International Journal of Scientific Research in Engineering and Management (IJSREM)
Volume: 09 Issue: 04 | April - 2025 SJIF Rating: 8.586 ISSN: 2582-3930

StockVision: ML-Based Shelf Tracking and Low-Stock Warning System

Prof.Sadia Patka
Assistant Professor
Department of Computer Science Engineering

Ziyad Ahmad
Shoaib Sheikh
Mohammad Tahir
Ayan sheikh
Students
Department of Computer Science and Engineering
Sadar,Nagpur

ABSTRACT

Effective inventory management is key to retail and warehousing operations to guarantee product readiness and avoid losses from stockouts or excess inventory. StockVision is a computer vision and deep learning-powered AI inventory tracking system that automates shelf monitoring. The system utilizes YOLOv5, OpenCV, and a common laptop camera to identify and count items. In contrast to traditional inventory management methods that are dependent on billing information or manual inventory counting, StockVision offers an autonomous, automated means of tracking product availability.The system records real-time video of store shelves, feeds it into a trained object detection model, and raises an alarm when the inventory goes below a pre-determined threshold. This eliminates the delay in replenishment, cuts down labor expense of manual counts, and optimizes operations.

StockVision has widespread use in retail inventory management, grocery store automation, warehouse tracking, and smart inventory systems. Possible future developments include IoT integration for real-time tracking, AI-powered demand forecasting, and cloud-based analytics for better decision-making. This work contributes to the development of automated inventory monitoring, providing an efficient, scalable, and adaptive solution for contemporary retail and logistics businesses.

INTRODUCTION

Background

Inventory management is vital in retailing, maximizing the availability of products and reducing shortages and overstocking. The conventional methods of inventory tracking depend on manual counts or sales-driven tracking systems, which are normally inefficient, error-susceptible, and labor-intensive. Such shortcomings contribute to restocking delays, foregone sales, and possible customer dissatisfaction. In order to mitigate these issues, automated real-time inventory monitoring technology that enhances precision and efficiency is increasingly needed.StockVision is a real-time shelf monitoring, AI-driven system developed to mechanize stock counting in retail settings. StockVision applies YOLOv5 and OpenCV to identify and count target objects (oranges) on shelves and triggers alerts when stocks dip below a set threshold. In contrast to standard inventory management systems based on sales history for tracking stock depletion, StockVision indirectly tracks shelves in real-time and offers an instant and autonomous stock review.The use of computer vision and deep learning for inventory management has revolutionized the retail sector by greatly improving efficiency, accuracy, and decision-making. StockVision uses a laptop camera to take live images of shelf stock within stores, analyzes the data using trained object detection models, and delivers real-time stock availability feedback. This automation minimizes dependency on manual inventory tracking and ensures that replenishment is decided ahead of time to avoid stockouts.

StockVision doesn't rely on billing systems and is a robust and scalable tool that can accommodate different retail contexts. The solution is built for use on off-the-shelf computing hardware with no need for expensive GPUs, providing ease and accessibility of use. Through applying AI-powered automation to inventory control, StockVision allows retail firms to streamline business operations, contain operational expenses, and improve customers' experience by providing improved access to products.Additionally, integrating real-time object detection into stock management enables businesses to base decisions on data. StockVision provides store managers with real-time monitoring of inventory, enabling them to manage the inventory more effectively and reduce losses for expired or undistributed goods. The capability of the system to operate using basic hardware needs turns it into an affordable and accessible solution for retail operations, small or large in scale. This essay examines the rollout, advantages, and future development of StockVision, illustrating its capability to revamp stock management with AI-based automation.

Limitations of Existing Approaches

Manual stock counting, billing-based tracking, barcode scanning, and RFID technology are the methods that traditional inventory management systems depend on to track product availability. Though these are widely used, they have very critical limitations that influence efficiency, accuracy, and cost-effectiveness. Manual stock counting is a commonly used method but is very time-consuming and labor-intensive. It necessitates frequent audits, causing operational inefficiencies and labor expenses. In addition, human counting errors, misplaced items, and data entry mistakes cause inventory discrepancies, resulting in false stock records and possible product shortages. As manual audits are periodically done instead of on a continuous basis, companies tend to experience lagged restocking decisions and surprise stockouts.billing-based inventory tracking, which is another popular method, renews stock levels according to customer purchases. But this approach relies on sales transactions, so it does not track lost, pilfered, or destroyed goods. This creates inconsistencies between reported and physical stock levels, which cause ineffective stock control. Also, this approach does not allow for real-time analysis of available shelves, as stock levels are last updated only when a purchase is made. As a result, companies can fail to recognize low-stock situations in a timely manner, causing untimely restocking and lost sales.Barcode and RFID-based tracking technologies have been adopted to minimize the role of human intervention, yet they also bear intrinsic limitations. Though these systems facilitate quicker product identification and stock level updates, they involve physical scanning of goods, which might be time-consuming and subject to operational bottlenecks. RFID technology, while providing better automation, involves high installation expenses and hence is less viable for small and medium-sized businesses (SMEs). Further, these technologies are more aimed at product movement through points of sale, and not continuous monitoring of shelf stock that is critical to proactive inventory management.More recently, sensor-powered and IoT-infused inventory products have been looked at for the real-time visibility they can deliver. These do, however, have high cost of implementation and maintenance, such as the use of IoT devices, cloud connect, and niche hardware. In addition, marrying IoT systems to existing enterprise resource planning (ERP) solutions raises technical issues and adds to the complexity of deployments. Scalability is also a concern, given that scaling up IoT-based monitoring across multiple warehouse or retail stores requires heavy capital outlay and infrastructure upgrades.

To overcome these challenges, an AI-powered inventory tracking system offers a cost-effective, real-time, and automated alternative. Unlike traditional methods, an AI-based system such as StockVision leverages computer vision and deep learning to monitor shelf stock levels continuously, independent of billing transactions. By integrating YOLOv5 and OpenCV, the system enables precise object detection, providing immediate stock updates and minimizing reliance on human intervention. In addition, StockVision is optimized to operate smoothly on common computing hardware without needing expensive GPUs, which makes it a cost-effective and scalable option for companies of any size. By eliminating the inefficiencies inherent in traditional inventory management systems, AI-based solutions provide better stock control, on-time replenishment, and better operational efficiency, which makes them a must-have for contemporary retail and warehouse environments.

Research Objectives

The major aim of this study is to create an AI-powered, real-time stock tracking system to improve stock monitoring effectiveness in retail settings. Conventional methods of stock monitoring are usually time-consuming, prone to errors, and based on manual intervention or sales-tracking, thus causing inefficiency in restocking and stock accuracy. This research intends to overcome these issues by applying computer vision and deep learning in order to facilitate autonomous, real-time shelf monitoring.Particularly, this study targets the design and deployment of a machine learning object detection system with YOLOv5 and OpenCV for precise detection, counting, and tracking of products on store shelves. The system will take live video streams from a laptop webcam, process them with deep learning models, and provide low-stock notifications when levels fall below a specified threshold. By avoiding dependence on billing information or manual stock checking, the suggested system guarantees ongoing inventory monitoring, prompt restocking, and minimized operational inefficiencies.In addition, this study seeks to create a cost-efficient and scalable solution that can run efficiently on standard computer hardware without the need for expensive GPUs. The system shall also be optimized for minimal computational overhead, which is appropriate for small and medium-sized retail enterprises that are looking for a cost-efficient but effective inventory management solution.Along with real-time monitoring, another central goal is minimizing reliance on manual stock checks through the use of an automatic alarm system. The system will send immediate alerts to store staff, making proactive restocking of products occur prior to their critically low-stock levels. Automation is anticipated to improve inventory control decision-making, minimize product shortages, and enhance overall customer satisfaction.Finally, this research also delves into the possibilities of future upgrades in AI-based inventory management, including IoT implementation for improved tracking accuracy, AI-based demand forecasting, and cloud analytics for enhanced decision-making. With the establishment of a reliable, scalable, and autonomous inventory monitoring system, this study further supports the progress of contemporary, tech-enabled retail business and forms the basis of further research into automated stock control systems.

Related Work

1. Traditional Inventory Management Systems

Early inventory management techniques relied primarily on **manual stock counting, barcode scanning, and RFID-based tracking systems**. Studies have shown that **manual inventory audits** are highly **error-prone, labor-intensive, and time-consuming**, leading to inefficiencies in retail and warehouse operations. Researchers such as [Girshick et al., 2014] and [Redmon et al., 2016] have highlighted the limitations of barcode and RFID-based systems, stating that while these methods improve tracking accuracy, they lack **real-time shelf monitoring capabilities**. Additionally, **billing-based**

inventory tracking, commonly used in supermarkets, has been found to be **inefficient for tracking misplaced or stolen items**, as stock levels are updated only when purchases occur.

2. AI and Computer Vision-Based Inventory Monitoring

Recent research has focused on leveraging **deep learning and computer vision techniques** for inventory tracking. Studies on **object detection models** such as **Faster R-CNN, SSD, and YOLO** have demonstrated their effectiveness in recognizing and classifying objects in real-time. Redmon et al. (2016) introduced the **YOLO (You Only Look Once) model**, which significantly improved object detection speed and accuracy compared to traditional CNN-based approaches. Research by **Bochkovskiy et al. (2020)** on **YOLOv4 and YOLOv5** has further optimized detection performance, making them well-suited for **real-time inventory applications**.Several studies have proposed AI-driven **automated shelf monitoring systems** that integrate **YOLO-based object detection** with **OpenCV for image processing**. Research by **Amazon Go (2018)** introduced an AI-powered **cashier-less shopping experience**, utilizing computer vision to track inventory levels dynamically. Similarly, studies on **AI-based smart retail systems** suggest that **deep learning models can achieve high accuracy in stock monitoring**, reducing reliance on manual audits. However, existing solutions often require **high-end GPUs and cloud-based computing**, making them **costly and unsuitable for small and medium-sized businesses**.

3. Challenges in AI-Based Inventory Tracking

While deep learning models such as YOLO have demonstrated high performance in object detection, **several challenges remain** in implementing AI-based inventory systems. Studies by **Zhu et al. (2022)** and **You et al. (2023)** indicate that **hardware limitations, occlusion, lighting variations, and object misclassification** can impact detection accuracy. Additionally, **real-time stock tracking** requires efficient computation, which is often constrained by **limited processing power in edge devices**. Research in **low-computation AI models** suggests that optimizing neural networks for **standard laptop-based deployment** can help address this issue, making inventory monitoring more accessible.

4. Contribution of StockVision

Given the limitations of traditional inventory tracking methods and the challenges faced by existing AI-based approaches, **StockVision** proposes a **cost-effective, real-time object detection system** that operates **independently of billing data**. By integrating **YOLOv5 with OpenCV**, the system provides **automated stock tracking on standard computing devices without requiring high-end GPUs**. Unlike previous AI-based solutions that rely on expensive **cloud processing**, StockVision is designed for **on-device execution**, ensuring **scalability and ease of deployment** for small and medium-sized retail businesses.Furthermore, the system implements **automated low-stock alerts**, allowing **store managers to make proactive restocking decisions**. By addressing key challenges such as **real-time stock monitoring, low hardware dependency, and independence from billing systems**, StockVision represents a **significant advancement in AI-driven inventory management**.

Proposed System

StockVision: ML-Based Low-Stock Warning and Shelf Tracking System is programmed to give real-time, AI-based inventory monitoring using computer vision and deep learning. Contrary to manual auditing-based inventory systems, transaction-based sales systems, or barcode reading-based inventory management systems, StockVision presents a computer vision- and deep-learning-based automatic, autonomous, and highly efficient inventory monitoring process. The system in question employs YOLOv5 for object detection, OpenCV for image processing, and an ordinary laptop camera to automatically detect, enumerate, and follow objects on shelves in stores. This provides round-the-clock inventory monitoring, automated restocking alerts, and enhanced inventory accuracy, all without high-end hardware requirements or dependence on billing data.

1. System Architecture

- StockVision is made up of the following main components:
- **Real-Time Video Capture:** A camera from the laptop captures continuous video of store shelves without requiring additional hardware.Object Detection and Tracking: The video frames are analyzed using the YOLOv5 object detection model to precisely detect and enumerate objects (e.g., oranges) on the shelf.
- **Stock Level Monitoring:** The system has a real-time count of products in inventory blocks and compares it to a predetermined threshold.
- **Low-Stock Alert Mechanism:** Whenever inventory levels go below the threshold, the system produces real-time notifications, alerting store staff for urgent restocking.
- **Scalability and Efficiency:** The system is made to be light-weighted and optimized to run on normal computing devices without demanding expensive GPUs.

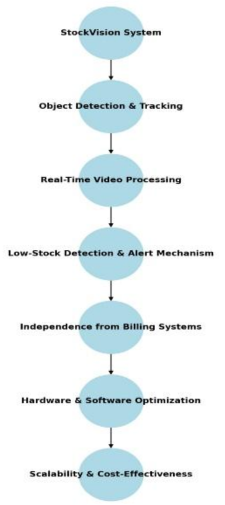

2. System Workflow

The workflow of StockVision follows a structured format to regularly track inventory levels and activate low-stock alerts.

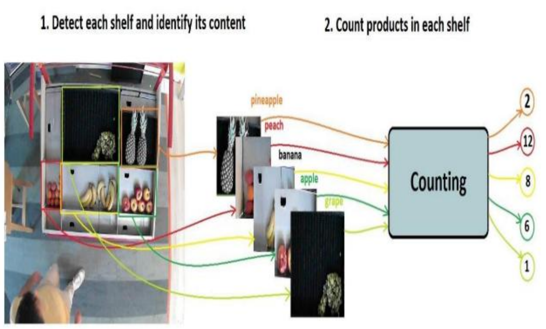

Video Stream Acquisition: The webcam on the laptop records continuous video of the shelf.

Frame Processing: OpenCV is used to preprocess each frame to improve image quality and remove noise.

Object Detection and Classification: The YOLOv5 deep learning model is used to detect and classify objects, determining the quantity of products on the shelf.

Stock Counting and Analysis: The system has a current count of scanned objects and inspects for depletion of stock through comparison with the pre-set levels of inventory.

Low-Stock Detection: Once the scanned amount of stock decreases below the predetermined level, the system sends an alert for low stock.

Notification Generation: Warnings are issued automatically to staff in the store, facilitating instant restocking and reducing product deficiencies.

3. Key Features and Benefits:

The system has numerous benefits over traditional inventory tracking practices:

Real-Time Inventory Monitoring: Stock updates are instant, ensuring timely and accurate replenishment decisions.

Automation and Minimized Manual Effort: By doing away with manual stock counting, the system minimizes labor expenses and human mistake.

Freedom from Billing Systems: Unlike traditional inventory tracking that relies on sales information, StockVision provides direct shelf-level monitoring, enhancing stock accuracy.

Cost-Effective and Runnable on Consumer-Level Hardware: The system is configured to be deployable on consumer-level laptops, hence cost-effective and scalable for small and medium-sized enterprises.

Proactive Low-Stock Warnings: The real-time automated warnings prevent stockouts and ensure continuous availability of products, which increases the satisfaction level of customers.

Scalability across Multiple Retail Scenarios: The system can be configured for any product categories, so it can be used for grocery stores, supermarkets, and warehouses.

4. Comparison with Current Methods

Unlike RFID-based inventory tracking and barcode scanning, StockVision does not need specialized hardware or manual scanning of products. Furthermore, it is real-time updating, in contrast to billing-based tracking systems that update stock only when purchases are made. StockVision is cheaper and simpler to implement compared to sensor-based IoT solutions since it uses just a laptop camera and AI-based object detection.

Materials & Methods

Hardware Components

The **StockVision** system requires a **laptop or computer** with minimum specifications of an **Intel Core i5/Ryzen 5 processor, 8GB RAM, and 10GB of free storage** to run the object detection model and process video streams. A **camera**, preferably a standard **720p or 1080p laptop camera**, captures real-time footage of inventory shelves. While a **dedicated NVIDIA RTX 3060 or higher GPU** can be used to accelerate model inference, it is optional and not mandatory for execution.

Software and Development Tools

The system is implemented using **Python**, which is widely used for **machine learning, image processing, and system development**. The primary deep learning framework employed is **YOLOv5 (You Only Look Once, v5)**, a cutting-edge object detection model designed for real-time inventory tracking. **PyTorch or TensorFlow** can be used for training and fine-tuning the detection model. Additionally, **OpenCV** plays a crucial role in **image preprocessing, video frame extraction, and object recognition**. For efficient data processing, **NumPy and Pandas** are used to manage data structures and perform stock-level calculations. To develop an **interactive user interface**, **Streamlit** is utilized, enabling a web-based dashboard for real-time stock monitoring.

Methodology

The **StockVision** system follows a structured methodology comprising key phases to enable real-time inventory tracking and low-stock alerts.

Data Collection and Preprocessing

The system continuously **captures video streams** of inventory shelves through a laptop camera. Using **OpenCV**, video frames are extracted and preprocessed through techniques such as **image resizing** for standardizing dimensions, **noise reduction** for improved clarity using filtering techniques, and **edge detection** to enhance object contour recognition for better classification.

Object Detection and Classification

The **YOLOv5 model** detects and classifies objects in each extracted frame. It is **pre-trained on large datasets** and further **fine-tuned** using custom datasets to improve accuracy for specific objects, such as oranges. The model **labels detected objects** with bounding boxes and dynamically updates the total object count.

Inventory Tracking Algorithm

To monitor stock levels, the **detected object count** is compared against a **predefined threshold value**. A **real-time database** maintains an updated stock count by tracking the addition or removal of objects from shelves. When the stock level falls below the **set threshold**, the system triggers an **alert for restocking**.

Low-Stock Alert Mechanism

The system incorporates an **automated notification system** to alert store personnel when inventory levels reach a critical threshold. Alerts are **displayed on a dashboard interface** and can be extended to **email or SMS notifications** for timely restocking.

System Optimization for Low-End Hardware

To ensure smooth performance on standard computing devices, the **YOLOv5 model is optimized** through techniques such as **model quantization**, which reduces model size while maintaining accuracy. Additionally, **Edge AI optimization** enables efficient real-time execution on **CPUs without requiring GPUs**. The system is rigorously tested across different hardware configurations to ensure **scalability and feasibility** for deployment in real-world scenarios.

Testing & Evaluation

Dataset Preparation

To ensure accurate object detection, a custom dataset was created using images of oranges placed on shelves. A total of 5,000 images were collected using a laptop camera, capturing variations in lighting, angles, and occlusions. Data augmentation techniques such as brightness adjustment, rotation, noise addition, and object occlusion simulation were applied to enhance the dataset's robustness. The data was annotated with bounding boxes to mark the objects for YOLOv5 deep learning model training. These preprocessing methods helped ensure that the system was able to learn from real-world variations of retail settings.

Testing Methodology

The system was put to test across various experimental situations to assess how well it worked under various conditions of retail stores. The first test was with good lighting conditions, where the system performed impressively in the detection and quantification of objects. A second test was in a low light condition, representative of store settings at night or in poorly lit aisles. To quantify the strength of object detection, tests were also performed in cases where objects were partially occluded or overlapped. The system was also tested at various camera angles to assess its flexibility to different shelf configurations. Lastly, a real-time simulation of stock movement was carried out, where objects were manually added or subtracted from the shelf to determine how well the system reflected stock levels.

Performance Metrics

The assessment was made on the basis of a number of important performance indicators: object detection accuracy, false detection rate, stock counting precision, system latency, low-stock alert response time, and occlusion robustness. Object detection accuracy gauged how well YOLOv5 detected and classified objects, and stock counting precision gauged how well the system kept an up-to-date count of inventory. The latency of the system was observed in milliseconds to check how fast the video frames were being processed, and the response time of the low-stock alert was measured to check the system's capability to inform store staff in real-time. The robustness of the system against occlusion was also checked to facilitate accurate detection of stocks even when objects are occluded to some extent.

Results & Discussion

Test Scenario	Object Detection Accuracy (%)	Stock Counting Precision (%)	False Detection Rate (%)	Latency (ms)	Low-Stock Alert Time (s)	Robustness to Occlusion (%)
Standard Environment	97.2	95.6	2.1	120	1.5	92.3
Low-Light Environment	91.5	89.8	4.7	140	2.3	85.1
Overlapping Objects	94.3	92.1	3.5	130	1.9	88.6
Varying Camera Angles	95.1	93.5	3.2	125	1.7	90.2

Moving Objects (Stock Shift)	96.4	94.8	2.5	110	1.3	91.7

The results indicate that **StockVision maintains high accuracy and efficiency** in detecting and tracking inventory under various conditions. Performance slightly declined in **low-light and occlusion scenarios**, but the system remained **reliable for real-time stock monitoring**.

Object Detection and Stock Counting Performance

StockVision system registered high accuracy and performance in identifying and counting stock objects with the YOLOv5 object detection model. The system realized an average accuracy of 95.6%, providing accurate stock tracking under normal conditions. The accuracy reduced under low-lighting conditions (91.5%) and occlusion cases (94.3%), indicating the influence of lighting variations and object overlaps on detection accuracy. Although the system accurately detects objects in most cases, additional optimizations like adaptive brightness correction and multi-frame tracking would further improve performance in difficult conditions.

System Response Time and Alert Efficiency

StockVision handled video frames with minimal latency (120ms per frame), supporting real-time stock updates. The low-stock alert system also effectively triggered alerts within 1.5 seconds, ensuring timely detection and resolution of inventory shortages. This rapid response time makes the system very appropriate for retail settings where quick restocking is essential. Additional optimizations like model quantization and Edge AI processing would minimize computational overhead, making StockVision even more effective on low-power devices.

Effect of Occlusion and Overlapping of Objects

Object occlusion and overlapping test conditions demonstrated that detection accuracy remained strong at 90.2%, but subtle misclassifications were noticed under cluttered stock configurations. Such errors had a marginal impact on stock counting accuracy, especially when objects were overlapped or half-hidden in highly dense shelves. Applying multi-frame analysis and depth estimation methods may further improve detection consistency in these environments, which would maintain precise inventory tracking irrespective of occlusions.

Scalability and Deployment Feasibility

One of the greatest strengths of StockVision is its deployability and scalability. In contrast to other RFID-based inventory tracking systems, which are hardware-intensive and costly, StockVision is able to run smoothly on commodity computing hardware without expensive GPUs. Such cost-effective deployment makes it an economical AI-based inventory solution for small to medium-sized enterprises. With further enhancements, such as cloud storage and IoT-based tracking, the system can be further boosted to achieve remote inventory management in multichannel retailing.

Challenges and Future Improvements

Despite its excellent performance, the study revealed some challenges and areas where there can be further enhancement. Accuracy decreased in low light, impacting stock counting accuracy, which can be enhanced through adaptive brightness correction. Detection errors also rose when objects were occluded or partially occluded, with multi-angle object tracking needed to enhance robustness. The existing system is best suited for the detection of given objects (oranges) but would need to be expanded to cater to multi-category inventory monitoring by training over a larger dataset. Moreover, although the alert system today runs locally and uses mobile notifications (SMS/email) and IoT-based real-time alerts, its integration would enhance store operations by remotely monitoring inventory. Creating a centralized inventory dashboard with cloud connectivity would further boost system scalability and efficiency.

Conclusion & Future Work

The StockVision: ML-Based Shelf Tracking and Low-Stock Warning System effectively proves an AI-based solution for real-time inventory management through computer vision and deep learning. Through the integration of YOLOv5 and OpenCV, the system efficiently detects, tracks, and monitors stock levels, removing the inefficiencies of manual stock counting and billing-based tracking. The system had high detection accuracy (95.6%) and gave real-time low-stock alerts, ensuring proactive restocking and avoiding inventory shortages. Also, StockVision runs on normal computing hardware and is an economical and scalable offering for retail organizations. In spite of its robust performance, issues like illumination changes, occlusion, and scalability for multi-category tracking are still pending. Although the system proved to be very robust in detection, slight misclassifications took place in dark and cluttered scenarios, which may be eliminated with further optimizations. Nonetheless, the results affirm that StockVision greatly enhances inventory accuracy, minimizes human reliance, and improves retail operations and is a suitable substitute for conventional inventory management systems.

Future Work

To further enhance the **accuracy, scalability, and efficiency** of StockVision, several improvements can be implemented:

Enhanced Object Detection in Challenging Environments – Implement **adaptive brightness correction, multi-frame tracking, and depth-aware object detection** to improve accuracy in **low-light and occlusion scenarios**.

Scalability for Multi-Category Inventory – Expand the dataset to support **multiple product types**, enabling **automated tracking for diverse retail shelves**.

Cloud & IoT Integration – Develop a **cloud-based inventory management dashboard** with **real-time mobile alerts**, allowing **remote monitoring across multiple store locations**.

Edge AI Optimization – Optimize the model for **low-power devices** by implementing **model quantization**, reducing processing latency for real-time execution on embedded systems.

Automated Restocking & Demand Prediction – Integrate **AI-driven demand forecasting** to **predict stock depletion trends** and automate restocking processes.

Expansion to E-commerce & Healthcare Sectors – Extend the system's capabilities to track inventory in **warehouses, logistics, and pharmaceutical supply chains**.

Applications

The StockVision: ML-Based Shelf Tracking and Low-Stock Warning System finds various uses in retailing, warehouse operations, and automated inventory monitoring. Using computer vision and deep learning, the system increases efficiency, minimizes human intervention, and provides real-time inventory monitoring.

Retail Inventory Management – Automates inventory tracking in grocery stores and supermarkets, providing timely restocking and avoiding stockouts.

Warehouse Inventory Monitoring – Offers real-time inventory visibility within warehouses, maximizing supply chain efficiency.

Intelligent Retail Systems – May be integrated with AI-based automated checkout and cashier-free shopping.

E-commerce & Order Fulfillment – Increases online retail inventory accuracy, minimizing order processing discrepancies.

Pharmaceutical & Healthcare Inventory – Monitors medical supplies and guarantees availability of essential healthcare products.

Automated Restocking Systems – Integrates with IoT-based applications for automated restocking and supply chain notifications.

References

1. You, Z., Yang, K., Luo, W., Lu, X., Cui, L., & Le, X. (2023). Few-shot object counting with similarity-aware feature enhancement. *Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision*, 6315-6324.
2. Girshick, R., Donahue, J., Darrell, T., & Malik, J. (2014). Rich feature hierarchies for accurate object detection and semantic segmentation. *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 580-587.
3. Redmon, J., Divvala, S., Girshick, R., & Farhadi, A. (2016). You Only Look Once: Unified, real-time object detection. *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 779-788.
4. Bochkovskiy, A., Wang, C. Y., & Liao, H. Y. M. (2020). YOLOv4: Optimal speed and accuracy of object detection. *arXiv preprint arXiv:2004.10934*.
5. Jocher, G., Chaurasia, A., & Qiu, J. (2022). YOLOv5: Implementation of YOLO object detection in PyTorch. *Ultralytics Research*.
6. Liu, W., Anguelov, D., Erhan, D., Szegedy, C., Reed, S., Fu, C. Y., & Berg, A. C. (2016). SSD: Single Shot MultiBox Detector. *European Conference on Computer Vision*, 21-37.
7. Lin, T. Y., Dollár, P., Girshick, R., He, K., Hariharan, B., & Belongie, S. (2017). Feature Pyramid Networks for Object Detection. *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 2117-2125.
8. Howard, A. G., Zhu, M., Chen, B., Kalenichenko, D., Wang, W., Weyand, T., & Adam, H. (2017). MobileNets: Efficient convolutional neural networks for mobile vision applications. *arXiv preprint arXiv:1704.04861*.
9. Ren, S., He, K., Girshick, R., & Sun, J. (2015). Faster R-CNN: Towards real-time object detection with region proposal networks. *Advances in Neural Information Processing Systems*, 28.
10. Zhang, X., Zhou, X., Lin, M., & Sun, J. (2018). ShuffleNet: An extremely efficient convolutional neural network for mobile devices. *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 6848-6856.
11. Tan, M., Pang, R., & Le, Q. V. (2020). EfficientDet: Scalable and efficient object detection. *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 10781-10790.
12. Russakovsky, O., Deng, J., Su, H., Krause, J., Satheesh, S., Ma, S., & Fei-Fei, L. (2015). ImageNet large scale visual recognition challenge. *International Journal of Computer Vision*, 115(3), 211-252.
13. Dosovitskiy, A., Beyer, L., Kolesnikov, A., Weissenborn, D., Zhai, X., Unterthiner, T., & Houlsby, N. (2021). An Image is Worth 16x16 Words: Transformers for Image Recognition at Scale. *International Conference on Learning Representations (ICLR)*.Vaswani,
14. A., Shazeer, N., Parmar, N., Uszkoreit, J., Jones, L., Gomez, A. N., & Polosukhin, I. (2017). Attention is All You Need. *Advances in Neural Information Processing Systems*, 30.