

Automated Negotiation and Multimodal Time-Series Forecasting for Efficient Procurement*

Demonstration Track

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

Haifeng Chen
 NEC Labs America
 Princeton, USA
 haifeng@nec-labs.com

ABSTRACT

Procurement is a key function in supply chain management that involves acquiring goods and services to meet organizational needs. Efficient procurement is crucial for minimizing costs, ensuring timely delivery, and maintaining quality standards. This paper explores the integration of automated negotiation and multimodal time-series forecasting to enhance procurement processes. Automated negotiation can streamline interactions with suppliers, while multimodal time-series forecasting can improve demand prediction accuracy by leveraging diverse data sources leading to better negotiation outputs. By combining these approaches, organizations can optimize procurement strategies, reduce costs, and improve overall supply chain efficiency. We present two case studies using simulations based on real-world data for procurement that show the effectiveness of the proposed framework.

KEYWORDS

Automated Negotiation; Time-series forecasting; Procurement

ACM Reference Format:

Yasser Mohammad and Haifeng Chen. 2026. Automated Negotiation and Multimodal Time-Series Forecasting for Efficient Procurement: Demonstration Track. In *Proc. of the 25th International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2026), Paphos, Cyprus, May 25 – 29, 2026*, IFAAMAS, 3 pages. <https://doi.org/10.65109/VKYZ9649>

1 APPLICATION DOMAIN

Procurement is one of the core functions in supply chain management. It involves the process of securing goods and services from external suppliers to meet the needs of an organization. Efficient procurement is essential for minimizing costs, ensuring timely delivery, and maintaining quality standards. Traditional procurement methods often involve manual negotiations with suppliers, which can be time-consuming and prone to errors. One important sub-problem of procurement is Delivery Date Adjustment (DDA) with suppliers. DDA operations are short-term (usually daily or weakly) negotiations between the procurement manager and suppliers to adjust the delivery dates of goods in order to match demand in a

*Demo video: <https://youtu.be/kDkVYAHQpFo>



This work is licensed under a Creative Commons Attribution International 4.0 License.

Proc. of the 25th International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2026), C. Amato, L. Dennis, V. Mascardi, J. Thangarajah (eds.), May 25 – 29, 2026, Paphos, Cyprus. © 2026 International Foundation for Autonomous Agents and Multiagent Systems (www.ifaamas.org). <https://doi.org/10.65109/VKYZ9649>




Figure 1: Proposed method overview

timely fashion while reducing carried inventory. DDA is a complex task that requires balancing multiple factors such as demand forecasts, inventory levels, and supplier capabilities [3]. This is the case especially in industries with long lead times and high demand variability, such as automotive and electronics manufacturing [2].

This demo presents ongoing work that combines automated negotiation with multimodal time-series forecasting to improve Demand Driven Acquisition (DDA) outcomes. Automated negotiation streamlines interactions with suppliers, reducing manual effort, lead-time, and carried inventory. Multimodal time-series forecasting improves demand prediction by leveraging diverse data sources, including historical sales, market trends, and external factors such as news articles and weather forecasts. The proposed approach is not limited to DDA and can be applied to other business tasks in which coordination between multiple players with different objectives is required.

2 AUTOMATED NEGOTIATION

Automated negotiation is a subfield of multi-agent systems and artificial intelligence that focuses on developing algorithms and protocols to enable autonomous agents to negotiate with each other to reach mutually beneficial agreements [9]. Automated negotiation has been applied in various domains, including e-commerce, supply chain management [8], and resource allocation [5].

Formally, a negotiation *scenario* λ is defined as a tuple $(\mathcal{A}, \mathcal{D})$ where \mathcal{A} is the set of agents (also called *negotiators*) numbered from 1 to n , and \mathcal{D} is the negotiation domain. The negotiation domain $\mathcal{D} \equiv (\Omega, \mathcal{U})$ is pair: (1) The outcome space $(\Omega$ of size m) comprising all possible agreements. A special outcome $\phi \notin \Omega$ is always assumed to exist to represent disagreement and we define the *extended outcome space* Ω^+ as the $\Omega \cup \{\phi\}$. (2) \mathcal{U} is a tuple of agent utility functions. Each agent utility function u_i is a mapping per agent from Ω^+ to the range $[0, 1]$. Time-pressure can be modeled by a discounting factor as in [9] or by a limit on the number of rounds/seconds allowed for the negotiation [1]. We assume

that each agent knows its utility function but not their opponents' utility functions. This paper uses the bilateral Alternating Offers Protocol (AOP) [1] in which agents exchange offers until either one is accepted leading to an agreement, an agent leaves the negotiation or a time-out condition is reached leading to a failure. For the this work, we trained an RL-based negotiator [10] for DDA. To negotiate effectively, the negotiator needs to know its utility function. In this work, we use future demand prediction through multimodal time-series forecasting to calculate an optimal procurement plan. This similarity between a delivery schedule and this plan is the measure of *utility* employed in this work.

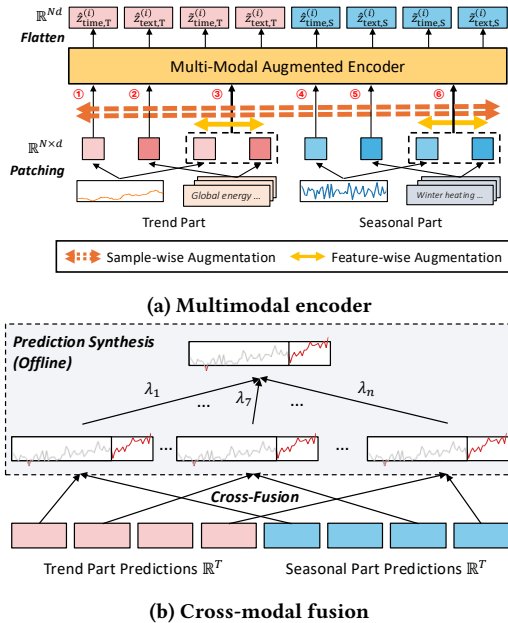


Figure 2: Two key components of multimodal time-series forecasting [4]

3 MULTIMODAL TIME-SERIES FORECASTING

Time series forecasting is crucial in fields like finance, healthcare, energy, and environmental monitoring, as it aids in decision-making, resource optimization, and strategic planning. However, challenges like high noise and limited training data often hinder the accuracy of forecasting. To address this, incorporating information from related data sources, known as *multimodal time series forecasting*, has emerged as a promising solution [4]. Let M be the number of modalities, and define the M -dimensional observation at time t as $\mathbf{x}_t = (x_t^{(1)}, x_t^{(2)}, \dots, x_t^{(M)})$, $t \in [1, L]$, where each $x_t^{(m)}$ comes from a different source such as numeric, text, image, or audio. We use $X_{1:L} = (\mathbf{x}_1, \mathbf{x}_2, \dots, \mathbf{x}_L) \in \mathbb{R}^{L \times M}$ to represent historical data from all modalities. The goal is to predict the future values of a time series using this multimodal data. Assuming $m = 1$ is the numeric target series and the forecasting horizon is T steps, we define the target variable $y_{L+1:L+T} = (x_{L+1}^{(1)}, x_{L+2}^{(1)}, \dots, x_{L+T}^{(1)}) \in \mathbb{R}^T$. We learn a parameterized mapping $F_\theta : X_{1:L} \rightarrow y_{L+1:L+T}$, trained by minimizing a task loss, typically the mean squared error (MSE). The function F_θ

leverages both intra-series dynamics (such as temporal patterns within modality 1) and cross-modal dependencies in $X_{1:L}$ (such as correlations between sensor signals and related text or images) to produce forecasts that outperform single-modality baselines.

Figure 1 illustrates the approach by using the time series and text data as two modality inputs. They are jointly decomposed into trend and seasonal components, and are then fed to the multimodal augmented encoder which incorporates both sample-wise and feature-wise multimodal data augmentation. The resulting outputs from Figure 1(a) are decoded using component-specific decoders, as shown in Figure 1(b). Finally, the predictions for trend and seasonality are combined through a cross-modal fusion scheme to generate the final predictions. For details, see [4].

4 SYSTEM ARCHITECTURE

The proposed method consists of a sequential pipeline (Fig. 2) in which the multimodal time-series forecasting module predicts procurement-plan control variables (e.g. future demand) based on historical data and textual information collected from relevant news sources, social media, etc. These predictions is then used to define the procurement plan through a business-model (this is the only stage that needs to be changed for a new application). The procurement plan defines the **utility function** which is then used to set the utility function used by the automated negotiation module to negotiates with suppliers to adjust delivery dates and quantities to meet the production plan while minimizing costs and risks.

Table 1: Inventory reduction results.

Product	Oracle	Negotiation	+Forecasting	+Multimodal
Petroleum	31%	10%	10%	24%
Gold	17%	8%	8%	17%
Gas	34%	10%	19%	29%
Silver	20%	8%	16%	20%
Indirect	23%	8%	12%	23%

5 EVALUATION

For our first use-case, we considered a trading scenario in which the buyer buys a product from its suppliers and sells it directly to customers (i.e. the business-modeling phase is trivial). We used datasets from [4] and conducted negotiations using a time-based strategy [6] implemented in NegMAS [7]. Performance was measured by inventory reduction. Table 1 shows results comparing: Oracle (perfect demand knowledge), Negotiation-only (mean-based forecasting), +Forecasting (DLinear [11]), and our approach (+Multimodal).

For our second use-case (indirect forecasting), multiple variables affect the procurement plan. We created a synthetic business-model with 20 factors (energy prices, metal prices, ICT stocks, bitcoin, Covid incidents). The last row of Table 1 shows that our method achieves 15% improvement over negotiation-only and matches Oracle performance. These results demonstrate that combining multimodal forecasting with automated negotiation significantly improves procurement outcomes across different scenarios. The approach is particularly effective when external textual information correlates with demand patterns.

REFERENCES

- [1] Reyhan Aydođan, David Festen, Koen V Hindriks, and Catholijn M Jonker. 2017. Alternating offers protocols for multilateral negotiation. In *Modern Approaches to Agent-based Complex Automated Negotiation*. Springer, 153–167.
- [2] Ruomeng Cui, Meng Li, and Shichen Zhang. 2022. AI and procurement. *Manufacturing & Service Operations Management* 24, 2 (2022), 691–706.
- [3] Brian Farrington and Kenneth Lysons. 2020. *Procurement and supply chain management*. Pearson UK.
- [4] Geon Lee, Wenchao Yu, Kijung Shin, Wei Cheng, and Haifeng Chen. 2025. Time-CAP: Learning to contextualize, augment, and predict time series events with large language model agents. In *Proceedings of the AAAI Conference on Artificial Intelligence*, Vol. 39. 18082–18090.
- [5] Cuihong Li, Joseph Giampapa, and Katia Sycara. 2006. Bilateral negotiation decisions with uncertain dynamic outside options. *IEEE Transactions on Systems, Man, and Cybernetics, Part C (Applications and Reviews)* 36, 1 (2006), 31–44.
- [6] Yasser Mohammad. 2021. Optimal Deterministic Time-Based Policy in Automated Negotiation. In *PRIMA 2020: Principles and Practice of Multi-Agent Systems: 23rd International Conference, Nagoya, Japan, November 18–20, 2020, Proceedings 23*. Springer International Publishing, 68–83.
- [7] Yasser Mohammad, Shinji Nakadai, and Amy Greenwald. 2019. NegMAS: A platform for situated negotiations. In *Twelfth International Workshop on Agent-based Complex Automated Negotiations (ACAN2019) in conjunction with IJCAI 2019* (Macau). <https://www.github.com/yasserfarouk/negmas>
- [8] Yasser Mohammad, Enrique Areyan Viqueira, Nahum Alvarez Ayerza, Amy Greenwald, Shinji Nakadai, and Satoshi Morinaga. 2019. Supply Chain Management World. In *International Conference on Principles and Practice of Multi-Agent Systems*. Springer, 153–169.
- [9] Ariel Rubinstein. 1982. Perfect equilibrium in a bargaining model. *Econometrica: Journal of the Econometric Society* 50, 1 (1982), 97–109. <http://www.jstor.org/stable/1912531>
- [10] Ayan Sengupta, Yasser Mohammad, and Shinji Nakadai. 2021. An Autonomous Negotiating Agent Framework with Reinforcement Learning based Strategies and Adaptive Strategy Switching Mechanism. In *Proceedings of the 20th International Conference on Autonomous Agents and MultiAgent Systems* (Virtual Event, United Kingdom) (AAMAS '21). International Foundation for Autonomous Agents and Multiagent Systems, Richland, SC, 1163–1172.
- [11] Ailing Zeng, Muxi Chen, Lei Zhang, and Qiang Xu. 2023. Are transformers effective for time series forecasting?. In *Proceedings of the AAAI conference on artificial intelligence*, Vol. 37. 11121–11128.