

Generative AI for Automated Negotiation

Authors:

Yasser Mohammad
NEC Corporation
y.mohammad@nec.com

Haifeng Chen
NEC Laboratories America
haifeng@nec-labs.com

Ryota Higa
NEC Corporation
r-higaryouta@nec.com

Tomohito Ando
NEC Corporation
t-andou_cq@nec.com

Satoshi Morinaga
NEC Corporation
mori-chin@nec.com

CONTENTS

1	Introduction	4
1.1	Motivation.....	4
1.2	Cross-Boundary Multi-agent systems.....	4
2	Technology Background	5
2.1	Automated Negotiation.....	5
2.2	Multimodal Time-Series Forecasting.....	9
2.3	Generative AI for Multimodal Time-Series Forecasting.....	9
3	Generative AI for Automated Negotiation	11
3.1	Generative AI for Interacting with Humans.....	12
3.2	Generative AI for Reasoning During Negotiation.....	13
3.3	Generative AI for Utility Function Evaluation.....	14
4	Use-Cases	15
4.1	Supply-Demand Matching.....	15
4.2	Indirect Matching.....	19
5	Success, Risks and Mitigation Measures	20
6	Conclusion	22
7	References	23
Annex A	24
A.1	Dataset used in this Paper.....	24
Acknowledgements	24

FIGURES

Figure 1	Components of an Automated Negotiation System.....	6
Figure 2	Possible subproblems of automated negotiation that can benefit from Generative AI technology.....	8
Figure 3	Two key components of our approach: (a) Multi-modal augmented encoder, and (b) Cross-modal fusion and offline synthesis [1].....	10
Figure 4	Structure of Generative AI-Powered Negotiation Agent.....	12
Figure 5	Example of Extraction from Email (based on real data).....	13
Figure 6	Example of output generation for Human-Agent negotiation using Generative AI.....	13
Figure 7	Processing Steps for the Supply-Demand Matching Use-Case.....	16
Figure 8	Using Generative AI to improve demand forecasting based on multimodal time-series forecasting (Direct forecasting use-case).....	16
Figure 9	Forecasting without Generative AI.....	16
Figure 10	Using enhanced future demand prediction to generate target procurement plan. Note that the target plan differs from the existing set of agreements which leads to negotiation.....	17
Figure 11	Negotiating for the new target procurement plan leading to lower carried inventory while guaranteeing all needs for manufacturing based on the new demand forecast.....	17
Figure 12	The achieved reduction in carried inventory based on negotiation alone and combining negotiation with Generative AI based multimodal time-series forecasting.....	18
Figure 13	The 20 factors used in the business-model in the second use-case. We only show the names of some of these factors.....	19

Generative AI for Automated Negotiation

Figure 14 Expected sales estimated in the second use-case.20
Figure 15 Procurement plan for one input product.....20
Figure 16 Results of the second use-case showing the value added by negotiation and Generative AI
based utility function modeling.20

TABLES

Table 1 Supply-Demand Matching Use-case. Inventory reduction for 4 products showing the value of the
proposed approach.19

Negotiation is ubiquitous in business and industrial operations. Automating these operations generates value by enhancing the speed and profitability of agreements between economic actors. Most advances in automated negotiation up to this point have focused on Game-Theoretic and Decision-Theoretic approaches leading to some drawbacks including (1) rigid negotiation protocols. (2) Context-blindness due to the difficulty of incorporating market information effectively for preference evaluation. This leads to suboptimal negotiation actions and eventually to lost value (money on the table). Generative AI allows us to tackle both problems leading to better negotiation protocols and strategies, easier adoption and higher business value. This paper introduces our vision and approach to achieve this result and presents a specific use cases in procurement.

1 INTRODUCTION

1.1 MOTIVATION

The ability to negotiate effectively is one of the cornerstones of success in business and industrial operations. Negotiation governs the agreements and transactions that drive value creation, operational efficiency, and profitability across the ecosystem. As organizations seek to accelerate and optimize these critical interactions, automating negotiation processes emerges as a powerful lever for enhancing both the speed and quality of business agreements.

However, most automated negotiation solutions to date have been built on Game-Theoretic and Decision-Theoretic foundations. While these methods offer mathematical rigor, they often result in inflexible negotiation protocols that are difficult to adapt to real-world complexities. This inflexibility stems from the need to rigorously define all aspects of the negotiation process in order to model it mathematically which is usually achieved by limiting the freedom of action of negotiators. For example, it is usually assumed that all exchanges are offers of complete contracts. In natural negotiations, other types of messages can augment offering like clarifications, arguments, information exchange about preferences, etc. Moreover, these traditional approaches struggle to incorporate dynamic market information and contextual factors, limiting their ability to evaluate preferences accurately and respond to changing business environments. This lack of context-awareness can lead to suboptimal negotiation outcomes and hinder the adoption of automation in practice.

Generative AI represents a transformative opportunity to overcome these limitations. By harnessing the adaptive and context-sensitive capabilities of generative models, organizations can develop negotiation protocols and strategies that are more flexible, responsive, and aligned with business objectives. This enables not only improved negotiation outcomes but also smoother integration and higher adoption rates across industrial settings. In this white paper, we outline our vision and approach for leveraging Generative AI in automated negotiation, with a particular focus on procurement use cases that demonstrate tangible business value.

1.2 CROSS-BOUNDARY MULTI-AGENT SYSTEMS

Multi-agent systems are rapidly gaining prominence in business and industrial environments, driven by the need for scalable, distributed solutions to complex operational challenges. These systems consist of multiple autonomous agents—software entities capable of independent

decision-making and action—that collaborate or compete to achieve organizational objectives. Their adoption enables businesses to automate intricate workflows, optimize resource allocation, and respond dynamically to changing market conditions, all while maintaining high levels of efficiency and resilience.

The increasing deployment of multi-agent systems is particularly evident in areas such as supply chain management, procurement, logistics, and manufacturing. Here, agents represent various stakeholders, departments, or even entire organizations, each with their own goals, constraints, and information sets. By delegating negotiation, scheduling, and coordination tasks to these agents, businesses can achieve faster response times, improved scalability, and greater adaptability in the face of uncertainty or disruption.

However, as multi-agent systems extend beyond the boundaries of a single organization, new challenges emerge in coordinating interactions between agents representing different business entities. Differences in objectives, information asymmetry, and varying negotiation protocols can lead to misalignment, inefficiencies, or even conflicts. Ensuring seamless coordination and effective communication across organizational boundaries becomes a critical concern, requiring robust protocols and adaptive strategies that can bridge diverse operational contexts. Addressing these challenges is essential for unlocking the full potential of multi-agent systems in driving business value and fostering collaborative innovation.

Generative AI-powered automated negotiation methods offer significant potential to address the coordination challenges inherent in multi-agent systems, especially when agents operate across organizational boundaries. By leveraging advanced language models and adaptive reasoning capabilities, generative AI can facilitate more flexible, context-aware negotiation protocols that dynamically accommodate diverse objectives, information asymmetries, and varying business practices. These systems can interpret and synthesize complex market signals, preferences, and constraints from multiple stakeholders, enabling agents to reach mutually beneficial agreements more efficiently. As a result, generative AI not only enhances the effectiveness of automated negotiations but also fosters greater interoperability and trust between autonomous agents representing different organizations, paving the way for more seamless and productive cross-boundary collaborations.

2 TECHNOLOGY BACKGROUND

2.1 AUTOMATED NEGOTIATION

Figure 1 shows the main components of an automated negotiation [2, 10]. The first component is the negotiation scenario/domain which defines the set of all possible agreements that could be reached by the negotiators. Usually this is described as the Cartesian product of the domains of a set of negotiation issues each representing a specific dimension of the negotiation problem. Each possible agreement is called an **outcome**, and the set of all possible agreements is called the **outcome-space**.

The second component is the negotiation protocol which dictates the types of actions available to negotiator and when they can execute them. The most used negotiation protocol is the alternating offers protocol (AOP) in which negotiators take turns offering complete outcomes from the outcome-space. When receiving an offer, a negotiator can accept it ending the negotiation with

agreement, reject it offering another outcome or leave the negotiation. A predefined number of offers or seconds can be used to ensure the termination of the protocol.

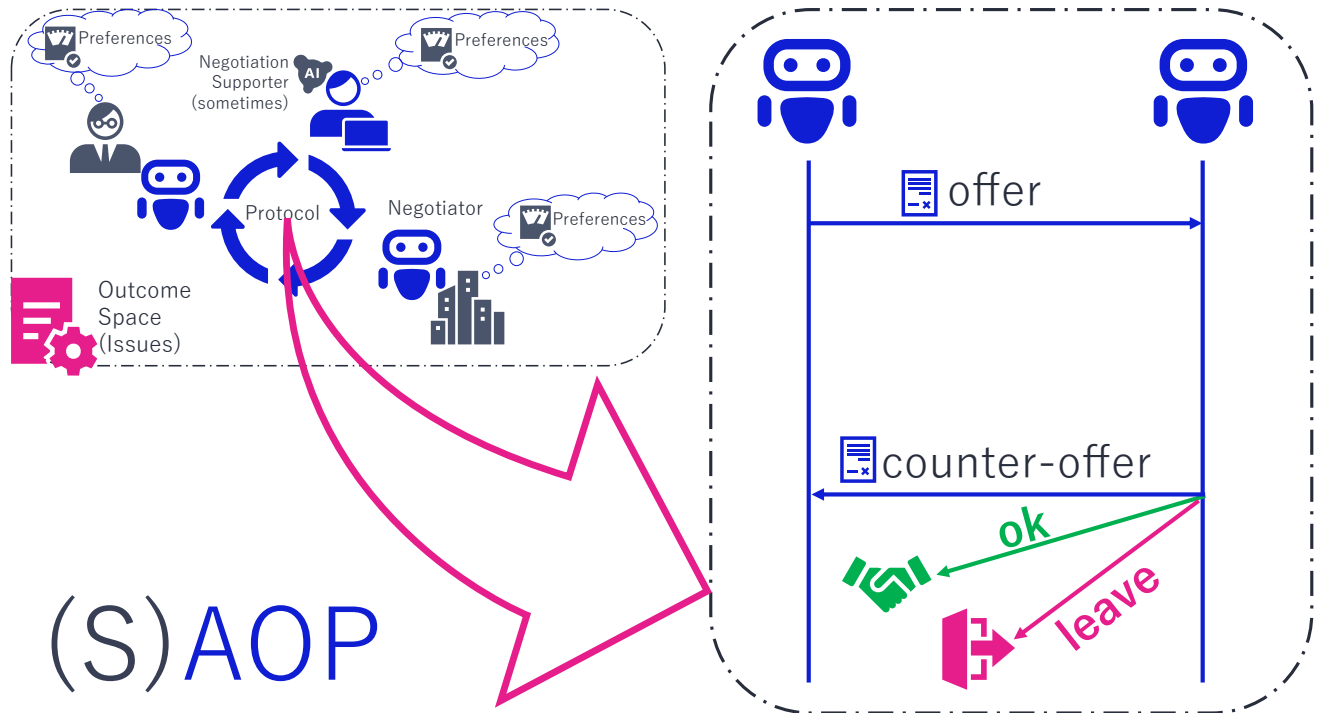


Figure 1 Components of an Automated Negotiation System. The left pane shows the relationship between the outcome-space, protocol, negotiators and their preferences. The right pane shows the flow of a negotiation between two negotiators using the Alternating Offers Protocol highlighting the three actions available in response of receiving an offer: acceptance, counter-offering or leaving the negotiation.

The third component is the set of preferences representing the interests of each negotiator in the output of the negotiation. Figure 1 shows an example of the three-party negotiation with one human negotiator, and two AI negotiators (one representing a human actor and another representing an organization). Even though everything discussed in this paper is applicable to multilateral negotiations with any number of negotiators; to simplify the exposition and notation, we will focus primarily on bilateral negotiations with two negotiators. The preferences of each negotiator should at minimum provide a partial ordering over outcomes in the outcome-space. Under widely applicable assumptions (i.e. completeness, transitivity, continuity, and independence), the Von Neumann–Morgenstern utility theorem ensures that we can ascribe an underlying utility function that maps any outcome in the outcome space to a real number ($u: \Omega \rightarrow \mathbb{R}$) to each negotiator for which the negotiator strives to maximize the expected utility value of the negotiation. Moreover, outside-options (i.e. alternatives available to the negotiator if the current negotiation fails) can be represented by a real-valued reservation value representing the

utility of disagreement for the negotiator. In this paper, we always assume that each negotiator has access to a utility function calculator that returns its utility for any given outcome but does not have any access to the utility function of its partners. This assumption ensures that the strategies used by the negotiator are not sensitive to incorrect information about partner utility functions. In practice, some information about the partner utility function may be available but this case is not considered in this paper.

Given the negotiation scenario, protocol, and preferences, each negotiator tries to reach an agreement that maximizes its own utility by employing a negotiation strategy. The structure of a negotiation strategy depends crucially on the protocol. In this paper we will focus on AOP because of its ubiquity and suitability for industrial application (e.g. being stateless, completely distributed, requiring no trusted third parties and ensuring privacy of preferences). Under AOP, we can describe a negotiation strategy using the following four components:

Bidding (offering) policy which maps the state of the negotiation to an outcome representing the offer of the negotiator

Acceptance policy which responds by either acceptance or rejection to an incoming offer from the negotiation partner.

Opponent model which adapts a model of partner strategy/preferences based on negotiation events. This component is optional as some strategies do not require such a model but availability of an opponent model has been shown to enhance negotiation outcomes (expected utility) in a wide-variety of domains.

These three components comprise the widely used BOA architecture [3] of automated negotiation. Another component that can come into play in more complicated scenarios is a strategy for ending the negotiation completely under some circumstances. The reason this is usually not considered a crucial part in automated negotiation research is that, under AOP, ending the negotiation is dominated by committing internally to always rejecting all offers if there are always some rational outcomes (i.e. outcomes with a utility higher than the reservation value) to offer. In practical business applications though this condition is not always satisfied, especially when concurrent negotiations are concerned. Under such conditions we add a leaving-policy is added leading to our BOLA architecture.

The BOLA architecture represents the negotiation strategy of a negotiator in a well-defined negotiation session with a specific scenario, protocol, preferences and partners. Nevertheless, practical applications of automated negotiation provide several other challenges that come before a negotiation session can start:

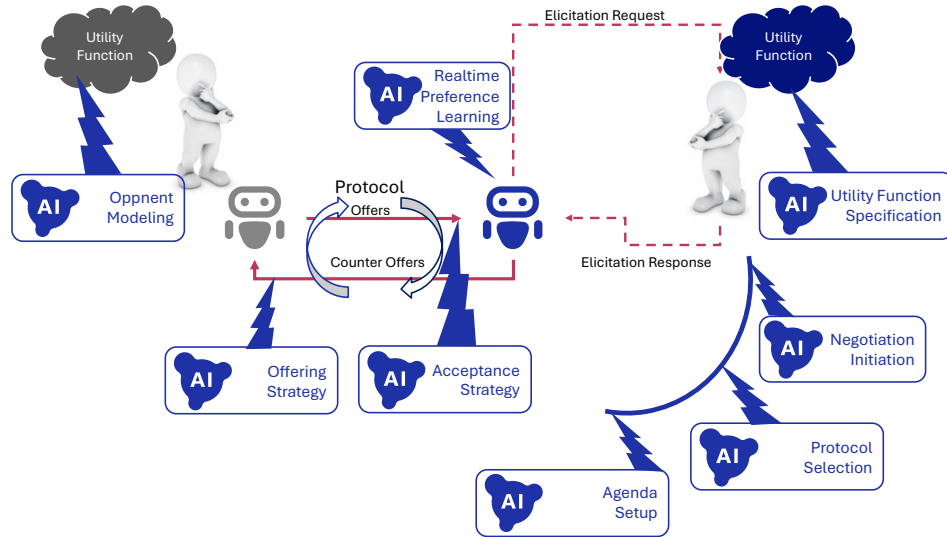


Figure 2 Possible subproblems of automated negotiation that can benefit from Generative AI technology.

Partner selection: It is not always obvious with whom are we supposed to negotiate. For example, a procurement manager may have to decide the supplier(s) to trade with based on previous performance and market conditions.

Protocol and Parameters selection: Even though most industrial negotiations use some form or another of the AOP protocol, some protocol parameters (e.g. time-limit) must be agreed upon before the negotiation can even start. Moreover, it may be beneficial in some cases to utilize other protocols like the single-text protocol, use auctions in some limited cases and so on.

Agenda setup: The discussion so far assumed that the negotiation agenda (i.e. outcome-space) is already defined and agreed upon between the negotiators. In practical applications this may not be the case. For example, the price range negotiated about is itself an important factor in determining the output of any trade negotiation and is usually a topic of negotiation itself. Moreover, in human-human negotiations, negotiators may come up with new negotiation issues (e.g. offering a discount on the next sale) that was not considered at the beginning of the negotiation. This kind of flexible agenda is rarely considered by traditional negotiation research. Agreeing on the partner list, protocol, parameters and the negotiation agenda is itself a coordination challenge and can be resolved using a meta-negotiation. Depending on the application this type of meta-negotiation may or may not be required. In this paper, we will assume that this type of meta-negotiation is either unnecessary or conducted manually. Many of the ideas presented carries through to meta-negotiations.

Figure 2 shows some of the ways Generative AI can enhance automated negotiation. Section 3 will touch upon these applications. Here we briefly describe simple examples of these enhancements. Preference elicitation is the process of learning about the user's preferences

which is a crucial challenge for practical application of automated negotiation [9]. Most research in this area uses predefined questions in the form of standard gambles that are difficult to understand and do not allow the user to express her knowledge effectively [9] as well as for extraction of offers from electronic messages (More details in Section 3.1). Generative AI can provide a possible way to extract well-defined preference models from natural language specifications. All aspects of meta-negotiation (agenda-selection, protocol and parameter selection, and partner selection) can benefit from Generative AI's ability to handle unstructured data. This becomes even more important when coordinating multiple negotiations concurrently [8]. All components of the negotiation strategy (offering and acceptance policies and the opponent model) can also benefit from Generative AI's reasoning ability (More details in Section 3.2). Utility function specification is the last application area in the Figure and is covered in details in Section 3.3.

2.2 MULTIMODAL TIME-SERIES FORECASTING

Time series forecasting is a crucial task with widespread applications across domains such as finance, healthcare, energy management, and environmental monitoring. In this paper, we will later discuss how to integrate it with automated negotiation technology to support novel industrial applications (Section 4.1). Accurate forecasting is vital for informed decision-making, optimizing resource allocation, and enhancing strategic planning. Recent advances in deep learning have led to significant progress in time series forecasting models. However, many time-series datasets are plagued by high noise levels, which pose substantial challenges in predicting future trends. For instance, stock prices are notoriously noisy due to the myriads of influencing factors like market sentiment, macroeconomic indicators, political events, and human behavior. To address these challenges, integrating information from related modalities has emerged as a promising approach. In healthcare, for example, clinical notes have been utilized to improve predictions of patient mortality. Similarly, in finance, incorporating text data from social media has been shown to enhance the accuracy of stock movement predictions. By combining time-series data with complementary sources of information, models can better capture complex patterns and reduce uncertainty while also addressing issues like data scarcity and noise. This integration ultimately enhances forecast reliability across a wide range of applications.

2.3 GENERATIVE AI FOR MULTIMODAL TIME-SERIES FORECASTING

The proposed method for multimodal time-series forecasting uses Generative AI to combine different types of data, such as numerical time-series (e.g., market data) and text (e.g., news reports), to create more accurate predictions. The process starts by breaking down the time-series data into its more predictable trend and seasonal parts. This decomposed data, along with the text information, is then augmented in several ways to create a richer dataset for the model to learn from. An encoder processes these various inputs to identify and learn the complex patterns and relationships that exist both within a single data stream and across the different modalities.

After the encoder processes the inputs, the system moves to a second stage where it generates multiple different predictions by creatively fusing the learned information from the various data sources. Rather than settling on a single output, this "cross-fusion" step produces a diverse set of possible forecasts. In the final step, these multiple predictions are combined and synthesized into a single, robust forecast. This aggregation process ensures that the final prediction is more

reliable and accurate, as it is based on a consensus derived from multiple perspectives on the data.

Let M denote the number of synchronized modalities and define the M -dimensional observation at time step t as $\mathbf{x}_t = (x_t^{(1)}, x_t^{(2)}, \dots, x_t^{(M)})$, $t = 1, \dots, L$, where each $x_t^{(m)}$ may originate from a heterogeneous source (numeric, text, image, audio, etc.). For instance, in market prediction involving both market time series and news data, $x_t^{(1)}$ could be the numerical market time series observed daily and $x_t^{(2)}$ might be an embedded representation of daily news data processed by a large language model (LLM). Other modalities can also be incorporated if they are converted into numerical representations.

We use $\mathbf{X}_{1:L} = (\mathbf{x}_1, \mathbf{x}_2, \dots, \mathbf{x}_L) \in \mathbb{R}^{L \times M}$ to describe historical data from all modalities. The objective is to predict the future trajectory of a time series using this multi-modal data. Assuming that $m=1$ is the numeric target series and the forecasting horizon is T steps, we define the target variable $\hat{\mathbf{y}}_{L+1:L+T} = (\hat{x}_{L+1}^{(1)}, \hat{x}_{L+2}^{(1)}, \dots, \hat{x}_{L+T}^{(1)}) \in \mathbb{R}^T$. Our aim is to establish a parameterized mapping $F_\theta: \mathbf{X}_{1:L} \mapsto \hat{\mathbf{y}}_{L+1:L+T}$, trained by minimizing a task loss, typically the mean squared error (MSE).

The function F_θ must exploit both intra-series dynamics (temporal patterns within modality 1 for example) and cross-modal dependencies encoded in $\mathbf{X}_{1:L}$ (e.g. correlations between sensor signals and accompanying text or images) to deliver forecasts that outperform unimodal baselines.

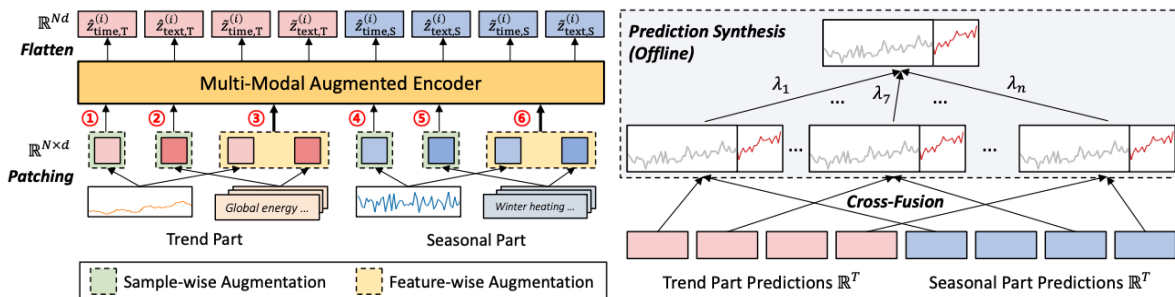


Figure 3 Two key components of our approach: (a) Multi-modal augmented encoder, and (b) Cross-modal fusion and offline synthesis [1].

We use the joint time series and text-based prediction as an example to illustrate our method, shown in Figure 1, which includes two main stages: (a) Multi-modal augmented encoder, and (b) Cross-modal fusion and offline synthesis. In the first stage, a lightweight STL (Seasonal and Trend decomposition using Loess) is applied to improve the predictability of raw time series data. The multi-modal augmented encoder is then utilized across six distinct input patches, labeled 1 to 6, which are obtained by augmenting multi-modal inputs either in a sample-wise or feature-wise manner. Note that inputs 1 to 6 are processed independently with the shared encoder. In the sample-wise augmentation, each modality's sliding window is treated as an independent sample. In this way, it expands the training set and combats data scarcity. In the feature-wise augmentation, we concatenate concurrent patches from different modalities before encoding. This exposes the encoder to joint cross-modal patterns [1].

Given the representations from multi-modal augmented encoder, the second stage in our method generates multiple predictions through cross-fusion. These predictions are then aggregated to generate the final prediction using an offline synthesis approach, as shown in Figure 1(b).

3 GENERATIVE AI FOR AUTOMATED NEGOTIATION

The advent of Generative AI heralds a new era for automated negotiation systems, offering capabilities that extend far beyond the remit of traditional approaches. By leveraging sophisticated models capable of understanding, generating, and reasoning about complex information, Generative AI unlocks a suite of applications designed to make automated negotiations more intuitive, adaptive, and ultimately, more effective. These applications address core challenges in both the human-computer interaction aspects of negotiation and the complex analytical underpinnings required for optimal decision-making in dynamic environments.

One of the most immediate and impactful applications of Generative AI is in **providing more natural and intuitive interfaces for human interaction with automated negotiation systems**. Traditional systems often require users to define strategies and preferences through rigid formalisms or complex programming. Generative AI, particularly through Large Language Models (LLMs), can facilitate interactions via natural language. This allows human negotiators to set objectives, convey constraints, query the system's reasoning, or even participate directly in a mixed-initiative negotiation using everyday language, significantly lowering the barrier to entry and enhancing user trust and adoption.

Beyond simplifying interfaces, Generative AI offers a superior method for **encoding and integrating nuanced human intuition and domain expertise into automated negotiation strategies**. Experienced human negotiators rely on a wealth of tacit knowledge, subtle cues, and qualitative assessments that are notoriously difficult to capture in explicit rules or mathematical functions. Generative models can be trained on historical negotiation transcripts, expert feedback, and best practice guidelines, allowing them to learn and replicate sophisticated human-like reasoning and strategic thinking. This enables the creation of negotiation agents that can employ more flexible, creative, and contextually appropriate tactics, mirroring the adaptability of skilled human counterparts.

A second major domain where Generative AI provides transformative potential is in **the evaluation and representation of preferences, often encapsulated as utility functions**. In many industrial applications, particularly procurement and supply chain management, defining an agent's utility function as a simple mathematical equation over the negotiation outcome space is an oversimplification. Real-world preferences are frequently complex, multi-dimensional, and highly dependent on a wide array of dynamic external factors such as fluctuating market conditions, evolving production plans, real-time sales forecasts, and the projected availability of raw materials, making static or easily defined functions inadequate.

Generative AI can address this complexity by **providing a more robust and adaptive framework for evaluating these intricate utility functions**. Instead of relying on predefined mathematical models, generative systems can learn to approximate or dynamically construct preference landscapes by analyzing vast amounts of diverse data. This includes historical data on past deals, current market intelligence reports, internal operational data, and even external news feeds. By identifying patterns and correlations within this data, Generative AI can develop a more holistic and context-aware understanding of what constitutes a "good" outcome under specific, evolving circumstances.

Furthermore, these advanced preference evaluation capabilities can be significantly enhanced using **multi-modal time-series analysis integrated within Generative AI models**. Industrial utility functions are not static; they evolve with the continuous inflow of new information from various sources over time. Generative AI can process and synthesize these diverse data streams—such as textual market sentiment analyses, numerical production schedules, and categorical supplier risk assessments—across different time horizons. This allows the negotiation agent to dynamically adjust its understanding of preferences and trade-offs, leading to more agile and strategically sound agreements that reflect the true, often implicit, objectives of the organization in a rapidly changing operational landscape.

The following subsections provide more details about each of these applications of Generative AI in automated negotiation.

3.1 GENERATIVE AI FOR INTERACTING WITH HUMANS

In an agent that negotiates with humans, generative AI can be used for two purposes: (A) interpretation of negotiation content and (B) creation of negotiation expressions. The overall structure and flow are as follows:

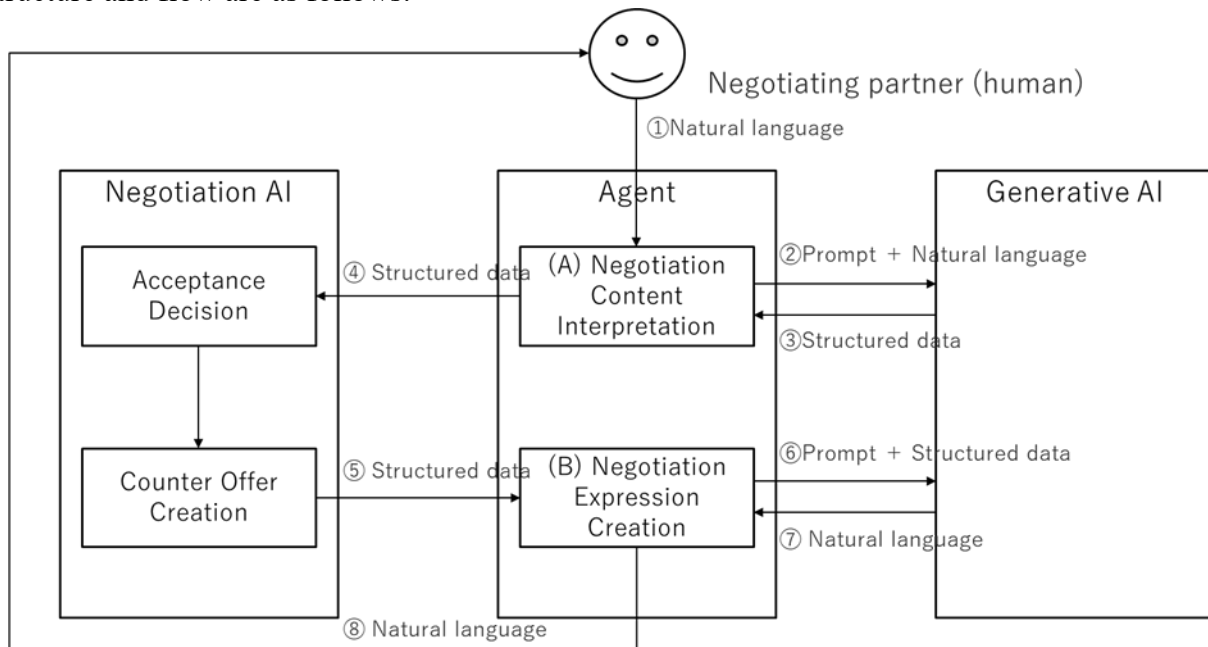


Figure 4 Structure of Generative AI-Powered Negotiation Agent

1. Receive the offer content from the negotiation partner (human) in natural language.
2. Provide a prompt requesting conversion into structured data interpretable by the negotiation AI, along with the offer content in natural language, to the generative AI.
3. Receive the offer content converted into structured data from the generative AI.
4. Provide the offer content from the negotiation partner to the negotiation AI in structured data.
5. Receive the offer content from the negotiation AI in structured data.

6. Provide a prompt requesting conversion into natural language, along with the offer content in structured data, to the generative AI.
7. Receive the offer content converted into natural language from the generative AI.
8. Deliver the offer content from the negotiation AI to the negotiation partner in natural language.

(A) Negotiation Content Interpretation Function

Extracts necessary information from the words of the negotiation partner (human) and communicates it to the negotiation AI. This allows the negotiation partner to negotiate with the AI using traditional communication methods.

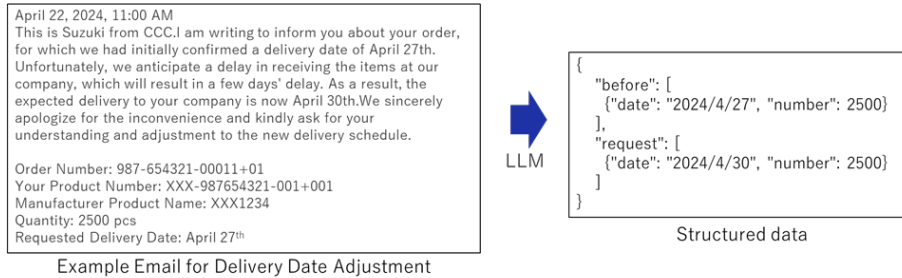


Figure 5 Example of Extraction from Email (based on real data).

(B) Negotiation Expression Creation Function

Enhances the negotiation expressions of the AI. Instead of using uniform, template-like expressions, it allows for more flexible and natural expressions depending on the negotiation situation (e.g., with gratitude, apologetically, confidently), which helps to build a good relationship with the partner and reach an agreement under better conditions.




Figure 6 Example of output generation for Human-Agent negotiation using Generative AI

The method proposed in Figure 4 is only one possible approach for enhancing interaction with the owner and partners of the automated negotiation agent. Other approaches have been proposed. For example Priya et. al. proposed GENTEEL-NEGOTIATOR, an LLM-enhanced method for polite negotiation dialogue [4].

3.2 GENERATIVE AI FOR REASONING DURING NEGOTIATION

The capacity for nuanced reasoning communicated through natural language is a hallmark of skilled human negotiators, enabling them to explore complex trade-offs, build rapport, and uncover hidden value. Traditional automated negotiation systems, reliant on predefined scripts or

formal logic, largely lack this crucial capability. Generative AI, particularly through the advancements in Large Language Models (LLMs), offers a paradigm shift by enabling automated agents to engage in negotiation dialogues using human-like language. This facilitates not only more accessible interaction but also unlocks sophisticated reasoning capabilities that were previously unattainable for machines.

However, the power of LLMs in negotiation extends far beyond the superficial ability to merely parse and generate grammatically correct sentences. The true revolution lies in the fact that these models, through their extensive training on vast textual datasets, develop intricate internal representations—often referred to as embeddings—that encode complex semantic relationships between words, phrases, and concepts. This means an LLM doesn't just process "price" as a string of characters; it has a learned understanding of its connection to concepts like "cost," "value," "discount," "budget," and "affordability," and how these relationships shift across different contexts.

This embedded conceptual knowledge is instrumental when it comes to **reasoning about offers received during a negotiation**. When an LLM processes an offer articulated in natural language, it leverages these learned associations to interpret the proposal's deeper implications. For instance, it can infer unstated assumptions, recognize subtle bargaining tactics, or understand the potential impact of a specific clause based on its relationship to broader business objectives learned from its training data. An offer proposing a "slightly extended delivery timeline in exchange for a volume discount" is understood not just as a set of terms, but as a trade-off leveraging the conceptual link between logistical flexibility and cost savings.

Consequently, the ability to generate **appropriate and strategically sound responses** is significantly enhanced. Armed with this nuanced understanding of the incoming offer and its own encoded knowledge of negotiation strategies and desirable outcomes, the LLM can formulate replies that are contextually relevant, persuasive, and aligned with overarching goals. It can craft counteroffers that propose alternative solutions by drawing upon related concepts, ask clarifying questions that probe specific conceptual ambiguities, or provide justifications for its positions that resonate with the other party's likely understanding, all articulated in coherent and persuasive natural language. This allows for a more dynamic and intelligent negotiation process that more closely mirrors human capabilities.

3.3 GENERATIVE AI FOR UTILITY FUNCTION EVALUATION

The third application of Generative AI in automated negotiation is solving the evaluation problem as described earlier. This solution depends on the specific computational structure of the utility function. In this section, we describe a common use-case in industrial applications in which Generative AI is used for multi-modal timeseries forecasting is used for utility calculation.

Generative AI offers several innovative approaches for utility function evaluation in automated negotiation systems. First, generative models can synthesize utility functions from heterogeneous data sources, such as historical negotiation records, market trends, and contextual business information. By learning from these diverse datasets, generative AI can construct utility functions that more accurately reflect real-world preferences and constraints, even when explicit mathematical formulations are unavailable or incomplete. This enables negotiation agents to

adapt their strategies dynamically as new information becomes available, resulting in more effective and context-aware decision-making.

Second, generative AI can be used to simulate and predict the likely utility functions of negotiation partners. By analyzing patterns in past interactions and external signals, generative models can generate plausible estimates of counterparties' preferences and reservation values. This capability allows agents to anticipate the responses of their negotiation partners, identify mutually beneficial agreements, and avoid deadlocks. Such predictive modeling enhances the agent's ability to propose creative solutions and reach agreements that maximize joint value.

Third, generative AI facilitates the continuous refinement of utility functions through feedback and reinforcement. As negotiation agents interact with human users or other agents, generative models can incorporate feedback from negotiation outcomes to update and improve their utility function representations. This iterative learning process ensures that the utility functions remain aligned with evolving business objectives and stakeholder preferences, supporting long-term optimization and sustained negotiation performance.

Several works follow this thread of research. For example, Bakhtin et. al. showed that LLM-based negotiators can achieve human-level negotiation ability in the well-known diplomacy negotiation game [5]. Kwon et al. provide a systematic evaluation of LLMs' ability as negotiators [6]. Schneider et al. analyze the skill gap and reasoning deficits of LLMs in the area of automated negotiation [7].

4 USE-CASES

In the previous section, we described three areas in which Generative AI can be used to support Automated Negotiation. The first two (i.e. interaction with human users and partners, and reasoning during negotiation) have been proposed perviously and we already referenced some recent works that showcases each of them. Due to lack of space, we will not describe any use-cases that showcase these two approaches here. The following two subsections provide use-cases for the third approach (i.e. using GenerativeAI to improve utility calculations) which are, to the best of our knowledge, the first applications of Generative AI for this purpose. These uses-cases use real-world data to ground the forecasting process but does not describe a deployed system.

4.1 SUPPLY-DEMAND MATCHING

Example

Supply-Demand Matching:

In this use-case, Generative AI is used to solve the evaluation problem (Section 3.3) using multi-modal time-series forecasting (Section 2.3) leading to better negotiation outcomes (Figure 12). This section details the processing stages highlighting the role of Generative AI and Automated Negotiation.

Generative AI for Automated Negotiation

Consider a trading company that buys and sells some product p . To maximize its profit, the company needs to match its supply and demand. Buying too much leads to increased risk and inventory carrying cost. Buying too little leads to lost business opportunities or even penalties.

Figure 8 shows the processing steps involved in this use-case. The first processing step uses historical demand data and related textual information curated from the web (e.g. tweets, economic reports, company return statements, etc) to provide an accurate forecast of future demand of the product p for some prediction horizon T . This demand time-series is called d hereafter with $d(t)$ representing the predicted demand on future time-step t . Figure 8 shows an example of this process. On the left, we can see the historical sales of Petroleum in the past year. The vertical line marks the 10th day in the past for which we have 8 news items (bottom of Figure 8). We also show one of these news items which happened to be a report that oil markets are moving from bull to bear. This kind of report is relevant both to Petroleum price and demand. On the right, we see the future demand generated by our multimodal time-series forecast.



Sample from 8 news items on 2025-05-26 (10 days ago): A couple of months ago, the bulls ran the oil markets. Now the bears have taken over.

Figure 8 Using Generative AI to improve demand forecasting based on multimodal time-series forecasting (Direct forecasting use-case).

Note the increases in demand that happens around the second and 8th day. Our forecasting method could predict these demand changes albeit with some delay.

Figure 9 shows the results of forecasting using the same underlying method but without employing Generative AI for analysis of news articles. In this case both increases were not correctly predicted. This example shows that, at least in this case, our proposed Generative AI enhanced multimodal time-series forecasting method can improve the final demand forecast which, as we will show later, provides a better utility estimate and ultimately leads to higher profits through better supply-demand matching.




Figure 9 Forecasting without Generative AI.




Figure 10 Using enhanced future demand prediction to generate target procurement plan. Note that the target plan differs from the existing set of agreements which leads to negotiation.

The second stage in the processing pipeline is shown in Figure 10. In this stage, the demand forecast generated in the previous stage is used to find the optimal procurement plan for the following seven procurement windows. On the left we show the demand forecast used and on the right, we show the resulting procurement plan in blue as well as the existing procurement contracts in yellow (based on Purchase Orders received from the customer). Moreover, we show the ideal procurement plan assuming perfect knowledge of future demand (red) and the procurement plan assuming a constant demand level (i.e. without forecasting). The fact that predicted demand does not match the procurement schedule based on existing POs (i.e. that the yellow and blue lines do not match) creates a demand-supply matching problem that we use automated negotiation to solve. In this case, we should reduce the total quantity delivered from 1500 to 1200 and change the delivery schedule to achieve this matching.



Figure 11 Negotiating for the new target procurement plan leading to lower carried inventory while guaranteeing all needs for manufacturing based on the new demand forecast.

Figure 11 shows the results of running an automated negotiation using our Reinforcement Learning based negotiator to achieve this supply-demand matching. The optimal procurement plan based on the forecast is shown in blue, the procurement schedule before negotiation is shown in yellow and the procurement plan based on the final agreement is shown in green. In this case, our method could match the optimal plan almost perfectly in 957ms after exchanging 684 offers with the supplier.




Figure 12 The achieved reduction in carried inventory based on negotiation alone and combining negotiation with Generative AI based multimodal time-series forecasting.

Figure 12 shows the added value from combining automated negotiation with Generative AI. These values are calculated by simulating the whole pipeline described in Figure 7. The KPI we use for evaluation here is the saving in carried inventory. On the right (red) is the ideal case assuming full knowledge of the future and a supplier that accepts any POs and executes them perfectly. In this unrealistic case, we could reduce the carried inventory by 31% compared with the case in which we just carry on with the POs without any negotiation (the zero value on the left). When applying automated negotiation using a simple forecasting method that just assumes that the demand will stay at its mean demand level in the past month, we get 10% reduction in carried inventory (gray). This shows that automated negotiation can add value to this procurement problem even without using Generative AI. When applying the three stages presented in Figure 7 with Generative AI enhancing the forecasting (first) stage and automated negotiation used to solve the matching problem (third), we can achieve 24% reduction in carried inventory (14% higher than automated negotiation alone and only 7% lower than the ideal reduction with full future knowledge and a perfect supplier).

So far, we used one specific example to show a concrete use-case in which Generative AI is used to improve utility calculation (through better demand forecasting) solving the evaluation problem. To provide a more rigorous evaluation of the proposed approach, we used the following Fuel and Metal datasets (See Annex A.1).

Table 1 shows the results for four conditions: Ideal (assumes full knowledge of the future and a perfect supplier), Negotiation (uses our automated negotiation only), Negotiation+Forecasting (using our automated negotiation with Dlinear as the forecasting engine), Negotiation+Gen AI enhanced Forecasting (using our automated negotiation with our proposed forecasting method employing GenAI). As described earlier, these values are calculated using simulation of the complete proposed pipeline (Figure 7) on real data (See Annex A.1). In two out of four products, the proposed method achieves a perfect score and in the other two it achieves the highest score among other variations. Moreover, the negotiation condition always shows some improvement (between 8 and 10% reduction). These results match the example given so far and shows that automated negotiation can add value through supply-demand matching in this domain with Generative AI further enhancing this value.

Table 1 Supply-Demand Matching Use-case. Inventory reduction for 4 products showing the value of the proposed approach.

Product	Ideal	Negotiation	Negotiation + Forecasting	Negotiation+ Gen AI enhanced Forecasting
Petroleum	31%	10%	10%	24%
Gas	34%	10%	19%	29%
Gold	17%	8%	8%	17%
Silver	20%	8%	16%	20%

4.2 INDIRECT MATCHING

Example

Indirect Matching:

In this use-case, also, Generative AI is used to solve the evaluation problem (Section 3.3) using multi-modal time-series forecasting (Section 2.3) leading to better negotiation outcomes (Figure 12). This time we consider a more complicated scenario in which we do not have direct data about demand which requires more complicated business modeling.

This use-case presents a more complicated scenario in which Generative AI-based multimodal time-series forecasting can be used to improve automated negotiation. This use-case uses the same stages shown in Figure 7 but with a slightly more complicated first stage and much more complicated second stage.

In this use-case, we are considering a manufacturer that tries to optimize its procurement plan and negotiate with suppliers to achieve this plan. The main difference between this use-case and the supply-demand matching trading use-case considered in Section 4.1 is the inability to calculate the procurement plan based on a direct forecast of the demand for the product. In this case, a more advanced business-modeling (second) stage is necessary to create a procurement plan based on multiple factors including price and demand forecasts for manufactured products, competition level, and other factors that may affect sales and in turn procurement decisions.

To give a concrete example, we developed a business model that predicts sales based on 20 factors (energy prices, metal prices, ICT stock prices, bitcoin prices, and covid incidents in multiple countries).




Figure 13 The 20 factors used in the business-model in the second use-case. We only show the names of some of these factors.

Other than the two datasets described in the previous use-case, we used the following datasets: Stock-Index, Bitcoin, Covid (See Annex A.1)

Generative AI for Automated Negotiation

Figure 13 shows the forecast of these factors generated by applying our multimodal time-series forecasting method.

The business model is then used to estimate future sales of the manufactured product (Figure 14) and then to create a procurement plan (Figure 15) for each input product based on an optimized manufacturing plan. This process can utilize a digital twin representing the manufacturing plant as well as simulations of customer behavior (e.g. response of demand to Covid cases).

Once this procurement plan is completed, the system uses automated negotiation with the supplier(s) to realize the plan.

Figure 16 shows the results of this use-case. Again, we measure the performance of different methods by the amount of inventory carried. The goal is to reduce this inventory which leads to reduced risk and costs. On the left we see the baseline with no negotiation or forecasting achieving 0 reduction and on the right we see the ideal situation assuming full knowledge of the future and a perfect supplier. In this case 23% reduction in inventory could be achieved. With only automated negotiation (gray), we could achieve an 8% reduction showing again the value of automated negotiation in procurement. Introducing forecasting using Dlinear we can achieve 12% (magenta) reduction, a modest improvement above relying on automated negotiation without a sophisticated forecast for any of the 20 factors. With the addition of Generative-AI enhanced time-series forecasting, we can now reach a reduction level of 19% (blue) which is only 4% below the ideal case.




Figure 16 Results of the second use-case showing the value added by negotiation and Generative AI based utility function modeling.

5 SUCCESS, RISKS AND MITIGATION MEASURES

A robust evaluation of automated negotiation (AN) technologies requires a multifaceted approach that moves beyond simple utility metrics to consider the distinct perspectives of both the individual negotiating agents and the overall system designer. Mohammad [13] provides a comprehensive set of criteria to assess performance from these two critical viewpoints.




Figure 14 Expected sales estimated in the second use-case.



Figure 15 Procurement plan for one input product.

From the perspective of a single agent participating in a negotiation, success is not just about maximizing its own outcome. A holistic evaluation includes three key factors:

- **Advantage:** The agent's primary goal is to maximize the value it receives above its minimum acceptable outcome (its "reserved value").
- **Partner Welfare:** For long-term strategic relationships, it is also beneficial to consider the advantage gained by negotiation partners, which can increase the likelihood of future trade.
- **Privacy:** A crucial consideration is the ability to protect sensitive information. This metric measures the fraction of an agent's preferences that are *not* revealed through the offers it makes during the negotiation.

From the viewpoint of the platform designer or a neutral third party, the goal is to evaluate the overall health, efficiency, and fairness of the negotiation mechanism itself. This involves the following measures:

- **Completeness:** The fraction of negotiations that successfully reach a "win-win" deal when such an outcome is possible. This is a key measure of the system's ability to avoid deadlocks and facilitate agreements.
- **Welfare:** The total sum of value achieved by all agents, measured relative to the maximum possible value that could have been achieved in an ideal agreement.
- **Optimality:** A measure of economic efficiency that calculates how close the final agreement is to the Pareto-optimal frontier—the set of outcomes where no party can be made better off without making another worse off.
- **Fairness:** This metric assesses how equitably the gains are distributed among participants by measuring the outcome's distance to well-known bargaining solutions, such as the Nash or Kalai-Smorodinsky solutions.

While Generative AI offers transformative potential for automated negotiation, its deployment is not without significant risks, chief among them being the phenomenon of "**hallucination.**" LLMs, by their nature as probabilistic language predictors, can generate outputs that are fluent and plausible-sounding but are factually incorrect, logically inconsistent, or entirely fabricated. In the high-stakes context of business negotiations, such hallucinations pose a critical threat. An agent acting on hallucinated information could agree to disadvantageous terms, misrepresent a company's position leading to legal liabilities, or damage long-term partner relationships by communicating nonsensical or untruthful information. The risk is amplified because these models are designed for coherence, not truthfulness, making their errors difficult to detect without rigorous verification.

This risk of hallucination directly impacts the three key integration points discussed in this paper. First, when interacting with humans, an LLM could misinterpret a natural language offer from an email, hallucinating a delivery date or quantity that was never stated and passing this incorrect structured data to the negotiation AI. Conversely, when creating a response, it might fabricate details or promises the company cannot honor. Second, during reasoning, an LLM's understanding is based on learned conceptual associations. It might hallucinate a connection between two terms—for example, incorrectly inferring that a "volume discount" automatically

implies "free shipping"—leading it to craft a suboptimal counteroffer based on a flawed premise. Finally, for utility function evaluation, a model performing multimodal time-series analysis could hallucinate correlations, inventing trends from news articles or market data that do not exist, resulting in a completely inaccurate demand forecast and a flawed utility calculation. The negotiation agent would then optimize for a reality that doesn't exist.

To mitigate these risks, we propose a Constrained Generation and Verification Loop architecture rather than allowing the LLM to operate with full autonomy. For interpretive tasks, the LLM's role should be constrained to extracting information into a rigid, predefined schema, with instructions to return a null value for any field not explicitly found in the source text. For generative tasks, a verification layer is essential. Before any natural language communication is sent to a human partner, a simpler, rule-based system must verify that all critical data points (e.g., prices, dates, quantities) in the generated text perfectly match the structured data produced by the core negotiation AI. For the most critical applications, especially utility evaluation based on forecasting, a "human-in-the-loop" system is necessary, where forecasts and their resulting procurement plans are flagged for human review and approval before the negotiation agent is authorized to act on them. This layered approach grounds the LLM's powerful capabilities in a framework of verifiable constraints and human oversight, harnessing its strengths while minimizing the potential for costly errors.

6 CONCLUSION

This paper explores the transformative potential of Generative AI in automated negotiation, particularly within business and industrial contexts. Traditional automated negotiation systems, grounded in game-theoretic and decision-theoretic approaches, often suffer from rigidity and lack of context-awareness, limiting their effectiveness in dynamic, real-world environments. By leveraging Generative AI, especially large language models and multimodal time-series forecasting, the authors propose more flexible, adaptive, and context-sensitive negotiation protocols and strategies that can better accommodate complex preferences, evolving market conditions, and nuanced human interactions.

A key contribution of the paper is the integration of Generative AI for three main purposes: facilitating natural language interaction between humans and negotiation agents, enhancing reasoning capabilities during negotiation, and improving the evaluation of utility functions through multimodal data analysis. The authors present concrete use-cases in procurement and supply-demand matching, demonstrating how Generative AI-powered forecasting and preference modeling can lead to more accurate demand predictions, optimized procurement plans, and significant reductions in inventory costs. These results are validated using real-world datasets spanning fuel, metals, stock indices, and other economic indicators.

Looking forward, the paper suggests several promising directions for future research and development. These include expanding the use of Generative AI to meta-negotiation tasks such as partner and protocol selection, further refining utility function modeling with richer and more diverse data sources, and developing more robust, trustworthy, and explainable negotiation agents. Additionally, the integration of digital twins, advanced simulation environments, and

continuous feedback mechanisms could further enhance the adaptability and business value of automated negotiation systems powered by Generative AI.

7 REFERENCES

- [1] Mohammad, Y. (2023). Generalized Bargaining Protocols. In *Australasian Joint Conference on Artificial Intelligence* (pp. 261-273). Singapore: Springer Nature Singapore.
- [2] Aydoğan, R., Festen, D., Hindriks, K. V., & Jonker, C. M. (2017). Alternating offers protocols for multilateral negotiation. *Modern approaches to agent-based complex automated negotiation*, 153-167.
- [3] Baarslag, T., & Baarslag, T. (2016). A Component-Based Architecture to Explore the Space of Negotiation Strategies. *Exploring the Strategy Space of Negotiating Agents: A Framework for Bidding, Learning and Accepting in Automated Negotiation*, 53-69.
- [4] Priya, P., Chigrupaatii, R., Firdaus, M., & Ekbal, A. (2025, April). GENTEEL-NEGOTIATOR: LLM-Enhanced Mixture-of-Expert-Based Reinforcement Learning Approach for Polite Negotiation Dialogue. In *Proceedings of the AAI Conference on Artificial Intelligence* (Vol. 39, No. 23, pp. 25010-25018).
- [5] Meta Fundamental AI Research Diplomacy Team (FAIR)[†], Bakhtin, A., Brown, N., Dinan, E., Farina, G., Flaherty, C., ... & Zijlstra, M. (2022). Human-level play in the game of Diplomacy by combining language models with strategic reasoning. *Science*, 378(6624), 1067-1074.
- [6] Kwon, D., Weiss, E., Kulshrestha, T., Chawla, K., Lucas, G. M., & Gratch, J. (2024). Are llms effective negotiators? systematic evaluation of the multifaceted capabilities of llms in negotiation dialogues. *arXiv preprint arXiv:2402.13550*.
- [7] Schneider, J., Haag, S., & Kruse, L. C. (2024, November). Negotiating with llms: Prompt hacks, skill gaps, and reasoning deficits. In *International Conference on Computer-Human Interaction Research and Applications* (pp. 238-259). Cham: Springer Nature Switzerland.
- [8] Mohammad, Y. (2021). Concurrent local negotiations with a global utility function: a greedy approach. *Autonomous Agents and Multi-Agent Systems*, 35(2), 28.
- [9] Mohammad, Y., & Nakadai, S. (2019, May). Optimal value of information based elicitation during negotiation. In *Proceedings of the 18th international conference on autonomous agents and multiagent systems* (pp. 242-250).
- [10] Morinaga, S., Nakadai, S., Ando, T., Imura, H., & Takahashi K. (2023). Negotiation Automation Platform, *OMG Journal of Innovation (23rd Edition)* (pp 50-81).
- [11] Jian Y., Ning, K., Pan, Z., Shen, X., Ni, J., Yu, W., ... & Song, D. (2025). Multimodal Time Series Analysis: A Tutorial and Survey, *arXiv preprint arXiv: 2503.13709*.
- [12] Lee, G., Yu, W., Shin, K., Cheng, W., & Chen, H. (2025, April). Timecap: Learning to contextualize, augment, and predict time series events with large language model agents. In *Proceedings of the AAI Conference on Artificial Intelligence* (Vol. 39, No. 17, pp. 18082-18090).
- [13] Mohammad, Y. (2023, December). Evaluating Automated Negotiations. In 2023 IEEE International Conference on Agents (ICA) (pp. 77-82). IEEE.

Annex A

A.1 DATASET USED IN THIS PAPER

1. **Fuel** is a monthly dataset consisting of gas 1 and oil 2 prices, spanning from January 2000 to September 2022. For each month, we extracted news articles from The New York Times 3 that contain relevant keywords {brent crude, crude oil, energy policy, gas demand, gas market, gasoline price, natural gas, oil demand, oil market, oil price, OPEC}, as suggested by ChatGPT.
2. **Metal** is a monthly dataset consisting of gold 4 and silver 5 prices, spanning from January 2000 to August 2022. For each month, we extracted news articles from The New York Times that contain relevant keywords {coinage, COMEX, currency strength, exchange market, exchange rate, Fort Knox, gold and silver, gold coin, gold industry, gold market, gold mining, gold price, gold reserves, gold silver, gold standard, hedging, inflation, karat, LBMA, mining output, precious metal, quantitative easing, recession, safe-heaven asset, silver coins, silver industry, silver institute, silver market, silver price, sterling silver, supply chain disruptions, world gold council}, as suggested by ChatGPT.
3. **Stock-Index** is a monthly dataset consisting of the commodity price index 6 ranges from March 2010 to February 2022. The news articles related to each month are gathered from S&P Global Commodity Insights. Stock-Tech is a weekly dataset of Microsoft (MSFT) and Apple (AAPL) stock prices ranging from December 1, 2006 to November 30, 2016. For each week, we aggregated news articles featuring these two companies retrieved from The New York Times API 7.
4. **Bitcoin** is a daily dataset consisting of Bitcoin (BTC), Ethereum (ETH), Tether (USDT), and Binance Coin (BNB) prices spanning from November 13, 2017 to November 23, 2019. For each day, we filtered tweets about Bitcoins 9 that received at least 100 likes and 50 retweets.
5. **Covid** is a daily dataset consisting of the number of confirmed COVID-19 cases in 10 countries: China, the United States, Italy, Singapore, South Korea, Georgia, Japan, Canada, Russia, and Australia 10, spanning from January 22, 2020 to July 27, 2020. For each day, we used news content that is related to Coronavirus which was aggregated, analyzed, and enriched by AYLIEN using the AYLIEN's News Intelligence Platform. Among a large number of articles, we filtered 20 articles per day based on the number of times the articles were shared on social media platforms including Facebook, LinkedIn, Reddit, and Google Plus.

GLOSSARY

Alternating Offers Protocol (AOP) A common negotiation protocol where participants take turns proposing offers. An offer can be accepted, ending the negotiation, or rejected, allowing the other party to make a counteroffer.

Automated Negotiation (AN) A field of artificial intelligence where software agents, representing individuals or organizations, interact to reach mutually acceptable agreements.

BOLA Architecture A framework for designing a negotiation agent's strategy, consisting of four components: a **Bidding** policy (what to offer), an **Opponent** model (learning about the partner), a **Leaving** policy (when to exit), and an **Acceptance** policy (when to agree).

Digital Twin A virtual model of a real-world physical system or process. In this paper, it is used to simulate manufacturing and customer behavior to inform procurement decisions.

Generative AI A class of artificial intelligence models that can create new, original content, such as text, images, or data, based on patterns learned from existing data.

Hallucination (AI) A phenomenon where a generative AI model produces output that is nonsensical, factually incorrect, or entirely fabricated, despite being presented in a confident and plausible manner.

Large Language Model (LLM) A type of generative AI specifically trained on vast amounts of text data to understand, generate, and reason about human language.

Multi-agent System A computerized system composed of multiple interacting, autonomous software agents. Automated negotiation systems are a type of multi-agent system designed for coordination and agreement-making.

Multimodal Time-Series Forecasting A technique for predicting future values in a sequence by analyzing historical data from multiple sources and types (modalities), such as numerical data (e.g., sales figures) and unstructured text (e.g., news articles).

Negotiation Protocol The set of rules that govern the interaction between negotiating agents. It defines what actions are allowed, in what order, and how an agreement is reached or the process terminates.

Negotiation Strategy The decision-making logic an agent employs within the rules of a protocol to achieve its objectives. This includes what offers to make, when to accept an offer, and when to walk away.

Outcome-Space The complete set of all possible agreements that can be reached in a negotiation. Each potential agreement is a single "outcome."

Preference Elicitation The process of determining and formalizing the preferences of a user or stakeholder to construct their utility function for the negotiation.

Reservation Value The minimum acceptable utility an agent is willing to accept from an agreement. Any offer with a utility below this value will be rejected, as the agent prefers to walk away (disagree).

Time-Series Forecasting A statistical method used to predict future values based on previously observed data points arranged in chronological order. In the context of this paper, it is used to forecast future demand or other control variables to inform utility calculations.

Utility Function A mathematical representation of a negotiator's preferences. It maps every possible outcome in the outcome-space to a numerical score (utility), allowing the agent to quantitatively compare different potential agreements.

Von Neumann–Morgenstern Utility Theorem A fundamental theorem in decision theory which states that under specific axioms of rational behavior (completeness, transitivity, continuity, and independence), a negotiator's preferences can be represented by a utility function. This guarantees that a numerical value can be assigned to any negotiation outcome, which the agent seeks to maximize.

ACKNOWLEDGEMENTS

The views expressed in the *OMG Journal of Innovation* are the author's views and do not necessarily represent the views of their respective employers nor those of the Object Management Group® (OMG®), an EDM Association community.