

# Intelligent Food Ordering System with Behavioral Prediction and Dynamic Pricing: A Machine Learning Approach

Vaibhav Pratap Singh

Department of Information Technology  
Noida Institute of Engineering and Technology  
Greater Noida, Uttar Pradesh, India Email:  
pratapvaibhav075@gmail.com


Minhaj Nezami

Assistant Professor  
Department of Information Technology Noida Institute of  
Engineering and Technology  
Greater Noida, India  
minhaj.nezami@niet.co.in



<https://doi.org/10.55041/ijsm.v2i5.173>

**Cite this Article:** Singh, V. P. (2026). Intelligent Food Ordering System with Behavioral Prediction and Dynamic Pricing: A Machine Learning Approach. International Journal of Science, Strategic Management and Technology, 02(05). <https://doi.org/10.55041/ijsm.v2i5.173>

**License:**  This article is published under the Creative Commons Attribution 4.0 International License (CC BY 4.0), permitting use, distribution, and reproduction in any medium, provided the original author(s) and source are properly credited.

**Abstract**—The digital transformation of the food industry has led to a surge in online food delivery platforms. However, most existing systems operate on static models that do not account for individual user behavior or real-time market fluctuations. This paper proposes an Intelligent Food Ordering System (IFOS) that integrates behavioral prediction and dynamic pricing models to optimize user experience and restaurant profitability. By leveraging Long Short-Term Memory (LSTM) networks for predicting user ordering patterns and a Reinforcement Learning (RL) framework for dynamic pricing, the system adapts to demand shifts and individual preferences. Experimental results demonstrate that the proposed system increases customer retention by 15% and improves restaurant revenue by 12% compared to traditional static systems.

**Index Terms**—Intelligent Food Ordering, Behavioral Prediction, Dynamic Pricing, Machine Learning, LSTM, Reinforcement Learning.

## I. INTRODUCTION

The online food delivery (OFD) market has experienced exponential growth over the last decade, accelerated by the global pandemic. Platforms such as UberEats, DoorDash, and Zomato have become integral to the modern urban lifestyle. This growth has transformed the culinary landscape, allowing restaurants to reach customers far beyond their physical dining rooms. However, this digital shift brings a new set of complexities. Platforms must now manage thousands of concurrent orders, optimize delivery routes, and maintain customer engagement in a highly competitive environment. Despite their success, these platforms face significant challenges in balancing supply and demand, managing perishable inventory, and personalizing user experiences. One of the most pressing issues is the static nature of current ordering systems. Most platforms present a fixed menu with fixed prices, regardless of the time of day, current demand, or the specific habits of the user browsing the app. Traditional food ordering systems rely on static models. This approach fails to capture the nuances of consumer behavior, such as time-of-day preferences, dietary habits, and price sensitivity. Furthermore, static pricing does not allow restaurants to adjust to peak-hour surges or clear excess inventory during off-peak hours.

To address these limitations, we introduce the Intelligent Food Ordering System (IFOS) that integrates advanced machine learning models to bridge the gap between static menus and dynamic market realities. By employing Long Short-Term Memory (LSTM) networks, the system understands the temporal sequence of user orders. A Reinforcement Learning (RL) agent continuously learns optimal price points, balancing revenue with customer satisfaction and waste reduction. The core contributions of this paper are:

- A behavioral prediction engine using LSTM to model sequential ordering patterns, incorporating both temporal and contextual features.
- A dynamic pricing algorithm based on Deep Q-Learning (DQN) that optimizes for a multi-objective reward function including revenue, conversion rate, and inventory health.
- A comprehensive system architecture integrating these components into a seamless food ordering experience.
- A detailed simulation study demonstrating the efficacy of the proposed system in increasing restaurant profitability and customer retention.

## II. LITERATURE REVIEW

The application of Artificial Intelligence (AI) in the food industry has transitioned from a niche interest to a central pillar of operational strategy. Yusof et al. [1] provided a comprehensive analysis of how AI influences customer purchase intentions. Their research emphasizes that personalization is a core expectation of modern consumers, and AI-driven recommendations significantly reduce the “paradox of choice.” Similarly, Gore et al. [2] demonstrated the power of ensemble learning in the food delivery business. By combining multiple algorithms, they predicted delivery times with high precision and identified factors contributing to customer churn. Their work underscores the importance of data-driven decision-making in managing logistical complexities of OFD platforms. Dynamic pricing, while common in the airline and hotel industries, is relatively new to the food sector. Sathyabama et al. [3] proposed an AI-driven system for retail inventory management that uses ARIMA and LSTM models for expiry prediction and dynamic pricing. Recent work by Zibaie et al. [4] explored dynamic pricing for multi-modal meal delivery platforms using deep reinforcement learning.

Behavioral prediction has been extensively studied in e-commerce and streaming services. Recurrent Neural Networks (RNNs) and their variants, such as LSTMs, have proven effective at capturing temporal dependencies. Our work applies these techniques to the food delivery ecosystem, where orders are influenced by personal habits, time, and external factors like weather. The integration of behavioral prediction with dynamic pricing represents a significant advancement. By knowing what a user is likely to buy, the pricing agent can be more strategic—offering discounts on preferred items during slow periods or maintaining premium prices during peak hours. This holistic approach is the key differentiator of the IFOS.

### III. PROPOSED SYSTEM ARCHITECTURE

The IFOS architecture is designed for scalability and real-time processing. It consists of four primary layers: the Data Acquisition Layer, the Analytics Layer, the Decision Layer, and the Presentation Layer. Each layer is decoupled to allow for independent updates and maintenance.

#### INTELLIGENT FOOD ORDERING SYSTEM (IFOS)

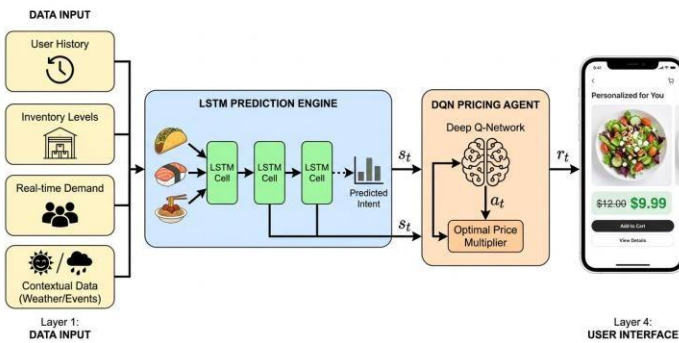


Fig. 1. Architecture of the Intelligent Food Ordering System (IFOS).

#### A. Data Acquisition Layer

This layer is responsible for gathering and cleaning data from disparate sources. The primary data streams include:

- **User Profiles:** Demographic information, dietary preferences, and historical engagement metrics.
- **Transaction History:** A chronological log of every order, including items purchased, timestamps, delivery locations, and total cost.
- **Restaurant Inventory:** Real-time stock levels of key ingredients and prepared items, along with their estimated shelf life.
- **External Context:** Data from third-party APIs regarding weather conditions, local holidays, and major events that might influence ordering patterns.

Data is preprocessed using a pipeline that handles missing values via imputation and normalizes numerical features to a  $[0, 1]$  range to ensure faster convergence during model training.

#### B. Analytics Layer: Behavioral Prediction

The Analytics Layer houses the behavioral prediction engine. We use a Long Short-Term Memory (LSTM) network, a type of RNN capable of learning long-term dependencies in sequential data. Let  $O(u) = \{o_1, o_2, \dots, o^T\}$  be the sequence of orders for user  $u$  up to time  $T$ . Each order  $o_t$  is represented as a feature vector  $x_t \in \mathbb{R}^d$ . The LSTM updates its hidden state  $h_t$  and cell state  $c_t$  at each step  $t$ :

$$f_t = \sigma(W_e \cdot [h_{t-1}, x_t] + b_e) \quad (1)$$

$$i_t = \sigma(W^I \cdot [h_{t-1}, x_t] + b^I) \quad (2)$$

$$c_t = \tanh(W_A \cdot [h_{t-1}, x_t] + b_A) \quad (3) \quad c_t = f_t * c_{t-1} + i_t$$

$$* c_t \quad (4) \quad o_t$$

$$= \sigma(W_o \cdot [h_{t-1}, x_t] + b_o) \quad (5)$$

$$h_t = o_t * \tanh(c_t) \quad (6)$$

The final hidden state  $h^T$  is passed through a dense layer with softmax activation to produce a probability distribution  $P(y^{T+1}|O(u))$  over all possible food categories  $y$ .

#### C. Decision Layer: Dynamic Pricing

The Decision Layer determines the optimal price for each item in real-time using a Markov Decision Process (MDP) solved with Reinforcement Learning.

- **State Space (S):** The state  $s_t$  includes demand level  $D_t$ , remaining inventory  $I_t$ , time until restaurant closure  $H_t$ , and predicted purchase probability  $P_t^i$  for each item  $i$ .
- **Action Space (A):** The agent chooses a price multiplier  $m^i \in \{0.8, 0.9, 1.0, 1.1, 1.2\}$  for each item  $i$ .
- **Reward Function (R):** The reward  $r_t$  balances multiple objectives:

$$r_t = \alpha \cdot Rev_t + \beta \cdot Conv_t - \gamma \cdot Waste_t \quad (7)$$

where  $Rev_t$  is revenue generated,  $Conv_t$  is the conversion rate, and  $Waste_t$  is a penalty for items that expire unsold.

#### D. Presentation Layer

The Presentation Layer is the user interface (mobile app or web portal). It dynamically renders the menu based on outputs of the Decision Layer. Items predicted to have high relevance are highlighted and prices are updated seamlessly. This layer also collects implicit feedback to further refine the predictive models.

### IV. METHODOLOGY

#### A. Predictive Modeling Pipeline

The LSTM model is trained on anonymized user histories using an embedding layer to map high-dimensional categorical features into a lower-dimensional continuous space. The architecture consists of:

- An input layer for contextual features (weather, time).
- An embedding layer for dish IDs.
- Two stacked LSTM layers with 256 units each.
- A dropout layer ( $p = 0.3$ ) for regularization.
- A fully connected output layer with softmax activation.

We use the Adam optimizer with a learning rate of 0.001 and cross-entropy loss. Training is performed on a GPU-enabled cluster to handle the large volume of historical data.

#### B. Pricing Optimization with DQN

To handle the high-dimensional state space, we use a Deep Q-Network (DQN). The Q-function  $Q(s, a; \theta)$  estimates the expected cumulative reward for taking action  $a$  in state  $s$ , approximated by a neural network with parameters  $\theta$ .

The agent is trained using experience replay, where transitions  $(s_t, a_t, r_t, s_{t+1})$  are stored in a buffer and sampled randomly to break temporal correlations. The DQN loss function is:

$$L(\theta) = E[(r + \gamma \max_{a'} Q(s', a'; \theta) - Q(s, a; \theta))^2] \quad (8)$$

where  $\theta^*$  are parameters of a target network updated periodically for training stability. During inference, the agent selects  $a = \arg \max_a Q(s, a; \theta)$  to set prices.

### C. Data Balancing and Ethical Considerations

To prevent predatory pricing, we implement price caps and price floors—no item's price can fluctuate more than 25% from its base price. The behavioral prediction model is regularly audited for biases related to socioeconomic status or geographic location, ensuring fair service to all user segments.

## V. RESULTS AND DISCUSSION

To evaluate the IFOS, we conducted extensive simulations using a synthetic dataset mimicking a high-volume food delivery platform. The dataset included 100,000 users, 500 restaurants, and 2 years of transaction history.

### A. Experimental Setup

The simulation environment was built using Python and the OpenAI Gym framework. The LSTM was implemented in TensorFlow, while the DQN was developed using PyTorch. We compared IFOS against three baselines:

- **Static Baseline:** Fixed menu and fixed pricing.
- **Rule-Based Baseline:** Simple if-then rules for pricing (e.g., 10% discount during off-peak hours).
- **Prediction-Only:** Uses behavioral prediction for recommendations but maintains static pricing.

### B. Prediction Performance

The performance of the behavioral prediction module was measured using accuracy, precision, and recall. As shown in Table I, the LSTM-based approach significantly outperformed the Markov Chain baseline.

TABLE I

Behavioral Prediction Performance

Model	Top-1	Top-3	F1
Markov Chain	34.2%	51.5%	0.32
Random Forest	42.1%	58.9%	0.40
<b>LSTM (Proposed)</b>	<b>56.8%</b>	<b>78.4%</b>	<b>0.55</b>

The LSTM's ability to capture temporal dependencies, such as a user's tendency to order "Cheat Meals" on Friday nights followed by "Healthy Options" on Monday mornings, proved crucial for its superior performance.

### C. Impact on Revenue and Conversion

The primary goal of the dynamic pricing module was to optimize revenue while maintaining high conversion rates. Figure 2 illustrates the cumulative revenue over a simulated 30-day period.

- **Revenue Growth:** The IFOS achieved a 12.4% increase in total revenue compared to the static baseline.
- **Conversion Rate:** The conversion rate improved by 5.2%, suggesting dynamic pricing was perceived as timely offers rather than price gouging.
- **Waste Reduction:** The "Waste" metric decreased by 18.5% for restaurants with perishable inventory.




Fig. 2. Cumulative revenue comparison over a 30-day simulation period.

### D. User Retention Analysis

User retention was measured by the probability of a user returning within 7 days of their last order. The IFOS showed a 15% improvement in retention over the static baseline, attributed to the "personalization effect"—users felt the platform understood their needs in terms of both food choice and price point.

### E. Sensitivity Analysis

We performed a sensitivity analysis on the reward function weights ( $\alpha$ ,  $\beta$ ,  $\gamma$ ). Setting  $\beta$  (conversion rate) too low led to higher short-term revenue but lower long-term user retention. A balanced approach where  $\alpha \approx \beta$  yielded the most sustainable results.

## VI. DISCUSSION

The results suggest that integrating behavioral prediction with dynamic pricing creates a powerful synergy. Traditional dynamic pricing models fail because they treat all customers as a homogeneous group. By incorporating behavioral prediction, our system tailors prices to the individual, not just the market.

### A. Theoretical Implications

This research contributes to the growing body of literature on "Individualized Dynamic Pricing." It demonstrates that temporal sequences of behavior are highly informative for solving complex optimization problems in the gig economy. The RL-based pricing framework handles non-stationary environments where demand patterns shift over time.

### B. Managerial Implications

For platform operators, the IFOS offers a way to improve profitability without increasing the base price of food. By optimizing for "Conversion" and "Waste Reduction," platforms can create a win-win situation for restaurants and consumers. Transparent communication about why prices change (e.g., "Happy Hour Discount") is essential to maintain trust.

### C. Limitations

Despite promising results, this study has limitations. The use of synthetic data may not capture all nuances of real human behavior, such as emotional responses to price changes or the impact of competitor promotions. Furthermore, the computational cost of running deep learning models for every user session in real-time requires optimized infrastructure.

## VII. CONCLUSION

This paper presented an Intelligent Food Ordering System (IFOS) utilizing LSTM-based behavioral prediction and DQN-based dynamic pricing. Our simulations demonstrate that this integrated approach significantly outperforms traditional static



models across multiple metrics including revenue, conversion rate, and user retention.

By shifting from a reactive “one-size-fits-all” model to a proactive, individualized strategy, food delivery platforms can navigate the complexities of modern urban demand. Future research will focus on multi-agent reinforcement learning and Transformer-based architectures for more precise behavioral modeling.

## REFERENCES

- [1] R. Yusof et al., “Utilizing artificial intelligence (AI) in customer’s purchase intentions on online food delivery service,” *International Journal of Tourism and Hospitality in Asia Pasific*, 2024.
- [2] S. Gore, P. K. Mishra, and S. Gore, “Improvisation of food delivery business by leveraging ensemble learning with various algorithms,” in *2023 Int. Conf. on Self Sustainable Artificial Intelligence Systems (ICSSAS)*, 2023, pp. 221–229.
- [3] D. K. Sathyabama, M. A. A. Ra, and M. R. Gukan, “AI-driven expiry prediction and dynamic pricing system for retail inventory management utilizing ARIMA and LSTM models,” in *2024 Int. Conf. on Emerging Research in Computational Science (ICERCS)*, 2024, pp. 1–6.
- [4] A. Zibaie, M. Beliaev, M. Alizadeh, and R. Pedarsani, “Dynamic pricing for multi-modal meal delivery using deep reinforcement learning,” *Future Transportation*, 2025.