



MAFIA-NeT: Multi-Agent Framework for Interactive Agricultural Negotiation and Trading Systems

Dung T. Phan^{1,2†}, Chi N. L. Phan^{1,2†}, Long S. T. Nguyen^{1,2†} ^a, Phuc T. Dao^{1,2}, Quan M. Bui^{1,2},
Tin T. Ngo^{1,2}, Thi T. Nguyen³ and Tho T. Quan^{1,2*} ^b

¹URA Research Group, Faculty of Computer Science and Engineering, Ho Chi Minh City University of Technology (HCMUT), Vietnam

²Vietnam National University Ho Chi Minh City, Ho Chi Minh City, Vietnam

³Agri Sung Joint Stock Company, Vietnam

Keywords: Multi-Agent Systems, Large Language Models, Agricultural Negotiation, Automated Bargaining, E-Agri Commerce, Decision Support.


Abstract: Large Language Models (LLMs) have demonstrated strong reasoning capabilities, yet they remain limited in sustained coordination and reliable real-world operation. These limitations have motivated the development of LLM-powered agents and LLM-based Multi-Agent Systems (MAS), which are better suited for domains that require adaptive reasoning and iterative decision-making. E-Agri commerce is one such domain, as agricultural negotiation involves dynamic pricing, diverse buyer requirements, and rapidly changing market conditions. Existing approaches mainly focus on prediction rather than interactive negotiation, making them insufficient for multi-turn bargaining scenarios. To address this gap, we propose MAFIA-NeT, a multi-agent framework for end-to-end agricultural trade negotiation. The system coordinates specialized agents for structured data parsing, market-informed price reasoning, and strategic decision support. A key contribution is the LLM-Guided Negotiation Subspace (LGNS), which uses an LLM-based anchor mechanism to identify a compact negotiation region and a lightweight regression model to estimate feasible pricing trajectories, enabling efficient negotiation without exhaustive search. MAFIA-NeT was deployed in collaboration with Agri Sung Joint Stock Company and evaluated on authentic B2B negotiation transcripts. Results show improvements in negotiation accuracy, price convergence, and operational practicality, demonstrating the promise of LGNS-guided multi-agent negotiation for scalable agricultural trade automation.


1 INTRODUCTION

Large Language Models (LLMs) have recently shown strong capabilities in reasoning, planning, and multi-step problem-solving (OpenAI, 2024; Chen et al., 2024). Despite this progress, LLMs used as stand-alone systems still face challenges in maintaining long-horizon coherence, coordinating complex sequences of actions, and operating reliably in practical environments (Wang et al., 2024). These limitations have led to interest in *LLM-powered agents*, which are autonomous components capable of perceiving, reasoning, and acting toward goals (Xi et al., 2023). Building upon this paradigm, LLM-based *Multi-*

Agent Systems (MAS) coordinate agents that communicate through structured protocols and leverage the reasoning strengths of LLMs (Botti et al., 2017). Such systems are suitable for domains that exhibit dynamic conditions, incomplete information, and the need for continuous, iterative decision-making.

E-Agri commerce represents one such domain. Agricultural trading is characterized by frequent market fluctuations, seasonal variations, heterogeneous buyer and seller requirements, and persistent uncertainty across supply-chain operations (Singh et al., 2022). As a result, negotiation processes often involve shifting constraints and rapidly changing information, making it necessary to reassess prices, quantities, and deal terms over multiple rounds. Prior work in E-Agri platforms has focused mainly on price forecasting and trend analysis using machine learning methods (Sharma et al., 2023; Reddy and Ku-

^a  <https://orcid.org/0009-0008-7488-4714>

^b  <https://orcid.org/0000-0003-0467-6254>

*Corresponding author.

†These authors contributed equally to this work.

mar, 2022). Although these approaches are valuable for prediction, they do not provide the autonomy, interactivity, or multi-turn reasoning required for negotiation-oriented workflows. Consequently, they fall short in real-world settings where stakeholders must frequently adjust their strategies and reconcile competing objectives.

To address these gaps, we introduce the *Multi-Agent Framework for Interactive Agricultural Negotiation and Trading Systems* (MAFIA-NeT), a multi-agent architecture for fully automated negotiation in agricultural commerce. MAFIA-NeT integrates specialized agents for structured data parsing, market-informed price reasoning, and human-aligned decision support. A key contribution is the *LLM-Guided Negotiation Subspace* (LGNS), which replaces full multi-stage negotiation exploration whose complexity grows rapidly with an LLM-driven anchor mechanism that identifies a compact and strategically meaningful negotiation region. A lightweight regression model then estimates feasible price trajectories within this region, enabling informed bargaining without exhaustive search. This hybrid mechanism shifts negotiation from unguided trial-and-error to a directed, analytically grounded process that improves convergence while preserving adaptability and interpretability. By combining decentralized collaboration, adaptive reasoning, and targeted negotiation-space reduction, MAFIA-NeT enhances responsiveness, stabilizes pricing, and supports scalable deployment in real-world E-Agri commerce environments. The main contributions of this work are as follows.

- We introduce MAFIA-NeT, the first multi-agent framework designed for end-to-end automated negotiation and trading in E-Agri commerce, integrating agents for data aggregation, market reasoning, and negotiation control within a modular and adaptive system.
- We develop the LGNS, which combines an LLM-based anchor mechanism with a regression estimator to prune the negotiation search space and guide price reasoning, resulting in more efficient negotiation trajectories compared with full-space exploration.
- We deploy and evaluate MAFIA-NeT in real-world settings through collaboration with Agri Sung Joint Stock Company¹. Experiments on authentic *Business-to-Business* (B2B) negotiation transcripts demonstrate significant improvements in price convergence, negotiation accuracy, and overall operational practicality.

¹<https://www.agrisung.com/>

2 RELATED WORK

Traditional MAS have long been used to automate complex tasks such as negotiation, recommendation, and decision-making in e-commerce environments. In these systems, agents typically represent stakeholders such as buyers, sellers, or intermediary services. They interact through predefined communication protocols to reach mutually beneficial agreements. Early work employed agent-based models to simulate buyer–seller interactions in B2B scenarios, supporting functions such as price bargaining, order scheduling, and delivery coordination (Koponen and Rytsy, 2020). Later efforts incorporated decision-support logic and integration with enterprise planning tools to optimize negotiation strategies across multiple issues (Calosso et al., 2017). Although these frameworks formalized structured negotiation procedures, they relied heavily on static rules and lacked the contextual flexibility needed in dynamic or uncertain environments.

Subsequent studies enhanced negotiation support through fuzzy logic and preference modeling. Systems such as (Alserhan et al., 2025) and (Haleema and Iyengar, 2022) allowed agents to evaluate qualitative factors, including brand or quality, alongside quantitative variables such as price and quantity. These models demonstrated multi-attribute reasoning and supported both bilateral and multilateral negotiation. However, their strategies remained predetermined, and the rule-based nature of these systems limited their ability to adapt to real-time negotiation dynamics or evolving user behavior.

In parallel, advances in conversational AI and LLM-driven agents shifted the focus to dialog coherence, contextual reasoning, and adaptive decision-making. Recent work includes a multi-agent conversational recommender system that incorporates act-planning and reflection to improve personalization (Fang et al., 2024), LLM-based task planning to decompose complex workflows (De Cicco et al., 2025), and context-based e-commerce assistants capable of citation-oriented responses (Zeng et al., 2025). Despite these developments, most LLM-powered systems are oriented toward single-turn recommendations or informational assistance rather than strategic, multi-turn negotiation.

Research on user behavior and technology adoption in *Small and Medium Enterprises* (SMEs) (Ikumoro, 2019) further highlights organizational and psychological considerations such as perceived usefulness, trust, and social presence (Chen et al., 2021). While these insights inform deployment practices, they do not address the design of adaptive, multi-

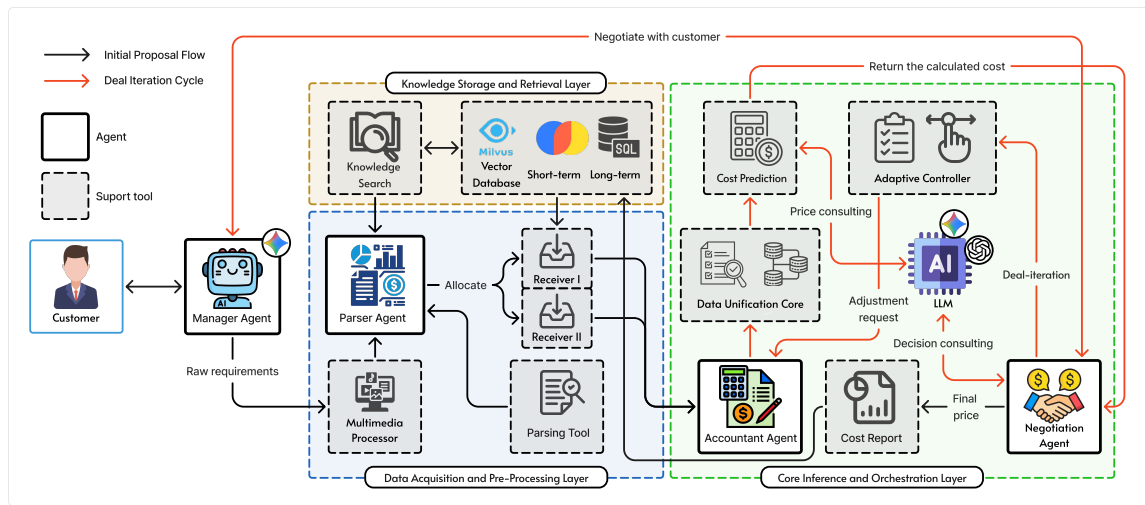


Figure 1: Overall workflow of the MAFIA-NeT architecture, showing data acquisition, knowledge retrieval, and core inference steps coordinated across multiple specialized agents.

agent conversational infrastructures suitable for B2B negotiation scenarios.

In summary, prior work has explored rule-based MAS for structured negotiation and LLM-driven systems for conversational intelligence, yet a gap remains in integrating these two directions to support fully automated, adaptive, and multi-turn negotiation in dynamic agricultural trade environments. Our work addresses this gap by combining MAS coordination with LLM-guided reasoning and negotiation-space reduction through LGNS.

3 MAFIA-NET ARCHITECTURE

3.1 Overview of MAFIA-NeT

MAFIA-NeT is a multi-agent negotiation and pricing framework designed for the E-Agri commerce domain. It integrates LLM-based reasoning with domain-specific tools within a generic agent orchestration paradigm, enabling multiple autonomous agents to interpret multimodal inputs, retrieve relevant knowledge, reason under commercial constraints, and collaboratively generate business-consistent responses. The architecture is organized into three tightly coupled layers: the *Data Acquisition and Pre-Processing Layer*, the *Knowledge Storage and Retrieval Layer*, and the *Core Inference and Orchestration Layer*. An overview of the end-to-end workflow is shown in Figure 1.

Data Acquisition and Pre-Processing Layer. This layer standardizes heterogeneous user inputs and pre-

pares them for downstream reasoning. Incoming content is first handled by the *Manager Agent*, which routes multimodal materials to the *Multimedia Processor*. High-fidelity textual representations are extracted through image-to-text components powered by an LLM service², providing a unified text interface for subsequent processing. The *Parser Agent* then converts these representations into structured requirement fields via a dedicated *Parsing Tool* and enriches them through semantic retrieval from the *Knowledge Search Tool*. Two validation modules, *Receiver I* and *Receiver II*, perform consistency and completeness checks before the data is handed off to the reasoning core.

Knowledge Storage and Retrieval Layer. This layer forms the long-term and short-term memory substrate of the system. A vector database³ supports semantic retrieval over past negotiations, reference documents, and product specifications, while a structured SQL database maintains persistent records such as finalized deals, pricing histories, and contract templates. During active negotiations, session-level information is cached in in-memory SQL tables for low-latency access. The *Knowledge Search Tool* bridges these sources by retrieving high-relevance embeddings and structured fields, ensuring that agents operate over a context that is both semantically rich and up-to-date.

Core Inference and Orchestration Layer. This layer coordinates negotiation reasoning, pricing estimation, and iterative deal refinement through the

²<https://deepmind.google/models/gemini/flash/>

³<https://milvus.io/>

Negotiation Agent, *Accountant Agent*, an LLM-based reasoning core, and control components including the *Data Unification Core*, *Cost Prediction Module*, and *Adaptive Controller*. To constrain negotiation search, the system operates over a compact LGNS, defined in Equation 1.

$$\text{LGNS} = \{\hat{s}_1, \hat{s}_2, \dots, \hat{s}_m\} \subseteq U. \quad (1)$$

Structured requirements from the upstream pipeline are unified with historical pricing data, accounting rules, and contractual constraints to produce an initial *Cost Report*, while the *Cost Prediction Module* estimates feasible price ranges under market and margin constraints. At the decision level, the LLM evaluates candidate adjustments within the LGNS, and the *Adaptive Controller* filters infeasible updates and guides the system toward utility-improving configurations.

Each proposed adjustment is treated as a state transition

$$\delta: s_t \rightarrow s_{t+1}, \quad (2)$$

and iteratively refined within a closed-loop deal-iteration cycle until convergence, after which the *Negotiation Agent* finalizes the agreement and returns validated terms to the customer.

3.2 Price Model

To support negotiation with numerically grounded justifications, MAFIA-NeT employs a hybrid statistical model to estimate the *Total Unit Cost* (TUC) of a product based on material, processing, shipment, and commercial terms. The computation is formalized in Equation 3:

$$\text{TUC} = \left(\text{RM} + \sum_{i=1}^n \text{PS}_i \right) \times (1 + \text{IF}) \times (1 + \text{PF}) + \text{SC}, \quad (3)$$

where:

- **RM**: Raw material cost predicted via a weighted ensemble of *Linear Regression* (LR) and *Bayesian Regression* (BR).
- **PS_i**: Cost of processing step *i*.
- **IF**: Incoterm factor (Ex Works⁴, Free On Board⁵, Cost, Insurance, and Freight⁶, etc.).

⁴The buyer assumes all transportation costs and risks from the seller's premises onward.

⁵The seller delivers the goods onto a vessel at the specified port; risk and cost then transfer to the buyer.

⁶The seller covers cost, insurance, and freight to the destination port; risk transfers when goods are loaded onto the vessel.

- **PF**: Payment factor (advance payment, Letter of Credit⁷, installment terms, etc.).
- **SC**: Shipping cost per unit.

By combining the interpretability of LR with the uncertainty modeling capability of BR, Equation 3 provides cost estimates that are both transparent and robust to noise. These estimates serve as reference signals for the *Negotiation Agent* when deciding whether proposed prices remain within acceptable margin and risk thresholds.

3.3 Agent Cooperation Mechanism

As illustrated in Figure 1, agents cooperate through a structured pipeline of extraction, validation, reasoning, and decision-making:

- **Manager Agent** manages user interaction, maintains dialog context, and routes multimodal inputs to appropriate tools and agents.
- **Parser Agent** performs structured requirement extraction and semantic enrichment, interfacing with the vector store and SQL records for retrieval and structured data access.
- **Accountant Agent** verifies data integrity, computes preliminary valuations based on Equation 3, and prepares detailed cost breakdowns.
- **Negotiation Agent** operates within LGNS by executing state transitions δ in Equation 2, thereby reducing unnecessary negotiation loops and accelerating convergence toward feasible deals.

3.4 Negotiation Process

Before entering the formal negotiation loop, the system invokes *Chain-of-Thought* (CoT) reasoning (Wei et al., 2022) to generate structured, multi-step rationales that support both conflict resolution and price justification. Specifically, CoT operates in two phases: (i) a *conflict-resolution phase*, where the model identifies feasible trade-offs (e.g., Incoterms or payment terms) when target prices violate cost constraints; and (ii) a *value-justification phase*, where the model validates the proposed price based on market signals and cost estimations. During negotiation, the agent adapts dynamically to revised buyer requirements and updates state transitions within LGNS accordingly. The complete decision-making procedure, comprising CoT-driven reasoning and validation, is summarized in Algorithm 1.

⁷A bank-issued guarantee ensuring the seller receives payment upon compliance with documentation requirements.

4 EXPERIMENTATIONS

We conduct two experiments to evaluate the effectiveness of the proposed MAFIA-NeT framework. The first experiment assesses the performance of the pricing prediction module trained on real transactional data. The second examines negotiation effectiveness using annotated real-world B2B dialogues between Agri Sung’s sales agents and international buyers.

4.1 Price Prediction Experiment

4.1.1 Experimental Setup

We evaluate six regression models for agricultural price prediction using Agri Sung’s historical transaction logs. Prior to training, pricing records are pre-

Algorithm 1: Deal Price Computation and Decision Workflow via CoT.

Input: Configuration files, pricing parameters, buyer requirements
Output: Final approved deal price with rationale

while *user requirements are valid* **do**

if *future delivery months specified* **then**

 Forecast and evaluate price accuracy;

else

 Apply latest market prices;

 Compute Base Cost based on specs;

if *Target Price < Computed Cost* **then**

CoT Phase 1: Conflict Resolution;

 Identify gap and propose trade-offs (e.g., Incoterms, payment terms, volume);

if *trade-offs bridge the gap* **then**

 Synthesize rationale for conceded offer;

return offer with trade-off justification;

else

 Mark as infeasible; request re-specification;

else

CoT Phase 2: Value Justification;

 Verify profit margins against market signals;

if *rationale is consistent* **then**

 Synthesize commercial justification;

return final price;

else

 Flag inconsistency; request modification;

processed using the LGNS-based filtering strategy introduced in Section 3, ensuring that only negotiation-relevant attributes are retained while redundant or noisy fields are removed.

Dataset Description. The dataset includes pricing entries recorded between January 2024 and January 2025 from Agri Sung’s operational transaction records across ten agricultural product categories such as nut butters, dried fruits, coffee, and spices. Although modest in size, the dataset captures realistic operational noise, seasonal variability, and market fluctuations, making it suitable for evaluating robustness rather than merely fitting a static corpus.

Model Configuration. We evaluate the following regression models:

- Linear Regression
- Bayesian Ridge Regression
- Polynomial Regression
- Random Forest Regressor
- XGBoost Regressor
- Equal-weight Ensemble (mean of all models)

All models are trained using an 80/20 train–test split, with hyperparameters tuned via grid search and cross-validation where applicable. This configuration allows us to compare simple linear baselines with more expressive tree-based and boosting models under the same data constraints.

Evaluation Metrics. To quantify predictive accuracy, we use four standard regression metrics: *Mean Absolute Error* (MAE), *Root Mean Squared Error* (RMSE), *Mean Absolute Percentage Error* (MAPE), and the *coefficient of determination* (R^2). Let y_i and \hat{y}_i denote the ground-truth and predicted prices, respectively, and let N be the number of test samples. The metrics are defined in Equations 4–7:

$$\text{MAE} = \frac{1}{N} \sum_{i=1}^N |y_i - \hat{y}_i|, \quad (4)$$

$$\text{RMSE} = \sqrt{\frac{1}{N} \sum_{i=1}^N (y_i - \hat{y}_i)^2}, \quad (5)$$

$$\text{MAPE} = \frac{100}{N} \sum_{i=1}^N \left| \frac{y_i - \hat{y}_i}{y_i} \right|, \quad (6)$$

$$R^2 = 1 - \frac{\sum_{i=1}^N (y_i - \hat{y}_i)^2}{\sum_{i=1}^N (y_i - \bar{y})^2}, \quad (7)$$

where \bar{y} is the mean of the ground-truth prices. Together, these metrics capture absolute error, penalization of large deviations, percentage deviation, and variance explained.

4.1.2 Results and Analysis

The performance of all evaluated models is summarized in Table 1. Among them, the *Weighted Ensemble LR-BR* model achieves the strongest overall performance across MAE, RMSE, MAPE, and R^2 .

Table 1 shows that the Weighted LR-BR ensemble yields the lowest MAE and RMSE, the lowest percentage error in terms of MAPE, and the highest R^2 . Although the dataset size constrains the achievable R^2 , the ensemble remains stable across product categories and maintains deviations below three percent from Agri Sung’s internal target prices. In practical terms, this level of accuracy provides a meaningful safety margin when negotiating with buyers, since even small percentage errors can translate into substantial changes in profit margin at scale.

By contrast, Random Forest and XGBoost exhibit negative R^2 scores and substantially higher errors. This behavior indicates overfitting to the limited training data and sensitivity to noisy attributes, even after LGNS-based filtering. The equal-weight ensemble performs better than the worst individual models but is still penalized by aggregating weaker predictors like Polynomial Regression and XGBoost. Overall, the results confirm that a carefully designed weighted combination of simple, well-regularized models is more effective than naively combining diverse regressors.

4.1.3 Discussion

The results in Table 1 and Equations 4–7 show that, in noisy and low-volume industrial settings, simple and well-regularized models outperform more complex ones. The LR-BR weighted ensemble preserves the transparency of linear regression while using a Bayesian prior to stabilize parameter estimates under uncertainty, resulting in stronger performance and a pricing model that is easier to audit. When integrated into MAFIA-NeT, this ensemble provides a reliable cost baseline that guides the Negotiation Agent toward economically consistent offers based on Equation 3 and the error metrics in Equations 4–7. As shown in Section 4.2, this integration improves deal closure rates and alignment with internal pricing expectations.

4.2 User Scenario Experiment

4.2.1 Experimental Setup

We further evaluate negotiation performance using 30 manually annotated transcripts from real-world B2B

interactions between Agri Sung Corp and international buyers. Each transcript captures a complete negotiation episode rather than a single turn, allowing us to assess how well the system operates over multi-step, context-dependent exchanges.

Dataset Description. Each transcript includes structured annotations for:

- buyer preferences (e.g., labeling language, certification requirements, preferred payment method),
- internal target prices defined by Agri Sung’s commercial and pricing teams.,
- final agreed unit prices and terms,
- key negotiation pivot points (for example, where concessions or counteroffers occur).

The annotations were independently validated to ensure consistency in marking deal outcomes and pricing decisions.

Evaluation Metrics. We quantify negotiation performance using four metrics. First, the *Deal Success Rate* (DSR) in Equation 8 measures the proportion of negotiations that conclude successfully:

$$\text{DSR} = \frac{N_{\text{WON}}}{N_{\text{Total}}}, \quad (8)$$

where N_{WON} is the number of negotiations marked as “Deal WON” and N_{Total} is the total number of scenarios.

Second, overall pricing alignment is measured by the *Target-to-Actual Price Gap* (TAPG) in Equation 9:

$$\text{TAPG} = \frac{1}{N_{\text{WON}}} \sum_{i=1}^{N_{\text{WON}}} \left(\frac{P_{\text{final}}^{(i)}}{P_{\text{target}}^{(i)}} - 1 \right), \quad (9)$$

where $P_{\text{final}}^{(i)}$ and $P_{\text{target}}^{(i)}$ denote the final agreed price and the internal target price for deal i , respectively.

Third, we distinguish between underpricing and overpricing by computing directional gaps. Let $S_{\text{neg}} = \{i \mid P_{\text{final}}^{(i)} < P_{\text{target}}^{(i)}\}$ and $S_{\text{pos}} = \{i \mid P_{\text{final}}^{(i)} > P_{\text{target}}^{(i)}\}$. We define:

$$\text{N-TAPG} = \frac{1}{|S_{\text{neg}}|} \sum_{i \in S_{\text{neg}}} \left(\frac{P_{\text{final}}^{(i)}}{P_{\text{target}}^{(i)}} - 1 \right), \quad (10)$$

$$\text{P-TAPG} = \frac{1}{|S_{\text{pos}}|} \sum_{i \in S_{\text{pos}}} \left(\frac{P_{\text{final}}^{(i)}}{P_{\text{target}}^{(i)}} - 1 \right). \quad (11)$$

Deals with zero deviation (final price exactly equals the target price) are included in TAPG (Equation 9) but excluded from N-TAPG and P-TAPG to avoid biasing directional averages.

Table 1: Performance of regression models for price prediction.

Model	MAE ↓	RMSE ↓	R ² ↑	MAPE ↓ (%)
WeightedEnsemble_LR_BR	1759.62	3063.61	0.4394	12.97
LinearRegression	1881.24	3538.18	0.2053	13.62
BayesianRidge	2179.30	3754.45	0.2726	16.37
Ensemble (Equal-weight)	3558.62	4043.55	0.2381	27.19
RandomForest	4102.04	4720.01	-0.0370	29.01
Polynomial	4681.83	5753.15	-0.5920	34.35
XGBoost	5258.91	6485.96	-0.9640	36.94

Table 2: Ablation study on major system components.

System Variant	DSR ↑	N-TAPG ↓	P-TAPG ↓	TAPG ↓
MAFIA-NeT (Full)	0.833	-0.0240	0.0146	-0.0176
w/o Manager Agent	0.400	-0.0543	0.0908	-0.0485
w/o Receiver Modules	0.500	-0.1079	0.0288	-0.0625
w/o Cost Estimator Agent	0.300	-0.3270	0.2460	-0.0722

4.2.2 Results and Analysis

Across 30 negotiation scenarios, 25 resulted in successful agreements, which corresponds to a DSR according to Equation 8. In terms of pricing alignment, we obtain:

- **DSR:** 83.33% (25 out of 30 scenarios),
- **N-TAPG:** -2.3984% ,
- **P-TAPG:** 1.46% ,
- **TAPG:** -1.7556% .

The negative TAPG value shows that the agreed prices are, on average, slightly below the target price, meaning the system leans toward conservative, buyer-friendly deals while staying within acceptable margins. Meanwhile, the sizes of both N-TAPG and P-TAPG indicate that final prices stay close to company expectations rather than shifting toward large discounts or markups. Qualitatively, the system is able to retain buyer preferences across multiple turns and reuse them when generating offers. For example, once a buyer expresses a preference for specific incoterms (such as FOB or CIF), a payment method (such as LC or advance T/T), or a particular labeling and packaging scheme, the Negotiation Agent consistently respects these constraints. This behavior indicates effective cooperation between the Manager Agent, Receiver modules, and the pricing model, rather than isolated one-shot recommendations.

4.2.3 Ablation Study

The ablation results in Table 2 show that each major component is essential for stable negotiation performance. Removing the Manager Agent sharply low-

ers DSR ($0.833 \rightarrow 0.400$) and worsens TAPG, highlighting the need for centralized orchestration and dialogue management. Disabling the Receiver Modules causes inconsistent or incomplete requirement parsing, increasing both N-TAPG and TAPG due to misaligned offers. The largest degradation occurs when the Cost Estimator Agent is removed: DSR drops to 0.300 and N-TAPG reaches -0.3270 , revealing systematic underpricing without cost-aware guidance. Overall, the study confirms that MAFIA-NeT’s effectiveness emerges from the coordinated interaction of all agent modules rather than any single dominant component.

4.2.4 Discussion

The user scenario experiment, grounded in Equations 8–11 and Tables 2 and 1, shows that MAFIA-NeT integrates predictive pricing, LGNS-guided search, and multi-agent coordination into a coherent negotiation system for E-Agri commerce. High DSR values and consistently small pricing gaps indicate that the framework closes deals at prices aligned with internal expectations while adapting to buyer preferences and constraints. The ablation study confirms that this performance does not stem from any single component but emerges from the interplay of agents: the Manager Agent preserves dialogue coherence, the Receiver Modules ensure reliable structured inputs, and the Cost Estimator Agent grounds negotiation decisions in the pricing model described in Section 3.2. Together, these results show that MAFIA-NeT is accurate in its predictions and operationally robust, positioning it as a strong candidate for deployment in enterprise-scale agricultural trading environments.

5 CONCLUSION

This paper presented MAFIA-NeT, a multi-agent framework for automated agricultural negotiation that integrates LLM-based reasoning, structured data processing, and coordinated decision-making. The core LGNS mechanism constrains negotiation to a compact, LLM-guided subspace, while a lightweight LR-BR ensemble provides stable and interpretable price estimates. Experiments on real transactional data and 30 B2B negotiation transcripts from Agri Sung show that MAFIA-NeT achieves an 83.33% Deal Success Rate with tightly bounded price deviations, demonstrating strong alignment with internal pricing goals. Ablation results confirm that these gains emerge from the interaction of all agent modules Manager, Receiver, and Cost Estimator, rather than any single component. Overall, the findings indicate that LGNS-guided multi-agent negotiation offers a practical and scalable direction for LLM-driven automation in E-Agri commerce. Future work will extend the framework to larger product portfolios and incorporate enhanced human-in-the-loop supervision.

ACKNOWLEDGEMENTS

The authors would like to acknowledge Ho Chi Minh City University of Technology (HCMUT), Vietnam National University Ho Chi Minh City (VNU-HCM), for institutional support of this study. We also thank Agri Sung Joint Stock Company for their collaboration and for providing real-world negotiation scenarios that supported the experimental evaluation.

REFERENCES

- Alserhan, H., Altarawneh, R., Alyami, N., Alsheyab, Y., Alrababah, R., and Alshamayleh, H. (2025). The challenges and opportunities of implementing predictive analytics in marketing strategies and e-commerce personalisation techniques. *Future Generation Computer Systems*, 150:746–757.
- Botti, V., Garrido, A., Giret, A., and Noriega, P. (2017). The role of MAS as a decision support tool in complex domains. *Information Sciences*, 375:233–248.
- Calosso, T., Cantamessa, M., and Villa, A. (2017). Negotiation Approaches in B2B E-Commerce Applied to Supply Chain Management. *Procedia CIRP*, 63:12–17.
- Chen, D., Zhang, Q., and Zhu, Y. (2024). Efficient Sequential Decision Making with Large Language Models. In *Proceedings of the 2024 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 9157–9170.
- Chen, J. V., Le, H. T., and Tran, S. T. T. (2021). Understanding automated conversational agent as a decision aid: Matching agent’s conversation with customer’s shopping task. *Journal of Strategic Information Systems*, 30(1):101–127.
- De Cicco, R., Elmashhara, M. G., Silva, S. C., and Hammerschmidt, M. (2025). The impact of providing non-human identity cues about sales agents on consumer responses: The role of social presence and speciesism activation. *Journal of Retailing and Consumer Services*, 78:103585.
- Fang, J., Gao, S., Ren, P., Chen, X., Verberne, S., and Ren, Z. (2024). A Multi-Agent Conversational Recommender System. arXiv preprint arXiv:2402.01135.
- Haleema, P. K. and Iyengar, N. C. S. N. (2022). JADE-TraderPal – JADE mobile agent embedded with flexible negotiation strategies for dynamic e-commerce environments. *Journal of King Saud University - Computer and Information Sciences*, 34(5):1905–1915.
- Ikumoro, A. (2019). Intention to use intelligent conversational agents in e-commerce among Malaysian SMEs: An integrated conceptual framework based on tri-theories including Unified Theory of Acceptance. *International Journal of Academic Research in Business and Social Sciences*, 9(11):99–117.
- Koponen, J. P. and Rytsy, S. (2020). Social presence and e-commerce B2B chat functions. *Journal of Retailing and Consumer Services*, 54:102016.
- OpenAI (2024). Large Language Models: Survey, Technical Frameworks, and Future Challenges. *Artificial Intelligence Review*, 57:260.
- Reddy, K. and Kumar, V. (2022). Machine learning approaches for agricultural commodity price prediction: A survey. *Journal of King Saud University - Computer and Information Sciences*, 34(9):6932–6944.
- Sharma, P., Kumar, R., and Jha, S. (2023). Deep learning-based price prediction and demand forecasting in agri-commerce. *Computers and Industrial Engineering*, 179:109076.
- Singh, R., Kumar, R., and Verma, A. (2022). Multi-agent systems for smart agriculture: A review. *Artificial Intelligence in Agriculture*, 6:87–104.
- Wang, J. et al. (2024). A Survey on Large Language Models: Applications, Challenges, and Future Directions. *IEEE Transactions on Neural Networks and Learning Systems*. Early Access.
- Wei, J., Wang, X., Schuurmans, D., Bosma, M., Ichter, B., Xia, F., Chi, E., Le, Q., and Zhou, D. (2022). Chain-of-Thought Prompting Elicits Reasoning in Large Language Models.
- Xi, Y., Song, Z., Wu, Y., Zhang, K., Pan, L., Chen, Q., and Zhao, W. X. (2023). Large Language Models as General-Purpose Agents: A Survey.
- Zeng, J., Liu, H., Dai, Z., Tang, X., Luo, C., Varshney, S., Li, Z., and He, Q. (2025). Cite before you speak: Enhancing context-response grounding in e-commerce conversational LLM-agents. arXiv preprint arXiv:2503.04830.