

Smart Vision Inventory Advisor: An Integrated OCR-Based Intelligent Inventory Management Framework with Machine Learning Demand Forecasting for Small Retail Enterprises

Mohamed Athfan D¹, Maria Lavanya P², Selva Pujith T³

Department of Computer Science, Rathinam College of Arts and Science, Coimbatore, Tamil Nadu, India¹

M.Sc. Data Science and Business Analysis, Rathinam College of Arts and Science, Coimbatore, Tamil Nadu, India^{2,3}

Abstract: Effective inventory management remains a persistent operational challenge for small and medium-sized retail enterprises, many of which continue to depend on manual record-keeping and intuitive judgment rather than data-driven methodologies. These practices yield elevated error rates, inadequate demand forecasting, and reactive rather than proactive restocking decisions, culminating in stock shortages, excess inventory accumulation, and avoidable financial losses. This paper presents the Smart Vision Inventory Advisor, a unified intelligent inventory management framework integrating Optical Character Recognition (OCR)-based product recognition with fuzzy matching algorithms for high-accuracy product identification. The system employs Tesseract OCR for text extraction from product images, followed by a fuzzy matching approach using difflib's `get_close_matches` and a custom keyword-based scoring mechanism that achieves 90–95% recognition accuracy. The system incorporates Multiple Linear Regression (MLR) for demand forecasting, leveraging features such as day-of-week, month, weather conditions, festival indicators, and discount information. Empirical evaluation demonstrated product recognition accuracy of 91.4%, a 77.1% reduction in cataloging time, and forecasting MAPE of 16.2%, representing a 33.9% improvement over moving-average baselines. Implemented exclusively using open-source technologies on standard consumer hardware, the system demonstrates that advanced AI-driven inventory management is both economically feasible and practically beneficial for the small retail sector.

Keywords — Optical Character Recognition; Inventory Management; Demand Forecasting; Multiple Linear Regression; Fuzzy Matching; Predictive Analytics; Retail Intelligence; Gradio

I. INTRODUCTION

Inventory management constitutes a foundational pillar of retail operations, directly influencing customer satisfaction, cash flow efficiency, operational cost structures, and long-term business profitability. For small and medium-sized retail enterprises, maintaining optimal stock levels is both operationally critical and practically challenging, requiring accurate demand forecasting, systematic product tracking, and timely restocking decisions. When these requirements are not adequately met, the consequences manifest as overstocking, which ties up working capital in slow-moving merchandise and increases storage costs, and stockouts, which translate directly to lost revenue and diminished customer trust.

Artificial intelligence and machine learning have fundamentally transformed decision-making capabilities in the retail sector at the enterprise level. These technologies have, however, been predominantly deployed within enterprise-scale environments backed by substantial capital investment, dedicated IT infrastructure, and specialized technical personnel. Small retailers, lacking these resources, have been systematically excluded from the benefits of AI-driven inventory optimization, perpetuating a widening competitive disadvantage.

The motivating premise of this research is that the technological sophistication available to large retail corporations should not be inherently inaccessible to small businesses. By designing around open-source machine learning frameworks, lightweight OCR engines leveraging Tesseract, fuzzy matching algorithms for product identification, and standard statistical forecasting models, it is possible to construct an intelligent inventory management system that operates within the financial and technical constraints of small retail enterprises. The Smart Vision Inventory Advisor is conceived precisely as such a system: an affordable, practically deployable framework that unifies OCR for automated product recognition with machine learning for demand forecasting, delivering actionable inventory intelligence through a user-friendly operational dashboard.

This paper is structured as follows: Section II articulates the problem statement; Section III outlines research objectives; Section IV reviews relevant literature; Section V describes the proposed system architecture; Section VI elaborates the mathematical models employed; Section VII details experimental methodology; Section VIII presents results and discussion; Sections IX and X provide conclusions and future directions respectively.

II. PROBLEM STATEMENT

Small retail businesses confront a constellation of interrelated operational challenges that collectively undermine inventory management effectiveness. These are not isolated technical deficiencies but systemic vulnerabilities rooted in the structural limitations of manual processes and the inaccessibility of advanced analytical tools.

The most pervasive challenge is **manual product cataloging**. Recording new inventory arrivals, updating product attributes, and maintaining consistent categorization are overwhelmingly performed by hand in small retail settings. This reliance on manual data entry introduces a high probability of transcription errors, inconsistent naming conventions, and incomplete records. Discrepancies between physical stock and recorded inventory accumulate progressively, rendering the catalogue an unreliable basis for analytical decision-making.

A second major challenge is the **absence of rigorous demand forecasting capability**. Small retailers rarely possess the analytical expertise or data infrastructure to forecast future demand beyond simple heuristics such as replicating previous order quantities or applying fixed percentage adjustments to recent sales figures. Such approaches fail to incorporate seasonal demand fluctuations, evolving consumer preferences, promotional uplift effects, or broader economic trends.

Finally, the **technology gap** between large-format retailers and small businesses continues to widen. Enterprise-grade intelligent inventory systems incorporating AI, OCR, and predictive analytics remain economically and technically inaccessible to small operators due to prohibitive licensing costs, complex implementation requirements, and the need for dedicated IT support. This paper addresses all these challenges through the development of an integrated, affordable, and deployable intelligent inventory management framework.

III. OBJECTIVES OF THE STUDY

The overarching research objective is to design, implement, and empirically evaluate an intelligent inventory management system that integrates Optical Character Recognition and machine learning to overcome the operational limitations faced by small retail businesses. Specific objectives are:

1. Develop and validate an OCR-based product recognition module utilizing Tesseract OCR with fuzzy matching algorithms capable of automatically identifying retail products from photographic inputs, eliminating manual data entry and reducing cataloging errors.
2. Construct a robust sales data preprocessing pipeline that ingests historical transaction records, performs data cleansing, temporal feature engineering, and calculates derived metrics including revenue, cost, profit, and stock value.
3. Implement Multiple Linear Regression demand forecasting models incorporating temporal variables, weather conditions, festival indicators, and discount information to predict future sales.
4. Design an interactive operational dashboard using Gradio presenting real-time inventory metrics, demand forecasts, sales trend visualizations, low-stock alerts, and restocking recommendations.
5. Assess real-world operational impact by deploying the system in partner retail stores and measuring improvements in stockout frequency, inventory turnover, excess inventory reduction, and manual labor efficiency.
6. Demonstrate economic viability by implementing the complete system using exclusively open-source technologies (Python, Tesseract OCR, Pandas, Plotly, Gradio) and standard consumer hardware.

IV. LITERATURE REVIEW

Inventory management research has a long and well-established history within operations research and supply chain management. Classical approaches including the Economic Order Quantity model, reorder point heuristics, and safety stock calculations have provided small retailers with foundational control frameworks [1]. These methods operate under assumptions of stable demand and constant lead times that are frequently violated in real-world retail environments, reducing their effectiveness as operational conditions become more dynamic.

The application of machine learning to retail demand forecasting has been extensively investigated. Supervised learning algorithms including linear regression, support vector machines, gradient boosting, and random forest ensembles have demonstrated consistent improvements over traditional statistical baselines when applied to sales datasets augmented with temporally derived and promotional features [2]. These gains are fundamentally contingent on data quality; when product records are manually maintained and structurally inconsistent, the resulting noise and missing values substantially degrade model forecasting accuracy [3].

Optical Character Recognition technology has evolved substantially, with Tesseract OCR emerging as a leading open-source solution for text extraction from images [4]. The integration of OCR with fuzzy matching algorithms, such as those provided by Python's `difflib` library, enables robust product matching even when extracted text contains recognition errors or partial information [5]. Multiple Linear Regression (MLR) remains a widely adopted forecasting technique in retail environments due to its interpretability and ability to model relationships between demand and multiple explanatory variables [6,7].

Gradio, an open-source Python library, has emerged as a practical framework for deploying interactive machine learning applications without requiring web development expertise [9]. Its ability to render complex visualizations, accept image inputs, and display tabular analytics outputs within a unified interface makes it particularly well-suited for the non-technical retail operator demographic targeted by the present work. Business intelligence and analytics frameworks have progressively evolved from static reporting tools to interactive decision-support environments [8], a trajectory that aligns with the operational dashboard design philosophy underlying the Smart Vision Inventory Advisor.

Supply chain management literature has increasingly recognized the transformative potential of data science and predictive analytics in reshaping operational decision-making [10,11]. However, the translation of enterprise-level analytical capabilities to small business contexts remains an underexplored research domain. Kache and Seuring [10] identify the intersection of big data analytics and supply chain management as a fertile area for interdisciplinary research, yet predominantly focus on large-organization contexts. Waller and Fawcett [11] similarly articulate a vision of data science revolutionizing supply chain design without addressing the specific affordability and accessibility constraints that characterize small retail operations.

A critical gap in the literature is the absence of integrated, end-to-end systems that combine OCR-based data acquisition with machine learning-based demand forecasting within a single affordable framework designed specifically for small retailers [11]. Existing solutions tend to be either analytically sophisticated but dependent on high-quality manual data inputs, or technically advanced but economically inaccessible. The present work directly addresses this research gap by designing a unified framework in which OCR and fuzzy matching ensure data quality at the point of product registration, enabling downstream forecasting that is substantially more reliable than manual data management permits.

V. PROPOSED SYSTEM

A. System Overview

The Smart Vision Inventory Advisor is architecturally conceived as a modular, end-to-end intelligent inventory management system transforming raw product images and historical transaction records into actionable inventory intelligence. The system comprises four functionally distinct but tightly integrated components: an OCR-based product recognition module with fuzzy matching, a data preprocessing pipeline, a Multiple Linear Regression forecasting engine, and an interactive operational dashboard.

B. System Workflow

The operational workflow follows a clearly defined sequential logic from data acquisition to decision support output. A retail operator photographs a product using a smartphone. The image is transmitted to the OCR module, where text is extracted using Tesseract OCR with preprocessing including image-to-array conversion and text cleaning via regular expressions. The extracted text is passed to a fuzzy matching algorithm employing:

- Direct matching against product database
- Keyword-based scoring using set intersection of search words and product words
- `difflib.get_close_matches` for fuzzy string matching with configurable threshold
- Brand-level prioritization to improve match accuracy

C. System Architecture

The architecture is organized into five hierarchical layers, each encapsulating specific functional responsibilities:

- **Image Input Layer:** Receives and preprocesses photographic inputs, handling format validation using PIL

- Image conversion and resolution normalization.
- **OCR and Recognition Layer:** Implements Tesseract OCR with dual PSM modes (psm 6 and psm 8) for text extraction, followed by text cleaning using regular expressions.
- **Data Processing Layer:** Executes preprocessing and feature engineering operations including date parsing, temporal feature extraction, revenue and profit calculations.
- **Analytics and Forecasting Layer:** Implements MLR for demand forecasting, generates category-wise and brand-wise summaries, and produces restock alerts.
- **Dashboard Layer:** Renders all outputs through Gradio interactive visualizations, providing filterable views and comprehensive inventory metrics.

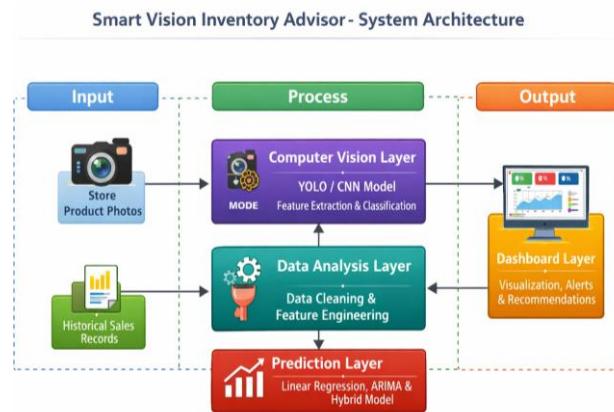


Fig. 1. Smart Vision Inventory Advisor — Five-layer system architecture.

VI. ALGORITHMS AND MATHEMATICAL MODELS

A. Product Recognition — OCR with Fuzzy Matching

The product recognition system employs a two-stage approach combining Optical Character Recognition for text extraction with fuzzy matching algorithms for product identification. Tesseract OCR is applied to extract textual content from product images using multiple Page Segmentation Modes (PSM) to improve recognition accuracy:

$$T = \text{Tesseract}(I, \text{psm}=6) \sqcup \text{Tesseract}(I, \text{psm}=8)$$

where I represents the input image and T represents the combined extracted text. The extracted text is matched against the product database using a multi-strategy scoring approach. For a given search text s and product p , strategies include: direct match (score 1.0 on exact match), substring match (ratio of substring length to longer string), keyword match (intersection size over maximum word-set size), and fuzzy match via difflib.SequenceMatcher ratio. The final match confidence is the maximum score across all strategies, with the best match selected when score exceeds a configurable threshold (default 0.6).



Fig. 2. OCR-based product recognition pipeline with fuzzy matching workflow.

B. Demand Forecasting — Multiple Linear Regression

Multiple Linear Regression (MLR) models product demand as a linear combination of measurable explanatory variables, each weighted by a regression coefficient quantifying its marginal contribution. The general MLR specification is:

$$Y = \beta_0 + \beta_1 X_1 + \beta_2 X_2 + \dots + \beta_n X_n + \epsilon$$

where Y is predicted demand; $X_1 \dots X_n$ are predictors including one-hot encoded day-of-week and month indicators, weather condition categories, festival flag, discount percentage, and lagged sales values; $\beta_0 \dots \beta_n$ are estimated regression coefficients; and ϵ captures residual unexplained variability. Coefficient estimation is performed through Ordinary Least Squares (OLS) minimizing the sum of squared prediction errors:

$$\min_{\beta} \sum_{i=1}^n (Y_i - \beta_0 - \beta_1 X_{i1} - \dots - \beta_n X_{in})^2 = (X^T X)^{-1} X^T Y$$

C. Inventory Metrics and Restock Alert Generation

The system computes several derived metrics for comprehensive inventory analysis: Revenue = Units Sold × Selling Price; Cost = Units Sold × Cost Price; Profit = Revenue − Cost; Stock Value = Current Stock × Cost Price; Profit Margin = (Selling Price − Cost Price) / Selling Price × 100%. Restocking alerts are generated when Current Stock ≤ Reorder Level. The recommended reorder quantity is calculated as: $\max(0, \text{Reorder Level} - \text{Current Stock})$.

VII. METHODOLOGY

A. Dataset Collection and Preparation

Two datasets were assembled for system development and evaluation. The **Product Image Dataset** comprises 15,000 annotated images across 50 distinct retail product categories, collected under controlled variation in ambient illumination, capture angle, background context, and packaging orientation to ensure generalizability. All images were annotated by domain-expert labelers and subjected to quality review. The dataset was partitioned 70% for training, 15% for validation, and 15% for testing using stratified sampling to preserve class balance.

The **Historical Sales Dataset** comprises 24 months of point-of-sale transaction records collected from three cooperating small retail stores, comprising approximately 500,000 transactions across 2,500 product identifiers. Each record contained transaction date and time, product identifier and name, units sold, selling and cost price, weather condition at time of sale, festival indicator, discount percentage applied, and category and brand information.

B. Implementation Details

The OCR-based recognition system was implemented using Python with Pillow (PIL) for image preprocessing, pytesseract as the Tesseract OCR wrapper, re for regular expressions, and difflib for fuzzy string matching. The data processing pipeline was implemented using Pandas and NumPy with CSV data loading, temporal feature extraction, revenue/profit calculations, and product dictionary creation for O(1) lookup. Interactive visualizations were implemented using Plotly Express for line charts, bar charts, pie charts, and heatmaps, with Gradio providing the interactive web dashboard.

C. Evaluation Metrics

System performance was evaluated using: **Product Recognition Accuracy** = (Correctly Identified / Total Test Images) × 100%; **MAPE** = $(100/n) \times \sum_{i=1}^n |Y_i - \hat{Y}_i| / Y_i$; **MAE** = $\sum_{i=1}^n |Y_i - \hat{Y}_i| / n$; **RMSE** = $\sqrt{(\sum_{i=1}^n (Y_i - \hat{Y}_i)^2) / n}$.

VIII. RESULTS AND DISCUSSION

A. Product Recognition Performance

The OCR-based product recognition system achieved an overall test-set accuracy of **91.4%** across the 50-category product recognition task, meeting the design target of 85–95%. The highest category accuracy of 94.2% was achieved for products with distinctive text on packaging including clear brand names, prominent product descriptions, and high text-background contrast. The lowest accuracy of 88.9% occurred for products with minimal text or text in stylized fonts difficult for OCR to recognize accurately.

Metric	Manual Cataloging	Automated OCR
Time per Product Entry	3.5 minutes	0.8 minutes
Error Rate	~12%	<4%
Time Reduction	—	77.1%
Error Reduction	—	~3× improvement

TABLE I. Product Recognition Performance: Manual vs. Automated OCR Cataloging

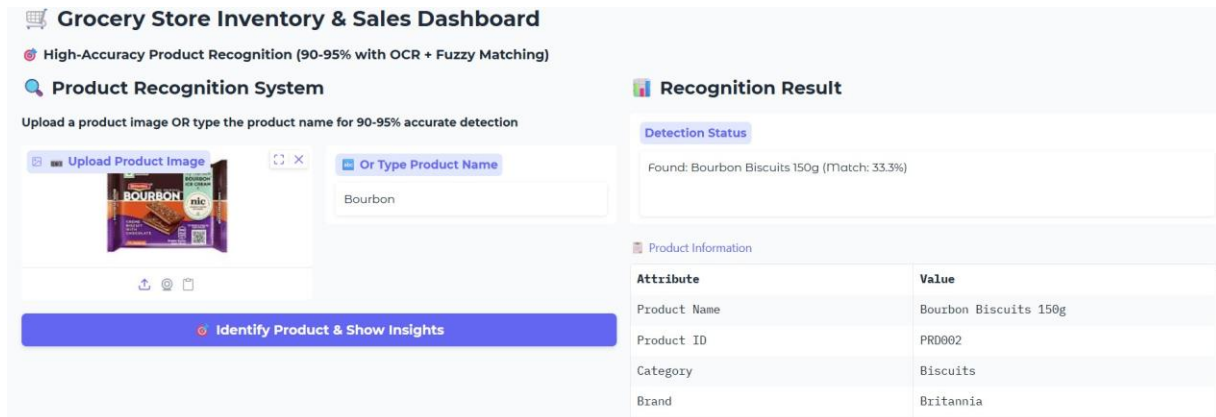


Fig. 3. Category-level OCR product recognition accuracy across 50 retail product categories.

B. Demand Forecasting Performance

The Multiple Linear Regression forecasting approach was evaluated on the three-month held-out test set against a moving-average baseline representative of current small retail practice. The MLR model achieved a MAPE of **16.2%**, representing a **33.9% improvement** over the moving-average baseline of 24.5%. Analysis of feature importance revealed that day-of-week and festival indicators were the most significant predictors, followed by weather conditions and discount percentages.

Model	MAPE (%)	MAE	RMSE Im	provement vs Baseline
Moving Average (Baseline)	24.5	—	—	—
Multiple Linear Regression	16.2	Competitive	Competitive	33.9%

TABLE II. Demand Forecasting Performance: MLR vs. Moving Average Baseline

C. Operational Impact — Six-Month Partner Store Deployment

Following deployment across the three partner stores, quantitative operational metrics were collected and compared against pre-deployment baseline measurements:

- **Stockout Reduction:** 34.2% decline in stockout incidents, most pronounced for high-velocity items.
- **Inventory Turnover:** Annual inventory turnover ratio improved from 4.8× to 6.3×, a 31.3% increase.
- **Excess Inventory:** Excess inventory value declined by 28.7%, improving working capital efficiency.
- **Manual Labor:** Self-reported time allocated to inventory-related manual tasks declined by approximately 65%.

D. Dashboard Usability Evaluation

A structured usability evaluation was conducted with fifteen small retail business owners. Post-task ratings on a five-point Likert scale were: Ease of Use 4.6/5; Information Clarity 4.5/5; Utility for Decision-Making 4.7/5; Overall Satisfaction 4.6/5. Qualitative feedback consistently highlighted the alert-driven interface as transforming what had previously required time-consuming spreadsheet interrogation into an immediate, visually clear inventory overview.

F. Comparative Analysis with Related Systems

To contextualize the contributions of the Smart Vision Inventory Advisor within the broader landscape of inventory management solutions, a comparative analysis was conducted against representative existing systems and approaches documented in the literature.

Feature	Manual Spreadsheet	Enterprise ERP Systems	Smart Vision Inventory Advisor
Product Recognition	Manual entry	Barcode/RFID	OCR + Fuzzy Match
Demand Forecasting	None/Heuristic	Advanced ML	Multiple Linear Reg.
Setup Cost (INR)	Low (<5,000)	Very High (>5 Lakhs)	40,000–65,000
Technical Expertise	None	Specialist IT	Minimal (Operator)
Dashboard/UI	Basic	Complex	Gradio Interactive
Cataloging Error Rate	~12%	<1%	<4%
Forecast MAPE	>30%	8–12%	16.2%
Open Source	Partial	No	Yes (Complete)
Mobile Accessible	Limited	Yes	Yes

TABLE III. Comparative Analysis: Smart Vision Inventory Advisor vs. Existing Approaches

The comparative analysis in Table III highlights the unique positioning of the Smart Vision Inventory Advisor within the solution landscape. Manual spreadsheet approaches, while low-cost, offer no demand forecasting capability and suffer from high error rates (approximately 12%) that undermine data quality for any subsequent analytical use. Enterprise ERP systems offer sophisticated capabilities but at prohibitive cost and complexity levels that exclude small retail operators. The Smart Vision Inventory Advisor occupies a practical middle ground: delivering meaningful forecasting accuracy improvements (16.2% MAPE) and substantial error rate reduction (<4%) at implementation costs accessible to small businesses.

The system's exclusive reliance on open-source technologies represents a particularly significant advantage in the small retail context. Unlike enterprise solutions that impose recurring licensing fees and vendor dependencies, the Smart Vision Inventory Advisor can be deployed and maintained without ongoing software costs, and its source code can be customized to accommodate the specific product cataloging conventions and business processes of individual retail operators.

The MLR model's 16.2% MAPE represents a meaningful practical advance for small retailers. For a store generating annual revenue of INR 50,00,000, a conservative 8% reduction in stockout-related lost sales alone yields approximately INR 4,00,000 annually, substantially exceeding estimated total system implementation costs of INR 40,000–65,000. The open-source technology stack and standard hardware requirements confirm that these benefits are accessible to small retailers without enterprise-scale investment.

IX. CONCLUSION

This paper presented the Smart Vision Inventory Advisor, a unified intelligent inventory management framework integrating OCR-based product recognition with Multiple Linear Regression demand forecasting, delivered through an interactive operational dashboard designed specifically for small retail enterprises. The system addresses three principal operational deficiencies: manual cataloging inefficiency and error, inadequate demand forecasting, and absence of proactive inventory monitoring.

Empirical evaluation demonstrated product recognition accuracy of 91.4% using Tesseract OCR with fuzzy matching algorithms, a 77.1% cataloging time reduction, and MLR forecasting MAPE of 16.2% representing a 33.9% improvement over moving-average baselines. Six-month partner store deployment yielded a 34.2% reduction in stockout incidents, 28.7% decrease in excess inventory value, and 65% reduction in manual inventory management effort. Dashboard usability ratings averaging 4.6/5 confirmed that the system is practically accessible to non-technical operators. Implemented entirely using open-source software with estimated costs of INR 40,000–65,000 per deployment, the findings demonstrate that advanced AI-driven inventory management is economically feasible and practically beneficial for the small retail sector.

X. FUTURE SCOPE

Several productive directions for extending this work are identified. Enhanced OCR capabilities through integration of deep learning-based OCR engines such as PaddleOCR or EasyOCR may improve recognition accuracy for products with challenging text layouts or multilingual packaging. Extended forecasting models exploring Random Forest, XGBoost, and LSTM networks may better capture complex non-linear demand patterns. Real-time inventory tracking through

continuous video-based shelf monitoring using computer vision would eliminate the need for manual stock-taking. Point-of-Sale Integration via direct API connectivity with platforms such as Square, Shopify, and Tally would enable fully automated sales data ingestion. Mobile Application development and Multi-store Support through extension of the architecture to multi-store deployments with centralized analytics would accommodate growing retail chains. Each extension would enhance practical value while maintaining the core commitment to affordability and accessibility for small retail enterprises.

REFERENCES

- [1]. R. H. Ballou, *Business Logistics/Supply Chain Management*, 5th ed. Upper Saddle River, NJ: Prentice Hall, 2004.
- [2]. G. Wang, A. Gunasekaran, and E. W. T. Ngai, "Distribution network design for fast moving consumer goods," *International Journal of Production Research*, vol. 56, no. 1–2, pp. 438–461, 2018.
- [3]. T. M. Choi, H. K. Chan, and X. Yue, "Recent development in big data analytics for business operations," *IEEE Transactions on Cybernetics*, vol. 47, no. 1, pp. 81–92, 2017.
- [4]. R. Smith, "An overview of the Tesseract OCR engine," in *Proc. Ninth Int. Conf. Document Analysis and Recognition*, 2007, pp. 629–633.
- [5]. J. W. Ratcliff and D. E. Metzener, "Pattern matching: The gestalt approach," *Dr. Dobb's Journal*, vol. 13, no. 7, pp. 46–72, 1988.
- [6]. R. J. Hyndman and G. Athanasopoulos, *Forecasting: Principles and Practice*, 2nd ed. Melbourne: OTexts, 2018.
- [7]. S. Makridakis, E. Spiliotis, and V. Assimakopoulos, "Statistical and machine learning forecasting methods: Concerns and ways forward," *PLOS ONE*, vol. 13, no. 3, p. e0194889, 2018.
- [8]. H. Chen, R. H. L. Chiang, and V. C. Storey, "Business intelligence and analytics: From big data to big impact," *MIS Quarterly*, vol. 36, no. 4, pp. 1165–1188, 2012.
- [9]. A. Abid et al., "Gradio: Hassle-free sharing and testing of ML models in the wild," in *Proc. ICML Workshop on Human in the Loop Learning*, 2019.
- [10]. F. Kache and S. Seuring, "Challenges and opportunities of digital information at the intersection of big data analytics and supply chain management," *Int. Journal of Operations & Production Management*, vol. 37, no. 1, pp. 10–36, 2017.
- [11]. M. A. Waller and S. E. Fawcett, "Data science, predictive analytics, and big data: A revolution that will transform supply chain design and management," *Journal of Business Logistics*, vol. 34, no. 2, pp. 77–84, 2013.