AI-Driven Automated Quality Inspection for Beverage Bottles: Leveraging Object Detection Models for Enhanced Supply Chain Efficiency

Deepak Raj R[1]; Sowmiya R[2]; Swathi K[3]; Harikaran G[4]; Gayathri K[5]; Ezhil Litta A[6]; Vishvash C[7]; Bharani Kumar Depuru[8]

[1,2,3,4,5,6] Research Associate, AISPRY Pvt Ltd, Hyderabad, India.
[7] Mentor, Research and Development, AISPRY Pvt Ltd, Hyderabad, India.
[8] Director, AISPRY Pvt Ltd, Hyderabad, India

Abstract: In the beverage industry, maintaining product quality during packaging and throughout the supply chain is critical to ensuring customer satisfaction and brand integrity. This research addresses the challenge of automating quality inspection for beverage bottles by leveraging cutting-edge AI-based object detection models. The study focuses on identifying and classifying six key quality defects particularly Cracked_Bottle, Misaligned_Label, Missing_Cap, Normal_ Bottle, Overfilled_Bottle, and Underfilled_Bottle. These defects, if undetected, can lead to customer dissatisfaction, increased return rates, and potential brand damage.

To tackle this problem, we implemented and evaluated three advanced object detection architectures—YOLOv8, YOLOv9, and YOLOv11—on a custom dataset comprising thousands of images of beverage bottles captured under diverse conditions, including varying lighting, angles, and backgrounds. Among the models, YOLOv8 emerged as the most effective, achieving an impressive 78% accuracy across all defect classes. The model demonstrated exceptional performance in detecting subtle defects such as misaligned labels and minor cracks, which are often overlooked in manual inspections.

The integration of AI-driven quality control systems into the beverage supply chain not only minimizes human error but also significantly enhances operational efficiency. By automating the detection of defects, this approach ensures that only products meeting stringent quality standards reach consumers. Furthermore, the system provides real-time feedback, enabling swift corrective actions and reducing waste. This research underscores the transformative potential of AI in revolutionizing quality assurance processes within the beverage industry, ultimately driving customer trust, reducing costs, and improving overall supply chain performance.

***Keywords**: Beverage Quality Assurance, Defect Detection, Yolov8, Yolov9, Yolov11, AI In Packaging, Supply Chain Optimization, Automated Quality Inspection.*

How to Cite: Deepak Raj R; Sowmiya R; Swathi K; Harikaran G; Gayathri K; Ezhil Litta A; Vishvash C; Bharani Kumar Depuru (2025)AI-Driven Automated Quality Inspection for Beverage Bottles: Leveraging Object Detection Models for Enhanced Supply Chain Efficiency *International Journal of Innovative Science and Research Technology*, 10(3), 2773-2782. https://doi.org/10.38124/ijisrt/25mar1796

I. INTRODUCTION

The beverage industry is a critical sector of the global economy, with millions of products manufactured and distributed daily. However, maintaining consistent product quality throughout the supply chain remains a significant challenge. Traditional quality control methods, which rely heavily on manual inspection, are not only labor-intensive but also prone to human error. These limitations often result in defective products reaching consumers, leading to increased return rates, financial losses, and a decline in brand trust.

To address these challenges, this study leverages Artificial Intelligence (AI) and computer vision technologies [2] to automate the quality inspection [3] process for beverage bottles. By incorporating AI-powered object detection models into the production line, manufacturers can efficiently identify and separate defective bottles from intact ones, ensuring that only high-quality products reach the market. This not only minimizes human error but also reduces the manpower required for quality control, leading to significant cost savings and improved operational efficiency.

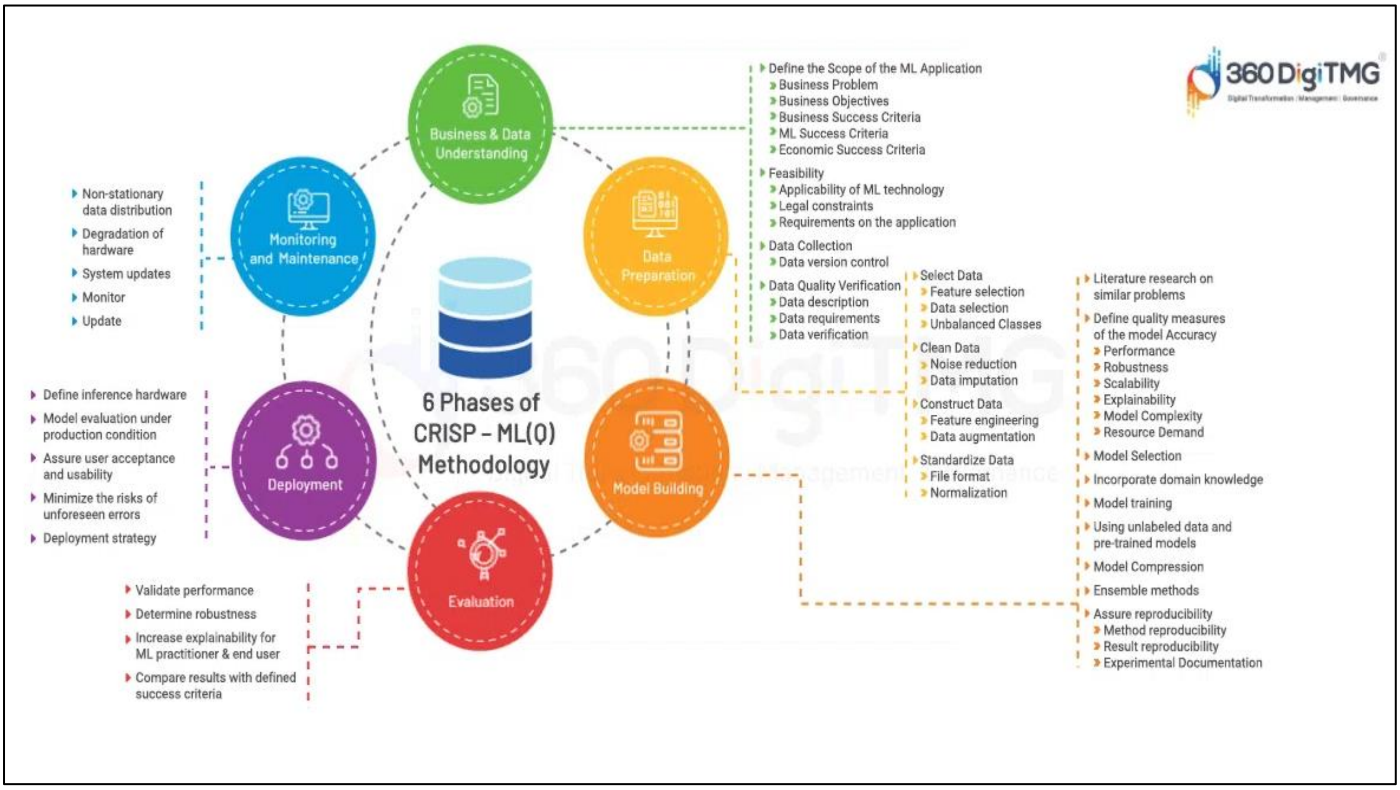

Fig 1 CRISP - ML(Q) - Project Methodology

This study adopts the CRISP-ML(Q) [1] methodology [Fig.1], a structured framework for developing machine learning solutions. The process begins with Business Understanding, where the primary goal is to automate defect detection in beverage bottles to enhance quality control. In the Data Understanding phase, a custom dataset is collected, comprising images of bottles with various defects, captured under different lighting, angles, and backgrounds. The Data Preparation phase involves preprocessing the dataset, including resizing, augmentation, and annotation, to ensure it is suitable for training. Key preprocessing steps include resizing images to 640x640, horizontal flipping, 90° rotation (clockwise and counter-clockwise), shear adjustments (±10° horizontal and vertical), saturation and brightness adjustments (±15%), and adding noise and blur (up to 0.5px and 0.5% of pixels, respectively).

In the Model Building phase, state-of-the-art object detection models [4] such as YOLOv8, YOLOv9, and YOLOv11 are trained and evaluated. These models are implemented using the Ultralytics framework [5], with training conducted on an AWS EC2 g4dn Xlarge instance running Python 3.10. The Model Evaluation phase uses metrics like precision, recall, and mean average precision (mAP) to assess performance [6]. The best-performing model is then deployed using the Streamlit framework on an AWS EC2 t2 Xlarge instance for real-time defect detection [Fig 2] talks on the overall flow of the project. Finally, the Monitoring and Maintenance phase ensures the system's ongoing performance and adaptability to new data.

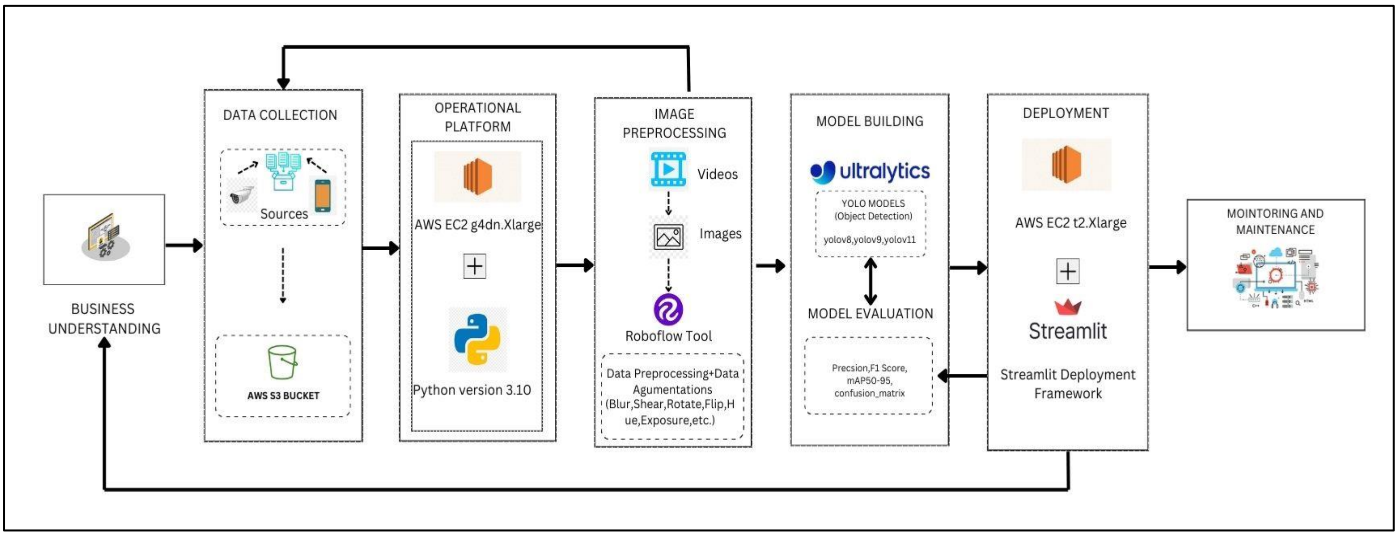

Fig 2 Project Architecture Diagram

ISSN No:-2456-2165 https://doi.org/10.38124/ijisrt/25mar1796

The integration of AI into the beverage supply chain offers numerous benefits. First, it reduces the reliance on manual labor, which is often inconsistent and error-prone. Second, it enables real-time detection of defects, allowing for immediate corrective actions and reducing waste. Third, it provides valuable data insights that can be used to optimize production processes and improve overall product quality. By addressing these challenges, this research aims to contribute to the advancement of AI-driven quality assurance systems in the beverage industry, ultimately leading to higher customer satisfaction and stronger brand loyalty.

II. BUSINESS UNDERSTANDING

The beverage industry is a critical sector of the global economy, with millions of products manufactured and distributed daily to meet consumer demand. According to recent market reports, the global beverage market is projected to grow significantly, driven by increasing consumer preferences for convenience and quality. However, maintaining consistent product quality throughout the supply chain remains a significant challenge. Issues such as defective packaging, misaligned labels, and improper filling can severely impact customer satisfaction and brand reputation. These defects, if undetected, can lead to increased return rates, financial losses, and a decline in consumer trust.

In the beverage industry, the physical construction of bottles and the many ways they travel during production and distribution make them susceptible to defects. These defects not only compromise the product's quality but also increase the likelihood of contamination, which can pose health risks to consumers. For instance, cracked bottles can lead to leakage, while misaligned labels can result in incorrect product information reaching the consumer. A strict quality control procedure is necessary since consumers depend on food and beverage products for their safety and quality.

However, the high volume of demand and the complexity of modern supply chains make it challenging to maintain the high quality of beverage production. Manual inspection of bottles is time-consuming and highly susceptible to errors, especially in large-scale operations. Human inspectors may miss subtle defects due to fatigue or inconsistency, leading to defective products reaching the market. This not only affects consumer trust but also increases operational costs due to returns and recalls.

To overcome this challenge [12], one viable solution is to automate the quality inspection process using computer vision technology [7]. Computer vision systems can extract information from images and videos in real-time, enabling the classification and detection of defects with high precision. These systems can be trained to identify and categorize defects [8], ensuring that only high-quality products reach the market. By automating the inspection process, manufacturers can significantly increase efficiency, accuracy, and quality assurance, while reducing the reliance on manual labor.

This study follows a structured approach to automate the quality inspection process for beverage bottles. 1) The first step involves preparing the dataset, which consists of images of bottles captured under varying lighting conditions, angles, and backgrounds. This ensures the model is trained to handle real-world scenarios effectively. 2) Next, the collected data is uploaded to Roboflow, where it undergoes thorough preprocessing, including resizing, augmentation, and annotation, to optimize it for object detection tasks. 3) Object detection models [9] are then built using state-of-the-art architectures such as YOLOv8, YOLOv9, and YOLOv11. These models are designed to detect and classify defects in bottles, ensuring high accuracy and reliability. 4) After the model building phase, the models are evaluated using various metrics such as precision, recall, and mean average precision (mAP). This step is crucial in the machine learning workflow to assess the model's performance and determine the most effective algorithm for the specific task of defect detection in beverage bottles.

III. DATA UNDERSTANDING

The dataset used in this research is collected from primary sources within the beverage production environment. The dataset primarily consists of images of beverage bottles captured under various real-world conditions to account for the natural variation in lighting, camera angles, and backgrounds encountered in a production line or distribution network. The dataset was created to help identify and classify six distinct defects in beverage bottles: Defective_Bottle, Defective_Label, Missing_Label, Normal_Bottle, Over_Filled, and Under_Filled.

➢ *The Initial Dataset Distribution for each Defect Class is as Follows:*

- Defective_Bottle 152 images
- Defective_Label 153 images
- Missing_Label 154 images
- Normal_Bottle 152 images
- Over_Filled 152 images
- Under_Filled 190 images

These image counts were relatively low, especially for some classes, which made it necessary to apply data augmentation techniques to ensure the dataset was large and diverse enough for effective model training. The goal was to augment each class to approximately 1000 images, ensuring the model had enough data to generalize well.

➢ *Dataset Sources*

The primary data for this study were captured using a variety of cameras in the production environment, under varying conditions including different lighting, angles, and backgrounds. The images were taken from various stages of the production and packaging lines, ensuring the dataset represented real-world production scenarios. The variations in lighting, camera distances, and angles ensure that the model is exposed to as many realistic conditions as possible.

The images were labeled according to the type of defect present in the beverage bottle, allowing for the classification of defects such as misaligned labels, cracked bottles, and overfilled or underfilled bottles.

➢ *Class Distribution and Augmentation*

To ensure balanced training and reduce the potential for model bias, several data augmentation techniques [11] were employed. These techniques helped increase the number of images in each class and also improved the robustness of the model by simulating different conditions that may occur during production.

The class distribution [Table 1] of the augmented dataset after applying the data augmentation techniques is as follows:

Table 1 Class Distribution Before and After Data Augmentation

Defect Class	Initial Count	Augmented Count	Final Count
Defective_Bottle	152	861	1013
Defective_Label	153	926	1079
Missing_Label	154	916	1070
Normal_Bottle	152	856	1008
Over_Filled	152	907	1059
Under_Filled	190	834	1024

By augmenting the images to a target of 1000 and above images per class, the final dataset was made sufficiently large to support training of the object detection models [10] effectively.

IV. DATA PREPARATION

In this phase, the gathered images were processed and prepared for model training. The dataset was uploaded to **Roboflow**, a platform that streamlines dataset management for machine learning projects. Roboflow provides tools for data annotation, augmentation, and preprocessing, making it a suitable tool for handling large image datasets efficiently.

➢ *Image Annotation*

The images were manually annotated [Fig 3] using **bounding boxes** to identify and label the defects present in each bottle. These annotations were crucial for the object detection models, as they allow the models to learn the spatial locations of the defects within each image.

Each image was labeled according to the defect present, and the following defect classes were used:

- **Defective_Label**: Bottles with misaligned or damaged labels.
- **Missing_Label**: Bottles with no label.
- **Over_Filled**: Bottles filled above the required level.
- **Defective_Bottle**: Bottles with visible deformations such as cracks or dents.
- **Normal_Bottle**: Bottles with nil defects.
- **Under_Filled**: Bottles filled below the required level

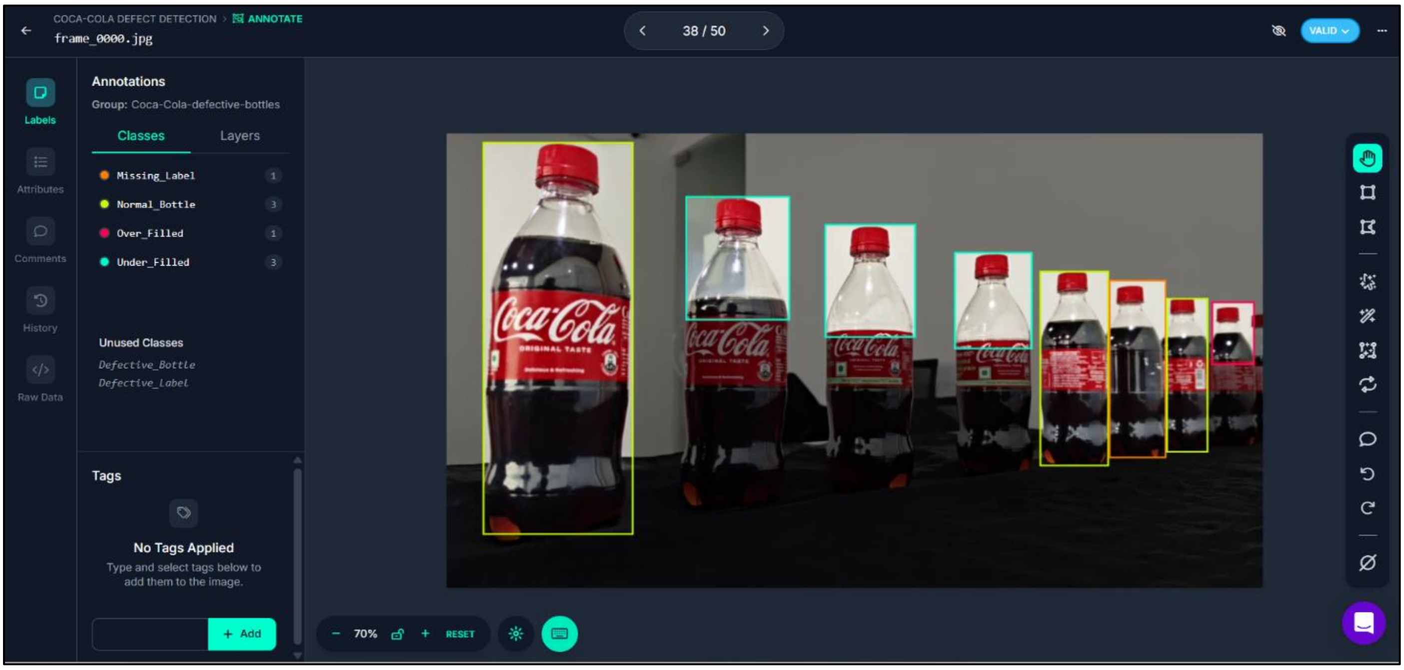

Fig 3 Classes Annotations using Roboflow

➢ *Preprocessing and Augmentation*

Once the images were annotated, several preprocessing and data augmentation techniques were applied to the dataset to enhance its diversity and improve the generalization capability of the models.

➢ *Preprocessing Steps:*

- *Resizing:*

All images were resized to a standard **640x640 pixels** to maintain consistency and facilitate faster processing.

ISSN No:-2456-2165 https://doi.org/10.38124/ijisrt/25mar1796

- *Normalization:*
  The pixel values of all images were normalized to a range of 0 to 1 by dividing each pixel by 255. This normalization ensures that the model can efficiently process the images.

- *Auto-orientation:*
  Images that were not aligned properly (e.g., rotated or upside down) were automatically corrected to ensure consistent orientation.

- *Null Filtering:*
  Images without valid annotations (null images) were filtered out to improve model performance and reduce errors during training.

- *Data Augmentation Techniques:*
  To augment the dataset and ensure the model could generalize well across different scenarios, the following augmentation techniques were applied:

- *Horizontal and Vertical Flipping:*
  Images were flipped horizontally and vertically to introduce variability in the orientation of bottles.

- *Rotation:*
  Images were rotated randomly by 90°, 180°, or 270° to expose the model to different angles.

- *Brightness Adjustment:*
  The brightness of the images was varied between -10% and +10% to simulate different lighting conditions.

- *Zooming and Cropping:*
  Random zooming and cropping were applied to focus on different parts of the bottle, especially the defect regions.

- *Noise Addition:*
  Gaussian noise (0.1%) was added to simulate imperfections in the images, such as sensor noise or interference.

- *Gaussian Blur:*
  A 1px Gaussian blur was applied to build resilience against slight focus variations.

After these preprocessing and augmentation techniques were applied, the total dataset size grew substantially, providing a more diverse and comprehensive dataset for training the models. The augmented dataset contains **7000 images**, with approximately **1000 images per class**.

➢ *Exporting the Dataset*

Once the dataset was fully prepared, it was exported from **Roboflow** in formats compatible with the training frameworks. The most commonly used formats for exporting datasets include:

- *JSON:*
  Used for compatibility with frameworks like TensorFlow Object Detection API and Detectron.

- *TXT*:
  A format suitable for YOLO models, which is particularly well-suited for object detection tasks.

With the dataset fully prepared and exported, the next step involved training the object detection models using YOLOv8, YOLOv9, and YOLOv11 architectures to evaluate their effectiveness in detecting and classifying defects in beverage bottles.

V. MODEL BUILDING

Model building is a crucial step in the CRISP-ML(Q) methodology, and it plays a significant role in object detection tasks. For this project, we explored various model architectures to automate the detection of defects in beverage bottles. The models we considered include different versions of **YOLO** (You Only Look Once) — YOLOv8, YOLOv9, and YOLOv11 — all of which are known for their accuracy, speed, and suitability for real-time object detection tasks.

We selected **YOLO** models due to their real-time object detection capabilities, accuracy, and performance in identifying defects like cracked bottles, missing caps, and misaligned labels. YOLO models have revolutionized object detection in computer vision by offering fast and efficient predictions.

➢ *YOLO V8*

YOLOv8, the 8th version of the YOLO family, was designed to provide a balance between speed and accuracy. It includes improvements in its backbone architecture and enhancements in the detection head. YOLOv8 is known for its speed and higher accuracy in detecting defects compared to previous versions. The model uses **CSPDarknet53** as its backbone as seem in its architecture diagram [Fig 3] and incorporates self-attention mechanisms to improve detection performance in cluttered environments.

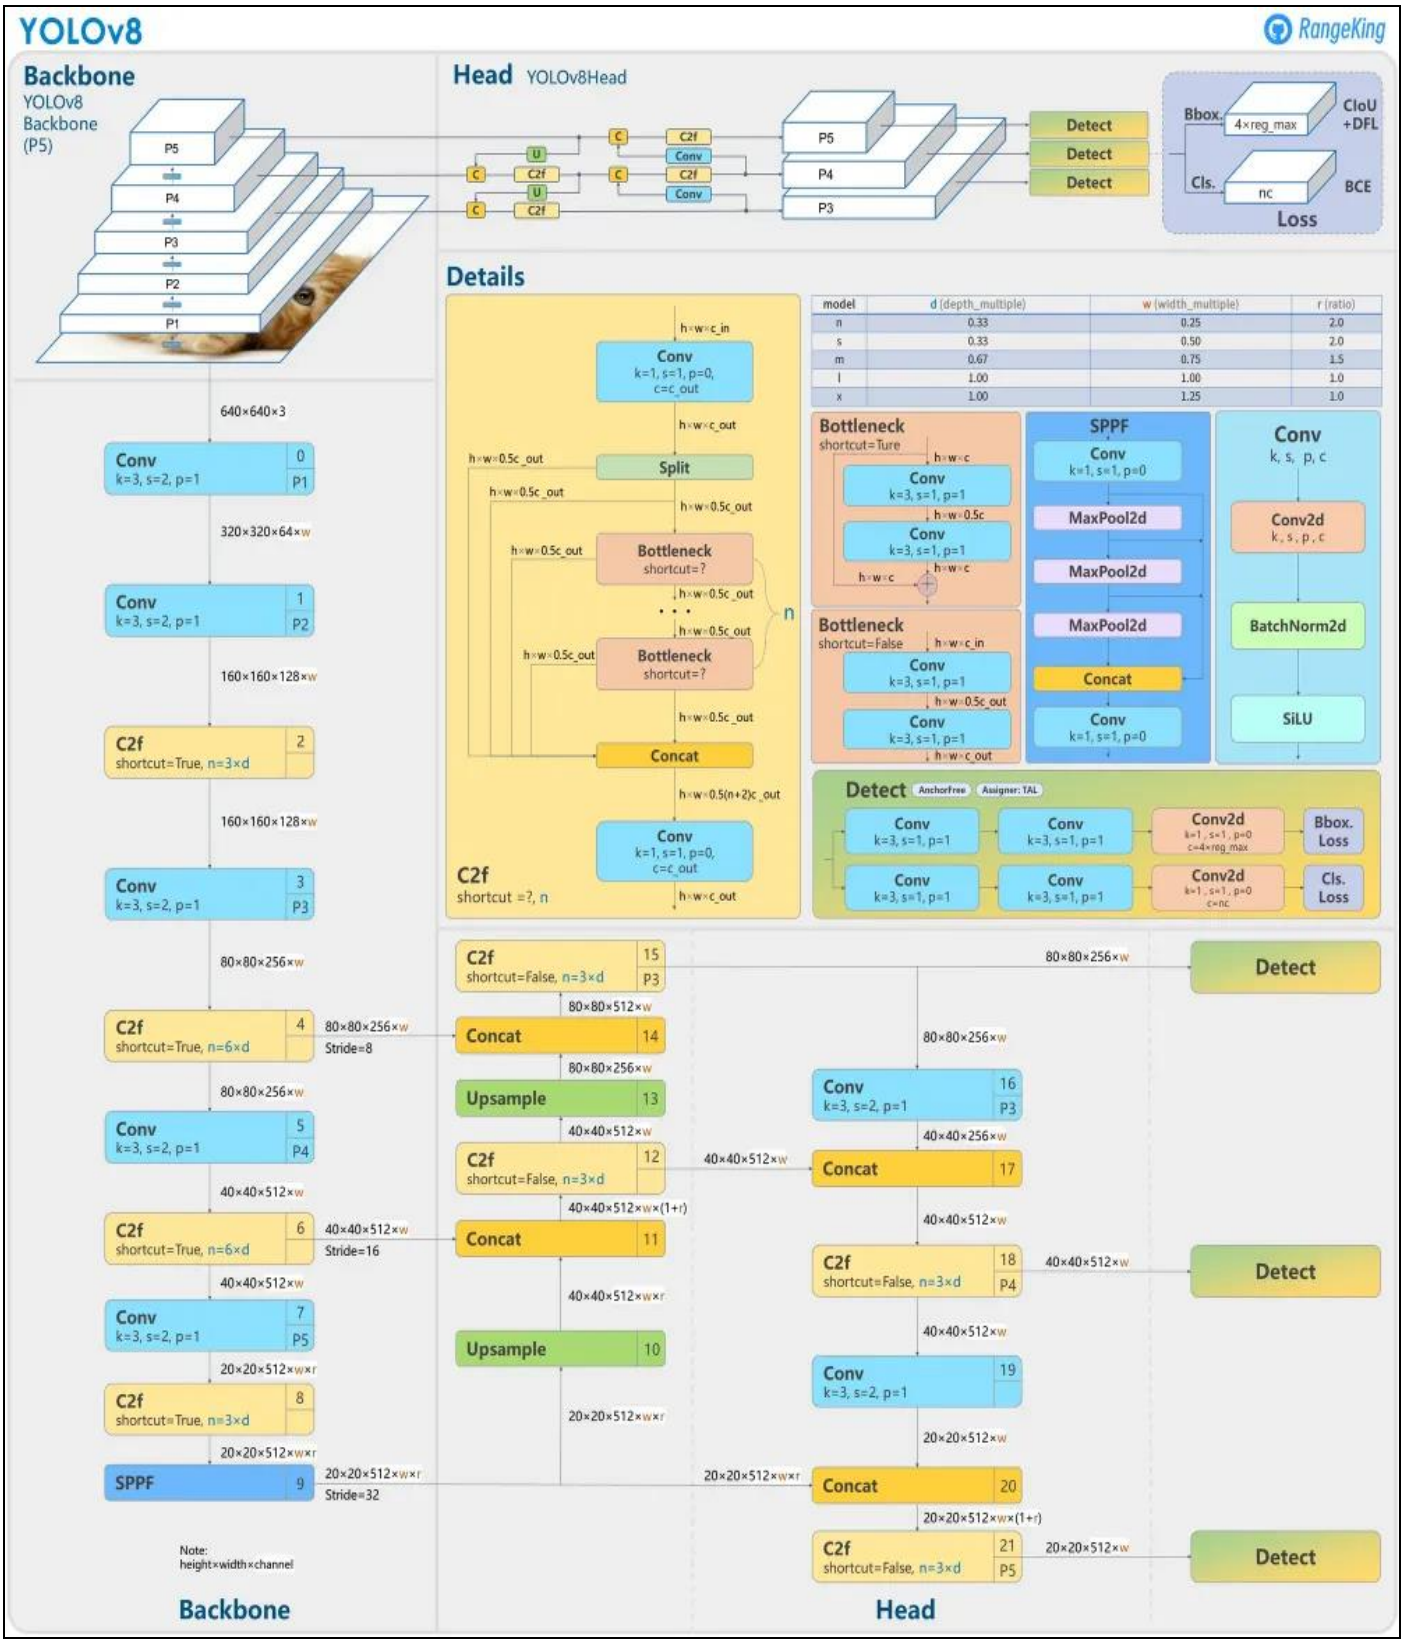

Fig 4 YOLOv8 Architecture, Visualization made by GitHub user RangeKing (Source:-YOLOv8 Architecture)

- *YOLO V9*

YOLOv9 is the latest state-of-the-art (SOTA) architecture in the YOLO family. It combines **Programmable Gradient Information (PGI)** and the **Generalized Efficient Layer Aggregation Network (GELAN)**, which significantly reduces the number of parameters while maintaining or improving accuracy. YOLOv9 is faster and more computationally efficient than previous versions, achieving up to **5-15% fewer calculations** while maintaining a significant performance boost in average precision (mAP). It is designed for real-time use, making it ideal for object detection in industrial settings.

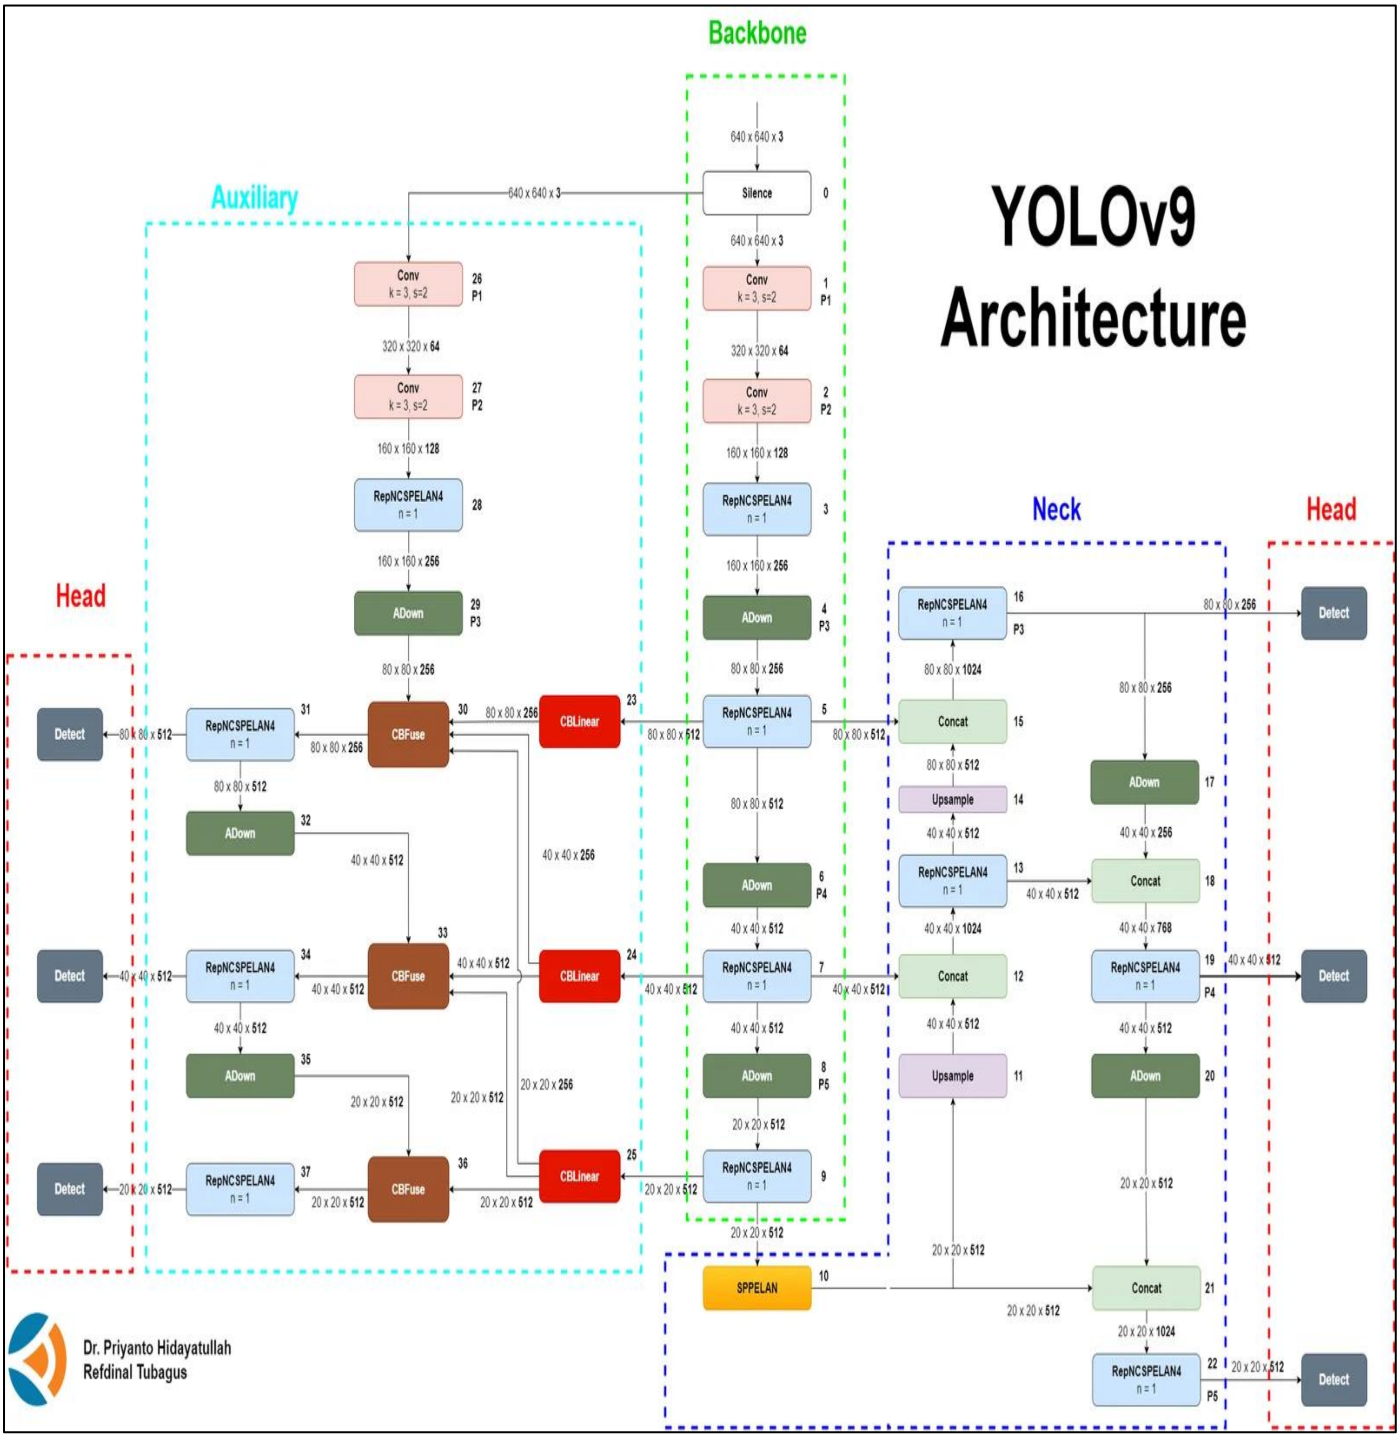

Fig 5 YOLOv9 Architecture, Visualization made by https://stunningvisionai.com author Dr. Priyanto Hidayatullah (Source:- YOLOv9 Architecture)

➢ *YOLO V11*

YOLOv11 builds upon the advancements of YOLOv8 and YOLOv9 by incorporating more advanced techniques for detecting subtle defects in object detection tasks. The model uses **multi-scale feature aggregation** and **contextual learning** to enhance its ability to detect and classify objects with higher precision. YOLOv11's architecture focuses on improving its **mean average precision (mAP)** across multiple defect classes, particularly in challenging conditions like low light or cluttered backgrounds.

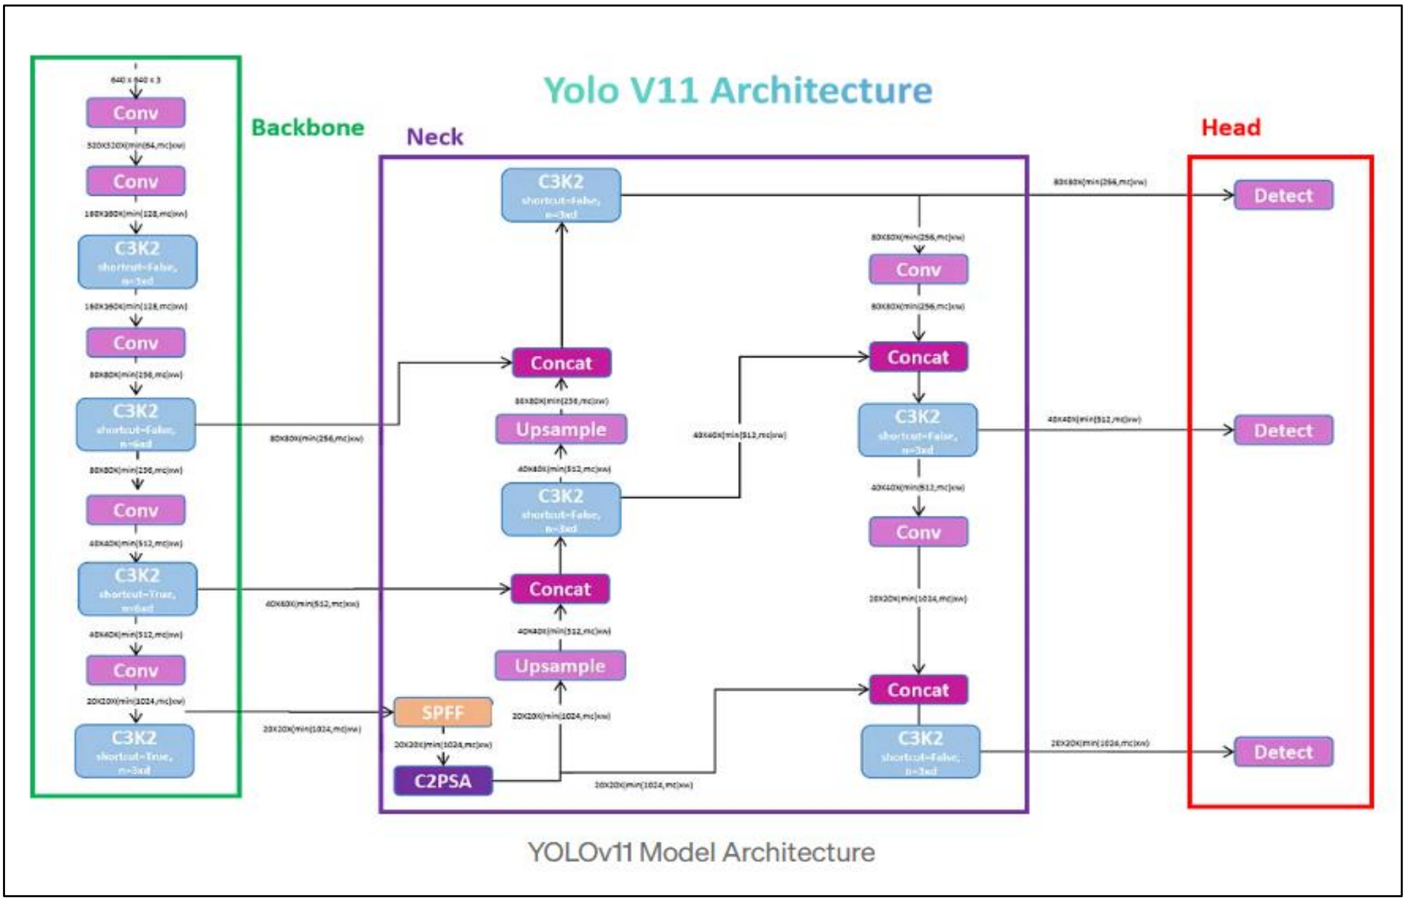

Fig 6 YOLOv11 Architecture, Visualization made by Medium user nikhil-rao-20 (Source:-YOLOv9 Architecture)

VI. HYPERPARAMETERS AND TRAINING DETAILS

In the training process, several key hyperparameters were consistent across all models to ensure a fair comparison. These hyperparameters included the learning rate, image size, batch size, optimizer, device, and number of epochs. The learning rate was set to 0.001, which is the default for most object detection models [13], to ensure stable convergence during training. The image size was standardized to 640x640 pixels, allowing the models to learn from images of a consistent size, which is particularly important for object detection tasks where spatial relationships between objects need to be preserved. The batch size was set to 16, balancing memory usage and computational efficiency. The Adam optimizer was chosen for its adaptive learning rate capabilities, which helps the model converge more efficiently. The training was performed on a CUDA-enabled GPU to accelerate the processing and minimize training time. The random seed was set to 42 to ensure reproducibility of results, and the models were trained for 30 epochs to allow sufficient training time for the models to learn the patterns in the dataset.

VII. MODEL EVALUATION

Model evaluation plays a pivotal role in assessing the performance of the trained models. To evaluate our models, we used standard metrics for object detection tasks [14], which are key indicators of how well the models can detect and classify defects in the beverage bottles.

➢ *Evaluation Metrics*

- *Precision:*
  Measures the proportion of true positive detections among all detected instances. Precision tells us how often the model correctly detects a defect without misclassifying a non-defective bottle.

- *Recall:*
  Measures the model's ability to identify all actual defects, focusing on how many true positive instances are found relative to all possible defects.

- *F1-Score:*
  A harmonic mean of precision and recall, providing a balanced measure of the model's accuracy.

- *Mean Average Precision (mAP):*
  This metric evaluates the model's overall ability to correctly identify and localize objects across different defect categories, with mAP@50 focusing on strict IoU thresholds (50%) and mAP@50-95 evaluating the model's robustness over a range of IoU thresholds.

The performance of each YOLO model variant was evaluated [15] based on mean average precision (mAP), which is the primary metric for object detection tasks. The mAP was calculated at two different thresholds: mAP@50 (evaluating the precision at an Intersection over Union (IoU) threshold of 0.50) and mAP@50-95 (which averages the precision over IoU thresholds from 0.50 to 0.95).

ISSN No:-2456-2165 https://doi.org/10.38124/ijisrt/25mar1796

- YOLO11s achieved a mAP@50 of 0.80 and a mAP@50-95 of 0.66, demonstrating solid performance, especially at the 50% IoU threshold, but with room for improvement at stricter IoU values.
- YOLOv8m performed better, with a mAP@50 of 0.88 and a mAP@50-95 of 0.75, indicating a more balanced performance across both IoU thresholds.
- YOLO11s achieved a mAP@50 of 0.85 and a mAP@50-95 of 0.75, demonstrating solid performance, especially at the 50% IoU threshold, but with room for improvement at stricter IoU values.
- YOLOv9c showed slightly lower performance, achieving mAP@50 of 0.79 and mAP@50-95 of 0.66, suggesting it was less effective than the other models, particularly at the stricter IoU threshold.
- YOLOv8s demonstrated similar performance to YOLOv8m, with a mAP@50 of 0.83 and a mAP@50-95 of 0.71, indicating strong accuracy across both evaluation metrics.
- YOLOv8m achieving a mAP@50 of 0.88 and a mAP@50-95 of 0.75, making it the most accurate and robust model in this comparison, with the highest performance at both IoU thresholds.

These results underline the importance of selecting the right model for the task at hand, with YOLOv8m being the top performer showing strong potential for real-world deployment.

VIII. MODEL PERFORMANCE COMPARISON

The following table compares the performance of the different models based on the **mean average precision (mAP)**:

Table 2 Model Performance Comparison (mAP)

Models	Epochs	map 50	map 50-95
Yolo 11s	30	0.8	0.66
Yolo V8m	30	0.88	0.75
Yolo 11m	30	0.85	0.75
Yolo v9c	30	0.79	0.66
Yolo V8s	30	0.83	0.71

Based on the results [Table 2], **YOLOv8m** achieved the highest performance, making it the model of choice for our defect detection system. It not only had the highest mAP scores but also demonstrated better generalization across different defect classes.

IX. CONCLUSION

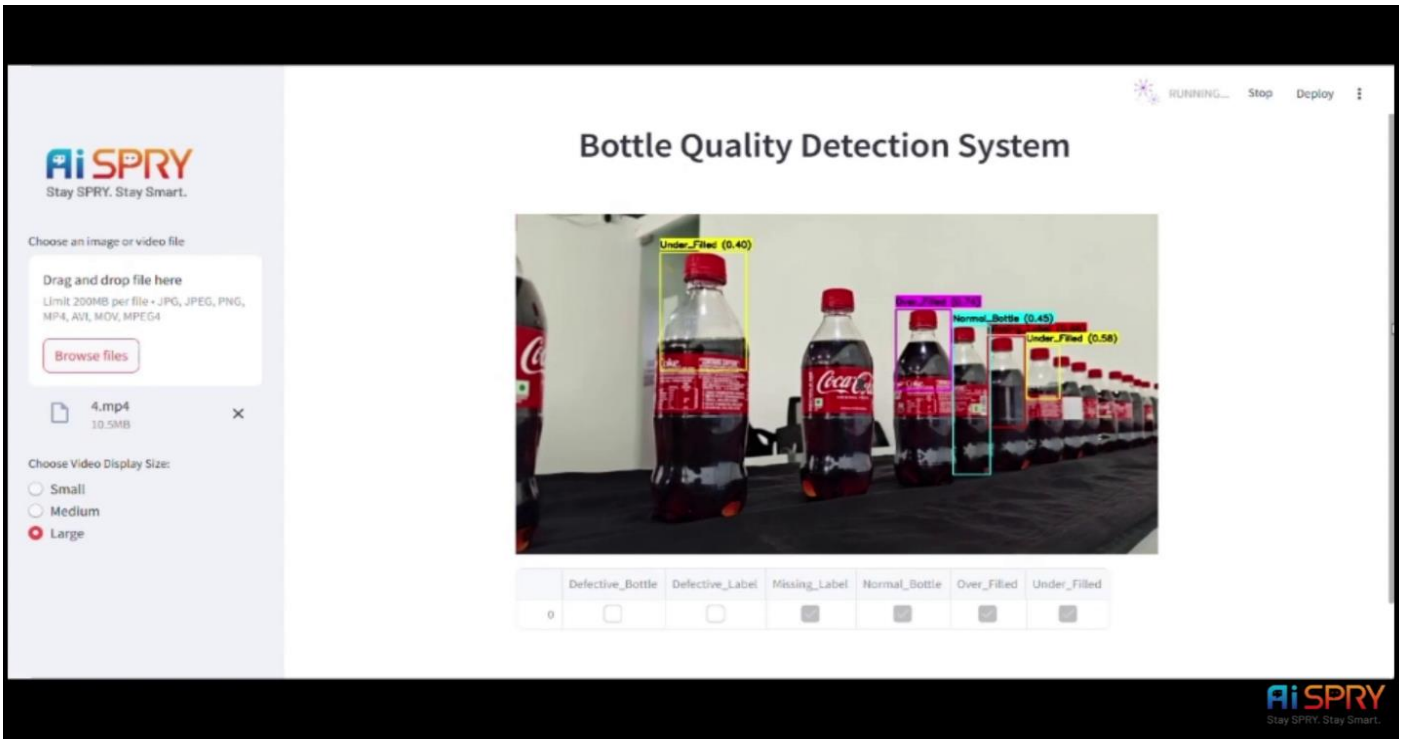

Fig 7 Defect Detection in Beverage Bottles

After conducting extensive training and evaluation, it is clear that YOLOv8m stands out as the best model for detecting defects [Fig 7] in beverage bottles within a production line. The model's superior performance in terms of both speed and accuracy makes it highly suitable for deployment in real-time industrial applications. By automating the defect detection process, we are able to significantly improve the quality control system, ensuring only high-quality products reach consumers. This research demonstrates the potential of AI-driven solutions in optimizing the beverage industry's quality assurance processes, ultimately contributing to better customer satisfaction and reduced operational costs.

FUTURE SCOPE

Future advancements for the defect detection system in beverage bottles could focus on expanding the range of detectable defects by integrating more specialized object detection models, such as segmentation models to identify even the smallest imperfections. Additionally, incorporating other sensors, such as infrared or ultrasonic sensors, could allow for detecting hidden defects that may not be visible to traditional camera systems. The system could also be adapted to monitor additional stages in the production line, such as labeling and packaging, to ensure that the quality control process covers the entire production cycle. Future work will focus on improving model generalization, ensuring robust performance across different production environments and conditions, and enhancing the system's ability to detect defects under varying lighting and camera angles. Integrating the system with IoT platforms will allow for real-time monitoring and predictive maintenance of production lines, reducing downtime and improving overall operational efficiency. Moreover, user feedback systems could be implemented to further fine-tune and personalize the defect detection capabilities.

ACKNOWLEDGMENTS

We affirm our use of the CRISP-ML(Q) and ML Workflow which are openly available on the official 360digitmg website with explicit consent from 360digitmg. We would like to express our sincere gratitude to the developers and contributors of YOLOv8, whose object detection models played a pivotal role in the success of this project. Additionally, we acknowledge the valuable support and resources provided by our industrial partners in the beverage production sector, whose collaboration made this research possible. We also thank the research team for their dedication and commitment throughout the project.

REFERENCES

[1]. Studer, S.; Bui, T.B.; Drescher, C.; Hanuschkin, A.; Winkler, L.; Peters, S.; Müller, K.-R. Towards CRISP-ML(Q): A Machine Learning Process Model with Quality Assurance Methodology. *Mach. Learn. Knowl. Extr.* 2021, *3*, 392-413. https://doi.org/10.3390/make3020020

[2]. Y. Li and Y. Zhang, "Application Research of Computer Vision Technology in Automation," *2020 International Conference on Computer Information and Big Data Applications (CIBDA)*, Guiyang, China, 2020, pp. 374-377. https://doi.org/10.1109/CIBDA50819.2020.00090.

[3]. Congcong Wang, Xiumei Wei, Xuesong Jiang, An automated defect detection method for optimizing industrial quality inspection, Engineering Applications of Artificial Intelligence, Volume 127, Part B, 2024, 107387, ISSN 0952-1976, https://doi.org/10.1016/j.engappai.2023.107387

[4]. Redmon, J., & Farhadi, A. (2018). YOLOv3: An Incremental Improvement. arXiv preprint arXiv:1804.02767. https://doi.org/10.48550/arXiv.1804.02767

[5]. Bochkovskiy, A., Wang, C. Y., & Liao, H. Y. M. (2020). YOLOv4: Optimal Speed and Accuracy of Object Detection. arXiv preprint arXiv:2004.10934. https://doi.org/10.48550/arXiv.2004.10934

[6]. Wang, C. Y., Bochkovskiy, A., & Liao, H. Y. M. (2023). YOLOv7: Trainable bag-of-freebies sets new state-of-the-art for real-time object detectors. arXiv preprint arXiv:2207.02696. https://doi.org/10.48550/arXiv.2207.02696

[7]. Lin, T. Y., Maire, M., Belongie, S., Hays, J., Perona, P., Ramanan, D., ... & Zitnick, C. L. (2014). Microsoft COCO: Common Objects in Context. In European conference on computer vision (pp. 740-755). Springer, Cham. https://doi.org/10.48550/arXiv.1405.0312

[8]. He, K., Zhang, X., Ren, S., & Sun, J. (2016). Deep Residual Learning for Image Recognition. In Proceedings of the IEEE conference on computer vision and pattern recognition (pp. 770-778). https://doi.org/10.48550/arXiv.1512.03385

[9]. 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. https://doi.org/10.48550/arXiv.1704.04861

[10]. Tan, M., & Le, Q. V. (2019). EfficientNet: Rethinking Model Scaling for Convolutional Neural Networks. In International Conference on Machine Learning (pp. 6105-6114). PMLR. https://doi.org/10.48550/arXiv.1905.11946

[11]. Shorten, C., Khoshgoftaar, T.M. A survey on Image Data Augmentation for Deep Learning. J Big Data 6, 60 (2019). https://doi.org/10.1186/s40537-019-0197-0

[12]. Everingham, M., Van Gool, L., Williams, C.K.I. et al. The Pascal Visual Object Classes (VOC) Challenge. Int J Comput Vis 88, 303–338 (2010). https://doi.org/10.1007/s11263-009-0275-4

[13]. Ren, S., He, K., Girshick, R., & Sun, J. (2015). Faster R-CNN: Towards Real-Time Object Detection with Region Proposal Networks. In Advances in neural information processing systems (pp. 91-99). https://doi.org/10.48550/arXiv.1506.01497

[14]. Girshick, R., Donahue, J., Darrell, T., & Malik, J. (2014). Rich Feature Hierarchies for Accurate Object Detection and Semantic Segmentation. In Proceedings of the IEEE conference on computer vision and pattern recognition (pp. 580-587). https://doi.org/10.48550/arXiv.1311.2524

[15]. Ali, U., Maizatul Akmar Ismail, Riyaz Ahamed Ariyaluran Habeeb, & Syed Roshaan Ali Shah. (2024). Performance Evaluation of YOLO Models in Plant Disease Detection. *Journal of Informatics and Web Engineering*, *3*(2), 199–211. https://doi.org/10.33093/jiwe.2024.3.2.15