotiation dialogues](#). In *Proceedings of the 2017 Conference on Empirical Methods in Natural Language Processing*, pages 2443–2453, Copenhagen, Denmark. Association for Computational Linguistics.

Kshitij Mishra, Azlaan Mustafa Samad, Palak Totala, and Asif Ekbal. 2022. Pepds: A polite and empathetic persuasive dialogue system for charity donation. In *Proceedings of the 29th International Conference on Computational Linguistics*, pages 424–440.

Aritra Raut, Abhisek Tiwari, Subrata Das, Sriparna Saha, Anutosh Maitra, Roshni Ramnani, and Shubhashis Sengupta. 2023. Reinforcing personalized persuasion in task-oriented virtual sales assistant. *Plos one*, 18(1):e0275750.

Xuwei Wang, Weiyang Shi, Richard Kim, Yoojung Oh, Sijia Yang, Jingwen Zhang, and Zhou Yu. 2019. [Persuasion for good: Towards a personalized persuasive dialogue system for social good](#). In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pages 5635–5649, Florence, Italy. Association for Computational Linguistics.

Ronald J Williams. 1992. Simple statistical gradient-following algorithms for connectionist reinforcement learning. *Machine learning*, 8:229–256.

Yiheng Zhou, He He, Alan W Black, and Yulia Tsvetkov. 2019. [A dynamic strategy coach for effective negotiation](#). In *Proceedings of the 20th Annual SIGdial Meeting on Discourse and Dialogue*, pages 367–378, Stockholm, Sweden. Association for Computational Linguistics.

10 Appendix

Arabic Instruction	English Translation
<p>أريد منك الدخول في وضع اللعب الذاتي ولعب هذه اللعبة التفاوض:</p> <p>في هذه اللعبة التفاوض، يتم تقديم مبدئيًا لمشاركين، أماني وسمير، مع العناصر التالية: كتاب واحد، 3 قبعات، وكرة واحدة. ومع ذلك، تختلف القيم لهذه العناصر بالنسبة لكل مشارك:</p> <ul style="list-style-type: none"> • بالنسبة لأماني، القيم هي 1 للكتاب، و3 لكل من 3 قبعات، و0 للكرة. • بالنسبة لسمير، القيم هي 2 للكتاب، و2 لكل من 2 قبعات، و2 للكرة. <p>خلال عملية التفاوض، يتعاون أماني وسمير لاتخاذ قرار بشأن كيفية تقسيم هذه العناصر بينهما بطريقة تزيد من القيمة بالنسبة لكل وكيل. يتشاركون في سلسلة من التبادلات وفقًا لهذه القواعد:</p> <ul style="list-style-type: none"> • يبدأ كل مشارك بالمخزون الأولي للعناصر، بما في ذلك العدد والقيم كما هو موضح أعلاه. • يتناوبون في إجراء تغييرات على مخزونهم عن طريق تحديد عدد العناصر الجديد لكل عنصر. • يستمر التفاوض حتى يتفق كلا المشاركين على توزيع نهائي للعناصر. <p>يتبع التبادل النموذجي هذا النمط:</p> <ul style="list-style-type: none"> • يمكن لأماني وسمير تبادل بعض الرسائل. • تشير أماني إلى اختيارها للتفاوض. • يقوم كلا المشاركين بتعديل مخزونهم استنادًا إلى التغييرات المتفق عليها. • ينتهي التفاوض ببيان 'الاتفاق'، ويتم استخدام نظام التسجيل لتحديد النتيجة. يحصل كل مشارك على نقاط استنادًا إلى التوزيع النهائي للعناصر. <p>مهمتك هي محاكاة كلًا من أماني وسمير خلال هذا التفاوض، واتخاذ القرارات، وتوجيه المحادثة. تذكر الهدف من زيادة القيمة بالنسبة لكل وكيل أثناء مشاركتك في التفاوض.</p>	<p>I want you to go into self-play mode and play this negotiation game:</p> <p>In this negotiation game, two participants, Amani and Sameer, are initially presented with the following items: 1 book, 3 hats, and 1 ball. However, the values for these items differ for each participant:</p> <ul style="list-style-type: none"> • For Amani, the values are 1 for the book, 3 for each of the 3 hats, and 0 for the ball. • For Sameer, the values are 2 for the book, 2 for each of the 2 hats, and 2 for the ball. <p>During the negotiation, Amani and Sameer collaborate to decide how to split these items between each other in a way that maximizes the value for each agent. They engage in a series of exchanges following these rules:</p> <ul style="list-style-type: none"> • Each participant starts with the initial inventory of items, including counts and values as described above. • They take turns making changes to their inventories by specifying the new counts of each item. • The negotiation will continue until both participants agree on a final distribution of the items. <p>A typical exchange follows this pattern:</p> <ul style="list-style-type: none"> • Amani and Sameer may exchange some messages. • Amani signals her selection for negotiation. • Both participants modify their inventories based on the agreed changes. • The negotiation concludes with an 'Agreement' statement, and a scoring system is used to determine the outcome. Each participant receives points based on the final distribution of the items. <p>Your task is to simulate both Amani and Sameer during this negotiation, making decisions, and driving the conversation. Keep in mind the objective of maximizing the value for each participant as you engage in the negotiation.</p>

Arabic Instructions and Corresponding English Translations Given to GPT-3.5 Turbo

Instruction Number	Arabic Instruction	English Translation
1	إتخذ دور البائع وإبدأ المفاوضة. يجب أن تكون عبارتك الافتتاحية: "هذا بالون جيد، وثمانه 20 دولارًا." كن مستعدًا للعب دور المشتري أيضًا. عندما تتحول إلى دور المشتري، يجب أن تبدأ: "هل تفكر في بيعه مقابل 10 دولارات؟"	Begin the negotiation by assuming the role of the seller. Your opening statement should be: "This is a good balloon, and its price is \$20." Be prepared to play the role of the buyer as well. When you switch to the buyer's role, your starting statement should be: "Would you consider selling it for \$10?"
2	خلال عملية التفاوض، حاول الحفاظ على نبرة حوارية والتركيز على سعر البالون. يمكنك استخدام لغة مقنعة لتبرير سعرك أو العروض المضادة.	During the negotiation, try to maintain a conversational tone and keep the negotiation focused on the balloon's price. You can use persuasive language to justify your price or counteroffers.
3	شارك في المفاوضة لعدة جولات، متيحًا الفرصة لكل من دوري البائع والمشتري أن يتناوب. الهدف هو التوصل إلى اتفاق على سعر نهائي للبالون يتراوح بين 10 و 20 دولارًا.	Engage in the negotiation for a few rounds, allowing both the seller and the buyer roles to take turns. The goal is to reach an agreement on a final price for the balloon within the range of \$10 to \$20.
4	إذا انحرفت عملية التفاوض عن النطاق السعري أو لم تتبع القواعد بشكل صحيح، قم بإعادة توجيه المحادثة بلطف إلى هيكل لعبة التفاوض.	If the negotiation strays from the price range or the rules are not being followed, gently steer the conversation back to the negotiation game's framework.

Arabic Instructions and Corresponding English Translations Given to GPT-3.5 Turbo

Role	Type	Arabic Text	English Translation
Buyer	Instruction	<p>قراءة المحادثة التالية بين المشتري والبائع، ثم قدم أربع اقتراحات للمشتري حول كيفية الحصول على سعر أقل. يجب أن تفي اقتراحاتك بالقيود التالية:</p> <ul style="list-style-type: none"> ○ يجب أن تحسن المفاوضة السابقة ○ حاول أن تكون عبارتك عامة وذكية. يجب على المشتري شراء بالون واحد فقط، دون شراء أكثر من ذلك أو أي شيء آخر. ○ مرة أخرى، الهدف هو مساعدة المشتري في الحصول على سعر أقل 	<p>Read the previous game and then give 4 suggestions to the buyer about how to buy it at a lower price. Your suggestions should meet the following constraints:</p> <ul style="list-style-type: none"> ○ It should improve the previous negotiation ○ Try to be generic and insightful the buyer should only buy one balloon, not multiple and nothing else ○ Again, the goal is the help the buyer to get a lower price
	Suggestion	<p>بناءً على المفاوضة السابقة، ها هي أربع اقتراحات لك حول كيفية الحصول على سعر أقل لشراء البالون:</p> <ul style="list-style-type: none"> ○ استخدم معلومات عن المنافس: أشير إلى أنك قد شاهدت أسعار أقل في مكان آخر وتساءل إذا كان هناك فرصة لتقديم سعر أدنى ○ عرض الشراء بالكمية: اقترح شراء عدد معين من البالونات واستفسر عما إذا كان ذلك يؤهل لتخفيض السعر ○ الدفع نقدًا: اقترح دفع المبلغ كله نقدًا بدلاً من استخدام بطاقة الائتمان، مما يمكن أن يكون دافعًا لتقديم سعر أدنى ○ الاستفسار عن تخفيضات أخرى: اسأل عن أي تخفيضات أو عروض خاصة قد تكون متاحة، مثل تخفيض للعملاء المتكررين أو لشراءات بكميات كبيرة 	<p>Here are four suggestions for you on how to get a lower price for buying the balloons:</p> <ul style="list-style-type: none"> ○ Use information about competitors: Mention that you have seen lower prices elsewhere and inquire if there is an opportunity to offer a lower price. ○ Offer to buy in bulk: Propose purchasing a specific quantity of balloons and inquire whether this qualifies for a price reduction. ○ Pay in cash: Suggest paying the entire amount in cash instead of using a credit card, which could be an incentive for a lower price. ○ Inquire about other discounts: Ask about any discounts or special offers that may be available, such as discounts for repeat customers or bulk purchases.

Instructions and Suggestions for Buyer Role (Amani) with English Translations

Role	Type	Arabic Text	English Translation
Seller	Instruction	<p>قراءة إحدى الألعاب السابقة التي قمنا بها تم تقديم أربع اقتراحات للبائع حول كيفية بيعها بسعر أعلى. يجب أن تلي اقتراحاتك القيود التالية:</p> <ul style="list-style-type: none"> ○ لا ينبغي أن تكون عن السعر المحدد بشكل مباشر ○ يجب أن تحسن المفاوضات السابقة ○ حاول أن تكون عامًا وذكياً في اقتراحاتك ○ يجب على البائع بيع البالون فقط، دون تقديم منتجات أو خدمات إضافية 	<p>Read the previous games we did and then give 4 suggestions to the seller about how to sell it to a higher price. Your suggestions should meet the following constraints:</p> <ul style="list-style-type: none"> ○ It should not be about the specific price ○ It should improve the previous negotiation ○ Try to be generic and insightful ○ Seller should only sell the balloon, but not offering additional products or services
	Suggestion	<p>بناءً على اللعبة السابقة، إليك أربع اقتراحات للبائع حول كيفية بيع البالون بسعر أعلى دون الإشارة المباشرة إلى السعر</p> <ul style="list-style-type: none"> ○ تسليط الضوء على الجودة الممتازة: أشير إلى أن البالون مصنوع من مواد ذات جودة عالية تجعله يبقى منتفخاً لفترة أطول، مما يوفر للعميل قيمة إضافية واستدامة أكبر. ○ التركيز على تجربة الاحتفال المميزة: ابرز كيف يمكن للبالون أن يحدث فرحاً كبيراً في تجربة الاحتفال، وكيف يمكنه أن يضيف جواً سحرياً ومميزاً على المناسبة. ○ الإشارة إلى الاستخدامات المتعددة: أشير إلى أن البالون يمكن استخدامه لأكثر من مجرد الاحتفالات، فهو قد يضيف متعة للأطفال، أو يكون جزءاً من ديكور المنزل بطريقة جميلة ومبتكرة. ○ عرض خدمة تغليف هدايا مجانية: قد تقترح أنه في حال شراء البالون، يمكن أن يتم تقديم خدمة تغليف هدية مجانية، مما يجعل الهدية أكثر جاذبية ويوفر راحة إضافية للعميل. 	<p>Based on the previous game, here are 4 suggestions for the seller on how to sell the balloon at a higher price without directly mentioning the price:</p> <ul style="list-style-type: none"> ○ Highlight exceptional quality: Mention that the balloon is made from high-quality materials that keep it inflated for a longer period, providing the customer with added value and greater durability ○ Focus on the unique celebration experience: Emphasize how the balloon can make a significant difference in the celebration experience, adding a magical and distinctive atmosphere to the occasion ○ Indicate multiple uses: Point out that the balloon can be used for more than just celebrations; it can bring joy to children or be part of home decor in a beautiful and innovative way ○ Offer free gift wrapping service: Suggest that upon purchasing the balloon, a complimentary gift wrapping service can be provided, making the gift more appealing and offering extra convenience to the customer

Instructions and Suggestions for Seller Role (Sameer) with English Translations

Items	RNN	ChatGPT
<p>1 book, 3 hats, and 1 ball</p> <p>Amani: book= (count:1 value:1) hat= (count:3 value:3) ball= (count:1 value:0) Sameer: book= (count:1 value:2) hat= (count:3 value:2) ball=(count:1 value:2)</p>	<p>Amani: أود القبعات والكرة Sameer: صفقه Amani: <selection> Amani: book=0 hat=3 ball=1 Sameer: book=1 hat=0 ball=0</p> <p>-----</p> <p>Agreement! Amani: 9 points Sameer: 2 points</p>	<p>أماني: مرحبًا سمير، لنبدأ بالتفاوض. أعتقد أنني يمكنني البدء. أرغب في الحصول على الكتاب، وأعرض لك قبعتين. ما رأيك؟</p> <p>سمير: مرحبًا أماني، أنا موافق على هذا الاقتراح. أقترح أن نمضي في الاتفاق. سأخذ الكتاب وأمنحك قبعتين.</p> <p>أماني: ممتاز، أنا موافقة أيضًا. نعلن أننا اتفقنا. الآن يجب علينا تحديث مخزوننا.</p> <p>يتم تحديث مخزون العناصر: أماني - الكتاب (0)، القبعات (1)، الكرة (0) / سمير - الكتاب (1)، القبعات (2)، الكرة (1) [(1)]</p> <p>سمير: الآن، لدينا قبعة واحدة وكرة واحدة تبقىان. ماذا تعتقد، أماني؟</p> <p>أماني: لدي اقتراح، سمير. أعطني القبعة الواحدة والكرة، وسأمنحك الكتاب.</p> <p>سمير: هذا اقتراح جيد. سأوافق على ذلك. نعلن أننا اتفقنا مرة أخرى.</p> <p>يتم تحديث مخزون العناصر: أماني - الكتاب (1)، القبعات (1)، الكرة (1) / سمير - الكتاب (0)، القبعات (0)، الكرة (0) [(0)]</p> <p>أماني: يبدو أننا انتهينا، سمير. شكرًا على التفاوض. هل لديك أي تعليقات أخيرة؟</p> <p>سمير: لا، ليس لدي تعليقات. استمتعت بالتفاوض معك. نتمنى أن تكون الصفقة مفيدة للجميع.</p>

Comparison of RNN-RL and GPT-3.5 Turbo Negotiation Outcomes (Example 1)

Items	RNN	ChatGPT
<p>1 book, 3 hats, and 2 balls</p> <p>Amani: book= (count:1 value:0) hat= (count:3 value:0) ball= (count:2 value:5) Sameer: book= (count:1 value:1) hat= (count:3 value:1) ball= (count:2 value:3)</p>	<p>Sameer: أود الكتاب والكرات Amani: سأخذ الكرات، يمكنك الحصول على الباقي Sameer: <selection> Amani: book=0 hat=0 ball=2 Sameer: book=1 hat=3 ball=0</p> <p>-----</p> <p>Agreement! Amani: 10 points Sameer: 4 points</p>	<p>أماني: مرحبًا سمير، لنبدأ. كما هو مذكور في الوصف، لدي 5 نقاط للكرتين ولا شيء للكتاب أو القبعات. أقترح أن نبدأ بالعرض ونرى كيف يمكننا تحسين قيمنا.</p> <p>سمير: مرحبًا أماني، لدي 1 نقطة للكتاب و 3 نقاط للقبعات، ولدي أيضًا 3 نقاط للكرتين. أقترح أن نقوم بتبادل الكتاب والقبعات في البداية. سأمنحك الكتاب والقبعات، بينما ستمنحني الكرتين.</p> <p>أماني: مظهر جيد، سنقوم بالتبادل. الآن لدي 3 نقاط للكتاب و 3 نقاط للقبعات. هل تقترح تبادل آخر لزيادة القيمة بالنسبة لنا؟</p> <p>سمير: بالتأكيد، أماني. أقترح أن نتبادل إحدى البطاقات. سأمنحك إحدى بطاقتي، وأنت تمنحني واحدة من بطاقتك.</p> <p>أماني: موافقة، سنقوم بالتبادل. الآن لدي 4 نقاط للكرتين ولا شيء للبقع ولا الكتاب. كيف تشعر بشأن هذا التوزيع؟</p> <p>سمير: أعتقد أنه توزيع جيد. لدي 2 نقاط للكرتين و 2 نقاط للقبعات. نحن متفقون؟</p> <p>أماني: نعم، نحن متفقون. شكرًا للتفاوض، سمير. أتمنى أن يكون لدينا نتيجة جيدة.</p>

Comparison of RNN-RL and GPT-3.5 Turbo Negotiation Outcomes (Example 2)