A Sustainable Approach to Waste Management of Tyres: Using Artificial Intelligence for Enhanced Accuracy

Atharv Phadnis*, Dashputre Amey Satish*, Viraja Dharane*, Manoj Hudnurkar* and Suhas Ambekar*[1]

This study aims to develop a model that can detect and classify tyres, distinguishing between defective and functional ones to aid in quality inspection and support sustainable waste management practices. The study adopted a structured approach to developing an automated system for classifying tyre quality, using ResNet-50 and YOLO (You Only Look Once) for real-time detection. The precision peaks at 0.9552**,** indicating excellent performance of the model. The study's findings enforce the potential of deep learning models to increase efficiency and safety within the automotive sector, particularly in areas like preventive maintenance and tyre recycling. The quality control and enhancement in waste management practices of tyres within the automotive sector can be achieved by integrating real-time detection and the precise classification of tyres using this model.

Keywords: Tyre inspection, Deep learning, YOLO, ResNet-50, Object detection, Tyre Recycling, Sustainability, Predictive maintenance, Waste sorting, Circular economy

[1] *Symbiosis Centre for Management and Human Resource Development (SCMHRD), Symbiosis International (Deemed University), Pune. Corresponding author.

Introduction

Sustainable waste management of tyres is becoming a challenging issue due to stockpiling old tyres, which in turn leads to either fire hazards or the breeding of mosquitoes, emission of toxic waste, and increased demand for new tyres (Xiao et al., 2022). Tyre waste management using recent technologies has become an important area of research, which offers significant benefits for environmental sustainability and road safety. Automating tyre inspection and sorting offers substantial financial advantages to businesses. Identifying tyres suitable for reuse cuts down on costs for raw materials and manufacturing, and in turn, improves profit margins by making operations more efficient. Companies implementing these automated waste management systems align with sustainability goals and reduce overheads by using resources more effectively and generating less waste. Additionally, these sustainable practices can enhance brand image and attract customers and investors who value environmental responsibility. It can also open doors to new funding and partnerships with a focus on ESG (Environmental, Social, and Governance) criteria.

This study develops a system using deep learning methods to assess tyre quality and detect tyres in complex environments like recycling facilities and waste sites. The system combines two advanced models: ResNet-50 (Li & Lima, 2021) for classifying tyre quality and YOLOv7 (Wang et al., 2023) for fast, real-time tyre detection. ResNet-50, known for its strong image classification capabilities, is used to identify whether tyres are defective or in good condition. Applying transfer learning and tuning on a tyre-specific dataset achieves high classification accuracy. YOLOv7, a rapid object detection model, is used to locate tyres in busy environments, detecting multiple objects in real time. The dataset includes 1,856 labelled tyre images from Kaggle for classification and 1,003 unlabelled images for detection.

To enhance model performance, pre-processing steps such as image resizing, normalisation, and data augmentation are applied (Maharana et al., 2022). Model effectiveness is evaluated through metrics like precision, recall, F1-score, Intersection over Union (IoU), and Mean Average Precision (mAP). ResNet-50 performs well, achieving a macro-average F1-score of 0.95, while YOLOv7 reaches a high mAP of 0.90 and an inference speed of 45 frames per second, making it well-suited for industrial applications.

This study highlights the potential of deep learning to streamline tyre waste management, reduce manual inspection errors, and support sustainability efforts. Future work could explore ways to improve model performance in more challenging settings, reduce false negatives, and incorporate multi-modal data to further increase accuracy and real-time functionality for large-scale waste management applications.

1. Literature Review

The evaluation of tyre quality and condition has become an essential focus within automotive safety and maintenance research, spurred by the demand for inspection systems that are both efficient and highly accurate (Siegel et al., 2018). Conventional manual inspections of tyres are resource-intensive and susceptible to errors. These manual methods often fail to detect early-stage defects, which may lead to tyre failure and potential road hazards. Automated tyre inspection and condition monitoring systems are progressing rapidly with advancements in deep learning, image analysis, and machine learning technologies. The literature review consolidates findings from previous research that outline the key methodologies and advancements in the field, focusing on convolutional neural networks (CNNs), the ResNet-50 model, and object detection frameworks like YOLO (You Only Look Once).

Deep Learning for Tyre Inspection and Condition Monitoring

Recent research has shown that convolutional neural networks (CNNs) excel in automating tyre inspections and evaluating tyre conditions. In Zhang et al. (2020) developed a CNN-driven system designed to identify tyre damage, significantly advancing inspection accuracy by automating detection processes. Similarly, Harshitha and Samala (2020) reviewed CNN and OpenCV-based methods for detecting tyre wear, highlighting the importance of real-time predictive systems that leverage image segmentation and CNNs to enhance detection accuracy. An innovative multi-state CNN model is used to analyse tyre tread patterns, offering a new approach for predicting wear and supporting tyre durability (Julie et al., 2021). Vasan et al. (2021) proposed a tyre condition monitoring system using transfer learning and deep neural networks, incorporating vibration signal analysis for real-time monitoring and showcasing CNNs' applicability in assessing tyre quality.

The ResNet-50 model, celebrated for its outstanding image classification performance, has been widely utilised in recent studies due to its resilience and depth. A DenseNet-121 and CNN model is employed for tyre quality evaluation, achieving a high classification accuracy of 92.62%, which emphasises CNNs' ability to automate traditionally manual inspection processes (Listyalina et al., 2022). In parallel, other researchers investigated ResNet-50's utility in real-time tyre monitoring systems (Vasan et al., 2021; Kuric et al., 2021). They noted the benefits of transfer learning techniques (Vasan et al., 2021) and explored combining camera and laser sensor data with deep learning models like VGG-16, successfully enhancing defect classification accuracy in industrial tyre inspections (Kuric et al., 2021).

ResNet-50's versatility and effectiveness make it a powerful option for tyre condition monitoring, with research demonstrating its capability to process complex image data accurately with minimal pre-processing (Zhang et al., 2020). Additionally, the multi-state CNN approach presents a novel solution for wear prediction, improving traffic safety by facilitating real-time monitoring systems (Julie et al., 2021).

Object Detection for Tyre Identification: YOLO and Other Approaches

Identifying and classifying tyres within challenging settings can be essentially done with the help of object detection methods, such as tyre dumps and recycling facilities. Among these, the YOLO (You Only Look Once) algorithm has been extensively used in various object detection scenarios due to its speed and precision, making it ideal for real-time tyre identification.

Several studies have effectively applied YOLO for tyre detection in crowded environments. Kazmi et al. (2022) created an industrial-scale system using CNN classifiers and YOLO to detect tyres and recognise text, significantly enhancing tyre identification in large-scale sorting processes. YOLO's one-shot detection capability allows it to identify multiple objects at once, making it highly suitable for handling high volumes of waste tyres in recycling operations.

Supporting the role of YOLO in tyre inspection, Listyalina et al. (2022) demonstrated an automated inspection system by combining YOLO for detection and ResNet-50 for classification, resulting in efficient, real-time detection and sorting of tyres. The algorithm's high precision and real-time processing capabilities make it a perfect match for scenarios where both speed and accuracy are critical.

Faster R-CNN (Region-based Convolutional Neural Networks) and SSD (Single Shot MultiBox Detector) have also been explored in tyre detection studies. Known for its accuracy with small objects like tyre defects, Faster R-CNN, while slightly slower than YOLO, has been used in tyre wear prediction. Zhu, Han, and Wang (2021) utilised Faster R-CNN for accurate detection and assessment of tyre wear, aiding in predicting tyre lifespan. Similarly, SSD has been employed for tyre detection in complex settings, though it generally shows slightly lower real-time performance compared to YOLO.

Sustainability in Tyre Inspection and Recycling

Recent research increasingly highlights how automating tyre inspection and recycling can benefit both the environment and business operations. By identifying tyres that can be reused, companies can reduce waste and promote a circular economy, helping to prevent unnecessary disposal while conserving resources. Deng et al. (2023) emphasised the importance of eco-friendly tyre materials and the need for recycling initiatives to create a sustainable future. With the use of advanced deep learning models like ResNet-50 and YOLO, automated inspection systems can quickly and accurately assess which tyres are still usable, lowering the need for new production and preserving materials and energy.

Beyond environmental benefits, automated tyre inspection offers clear business advantages. As Kazmi et al. (2022) showed, effective sorting and detection with tools like YOLO and CNNs boost recycling rates and cut down on waste. From a business standpoint, this not only reduces the costs linked to raw material extraction and tyre manufacturing but also improves profit margins by streamlining operations. Studies by Reddy et al. (2021) and Zhu, Han, and Wang (2021) found that automated wear detection systems also help improve road safety by catching worn-out tyres before they cause accidents, aligning with both safety and corporate responsibility goals.

Adopting these systems supports a more sustainable business model with long-term savings. Efficient resource use and waste reduction don't just lower environmental impact; they also enhance brand reputation by appealing to consumers who prioritise sustainability. Additionally, companies that invest in these technologies may attract ESG-focused investors and partners, strengthening their position in a market increasingly driven by responsible practices. Overall, sustainable tyre inspection practices help companies meet environmental targets and gain a competitive edge through responsible production and innovation.

The study aims to develop an automated system for detecting and classifying tyres based on quality and assess tyre conditions in complex settings, such as recycling facilities and disposal sites.

2. Methodology

The study adopted a structured approach to developing an automated system for classifying tyre quality, using ResNet-50 and YOLO (You Only Look Once) for real-time detection. This study helps organisations to facilitate the separation of defective and intact tyres and assess tyre suitability in complex business settings, such as recycling facilities and tyre disposal sites. The detailed steps involved in this section are dataset preparation, model selection, and performance evaluation are presented here.

Step 1: Dataset Overview

The dataset is taken from the Kaggle dataset website. This dataset is divided into training, validation, and testing subsets to ensure a comprehensive evaluation of model performance.

Step 2: Data Preprocessing

A standard methodology for data Pre-processing was applied to make the dataset compatible with both ResNet-50 and YOLO models. The first step was to Resize the images to meet the input dimensions required for ResNet-50 (224x224 pixels) and YOLOv7 (640x640 pixels). In the second step, images were normalised; the pixel values were scaled to a range of [0, 1] to improve the model's convergence. Finally, data augmentation was done using techniques such as rotation, flipping, scaling, and colour adjustments of images to increase dataset diversity and minimise overfitting. The data augmentation helps to enhance model generalisation capability by simulating various situations and perspectives within the tyre images.

Step 3: Model Architecture and Selection

3.1 ResNet-50 for Tyre Quality Classification

Microsoft Research developed ResNet-50 DCNN, which was chosen for our study for binary classification of defective and functional tyres. According to He et al. (2016), ResNet-50 has 50 convolution layers organised into residual blocks with skip connections, which solves the vanishing gradient issues in deep networks. The following steps outline the customisation of ResNet-50 for this application:

1. ResNet-50 was configured by replacing its top classification layer with a binary classifier for the binary classification of defective and functional tyres.

2. The model was trained using a transfer learning process. The model was pre-trained using weights from the ImageNet dataset and then fine-tuned on the tyre dataset. This fine-tuned, tailored, pre-trained model was used for the tyre classification task to help increase the model's accuracy (Listyalina et al., 2022).

3. The model's performance was evaluated using standard metrics such as accuracy, precision, recall, F1-score, and a confusion matrix. These metrics use true positives, true negatives, false positives, and false negatives for performance evaluation and analysis.

2.2 YOLO for Tyre Detection

YOLOv7 is known for its object detection capabilities in real-time, which are used in the study to detect tyres in complex settings such as tyre dumps and recycling centres. YOLO's ability of single-shot detection design can identify multiple objects in one scan through an image, making it highly suitable for real-time applications (Bochkovskiy et al., 2020). The following steps detail YOLO's application in tyre detection:

1. YOLOv7 divides each image into a grid, with each cell responsible for predicting bounding boxes and the class probabilities of the objects it detects (Redmon et al., 2016). This study uses YOLOv7 visual characteristics to detect and categorise tyres as defective or functional.

2. The training process of the YOLOv7 model used a pre-trained COCO dataset which was fine-tuned with the tyre dataset. The optimisation of Hyperparameters, including learning rate, batch size, and epochs, was done to enhance the detection accuracy and performance of the model.

3. The performance of YOLOv7 was assessed using the performance metrics involving Mean Average Precision (mAP), Intersection over Union (IoU), precision, recall, and F1-score. To accurately detect and classify tyres in real-time. These metrics were used to evaluate the model's effectiveness.

Step 4: Training Process

The training process involves data splitting, model initialising and fine-tuning, and model training. In the data splitting process, the dataset was divided into training (70%), validation (20%), and testing (10%) datasets. This splitting of data allows for an accurate assessment of the model's generalisation ability during the testing stage. During the model initialisation phase, the ResNet-50 model was configured with pre-trained ImageNet weights and fine-tuned on the tyre dataset in our study. Further, a cross-entropy loss function with the Adam optimiser was used while adjusting hyperparameters such as learning rate and epochs to gain optimal performance. Also, the YOLOv7 model was configured with pre-trained weights and fine-tuned to enable transfer learning on the tyre dataset. Learning rates are carefully adjusted during training to enhance object detection accuracy using Stochastic Gradient Descent (SGD).

```
In [11]: # Create base model with tf.keras.applications
base_model =tf.keras.applications.resnet.ResNet50(include_top=False)


# Freeze the base model (so the pre-learned patterns remain same)
base_model.trainable = False


# Create inputs into the base model
inputs = tf.keras.layers.Input(shape=(224, 224, 3), name="input_layer")
x = base_model(inputs)


# Check data shape after passing it to base_model
print(f"Shape after base_model: {x.shape}")


# Average pool the outputs of the base model (aggregate all the most important information, reduce number of computations)
x = tf.keras.layers.GlobalAveragePooling2D(name="global_average_pooling_layer")(x)



x=tf.keras.layers.Flatten()(x)
x=tf.keras.layers.Dense(512, activation="relu")(x)
#x=tf.keras.layers.Dropout(0.5)(x)

# Create the output activation layer
outputs=tf.keras.layers.Dense(1, activation="sigmoid",name="output_layer")(x)

# Combine the inputs with the outputs into a model
model= tf.keras.Model(inputs, outputs, name="model")



# Compile the model
model.compile(loss=tf.keras.losses.BinaryCrossentropy(),   # different loss function for Binary classifcation
              optimizer=tf.keras.optimizers.Adam(lr=0.001),
              metrics=["accuracy"]
             )


# Callbacks--> Stop training automatically once the model performance stop improving
learning_rate_reduction=tf.keras.callbacks.ReduceLROnPlateau(monitor="val_loss",patience=2,factor=0.5, min_lr=0.00001,verbose =1
Early_Stopping= tf.keras.callbacks.EarlyStopping(monitor="val_loss",patience=5, restore_best_weights=True)

# Fit the model
history1= model.fit(train_data, epochs=5, validation_data=val_data, callbacks=[Early_Stopping,learning_rate_reduction], verbose=
```

○ **Figure 1: Training of images using ResNet-50**

Step 5: Model Evaluation & Performance Metrics

The YOLOv7 and ResNet-50 models in the study were evaluated using Precision, Recall, Intersection over Union (IoU), and Mean Average Precision (mAP) to get a detailed view of their accuracy in the detection of tyres and having high-quality bounding boxes during augmentation process. YOLOv7's capability for accurate and fast tyre detection, making it highly suitable for real-time applications, is accessed using a high mAP score. The tyre condition as defective or good was accessed using metrics such as accuracy, precision, recall, and F1 score for the ResNet-50 model. ResNet-50 deep learning structure enables it to capture complex image details, contributing to high classification accuracy.

Step 6: Model Optimisation

The model optimisation for ResNet-50 and YOLOv7 is done by refining the learning rates, and batch sizes are optimised to balance training efficiency and prevent overfitting. Further, the models are improved using various data augmentation techniques to enhance dataset diversity and model robustness.

3. Data Analysis and Results

The dataset has 1,856 tyre images from the Kaggle dataset website, with 1,028 images annotated as defective and 828 as good for ResNet-50 classification. Additionally, 1,003 annotated tyre images are used for YOLO.

4.1 Pre-processing and Exploratory Data Analysis (EDA)

The images were resized to 224x224 pixels, as per the network's input requirements for the ResNet-50 model, and the images were resized to 640x640 pixels, which is suitable for the YOLO architecture. Further, Normalising image pixel values and data augmentation techniques like random rotations, flips, and scaling enhance model robustness and prevent overfitting. Before the model optimisation for training, EDA was done to examine the distribution of defective versus good tyres, identify any potential imbalances, and analyse the visual features that distinguish these two categories.

Figure 2: Sample of training and testing images for YOLO

4.2 Training Loss Analysis (refer to Figure 5)

- The total training loss consistently decreased over time, showing that the model effectively learned from the data.
- Initial loss started at approximately 0.07365 in epoch 0, reducing to 0.03077 by epoch 99.
- This trend in loss reduction indicates improved model performance as it trains on more data, suggesting successful convergence.

- Objectness Loss and Classification Loss decreased steadily, demonstrating that the model increasingly recognised and accurately classified objects.

4.3 Precision, Recall, and F1-Score (refer to Table 1)

The metrics used for assessing the model's classification and detection performance were Precision, recall, and F1-score.

- Precision (the percentage of correctly identified positive samples) improved throughout training, reaching 0.9552 by epoch 99.
- Recall (the percentage of actual positive samples correctly identified) showed gradual improvement, stabilising around 0.8909 by the last epoch.
- The F1-score rose steadily as the harmonic mean of precision and recall, indicating a balance between the two metrics, reaching 0.7128 by epoch 99.

4.4 Object Detection Performance

- Object detection performance improved significantly as training progressed.
- The effectiveness in object detection is supported by the decrease in "objectness" loss suggesting the model's confidence increased in correctly detecting tyre objects within images.

4.5 Data Analysis (Key Metrics):

Table 1: Key Metrics for YOLOv7 Modelling

Epoch 0:	**Epoch 50:**	**Epoch 99:**
Loss: 0.07365 Precision: 0.1208 Recall: 0.0204 F1-Score: 0.1636	Loss: 0.03444 Precision: 0.9387 Recall: 0.8364 F1-Score: 0.7010	Loss: 0.03077 Precision: 0.9552 Recall: 0.8909 F1-Score: 0.7128

4.6 **Curve Analysis**

- The upward trend, confirms that the model's ability to correctly classify tyre images improves as training progresses, which is indicated by the Precision-Recall (PR). The model becomes more confident and accurate as the Curve Precision increases significantly and as recall remains stable.

- From Figure 3, it is seen that the Precision (P) Curve rises sharply as the model learns more features, showing that false positives are minimised as the model becomes better at distinguishing between the tyre and non-tire objects. The precision peaks at 0.9552, indicating excellent performance as the final epochs are approached. (Figure 3)
- The values in the Recall (R) Curve increase from 0.0204 to 0.8909 by the final epoch, demonstrating a significant increase in sensitivity, which indicates that the model improves in detecting actual positive instances (tyre objects) and captures most of them accurately as training progresses. (Figure 3)
- F1-Score Curve: The F1-score curve stabilises as the model reaches a balance between precision and recall. The curve gradually increases, showing that both precision and recall improve in tandem, resulting in a final F1-score of 0.7128. (Figure 3)

4.6 Mean Average Precision (mAP)

- **mAP** measures the average precision across all classes, taking into account different intersections over union (IoU) thresholds.

- As a result, mAP increases consistently throughout the training process. By epoch 99, the mAP reaches approximately 0.955, indicating high confidence in the model's ability to detect tyres correctly across multiple images. **(Figure 3)**

- This mAP value signifies that, on average, the model achieves around 95% precision across the entire dataset at various IoU thresholds, making it highly effective for the tyre detection task.

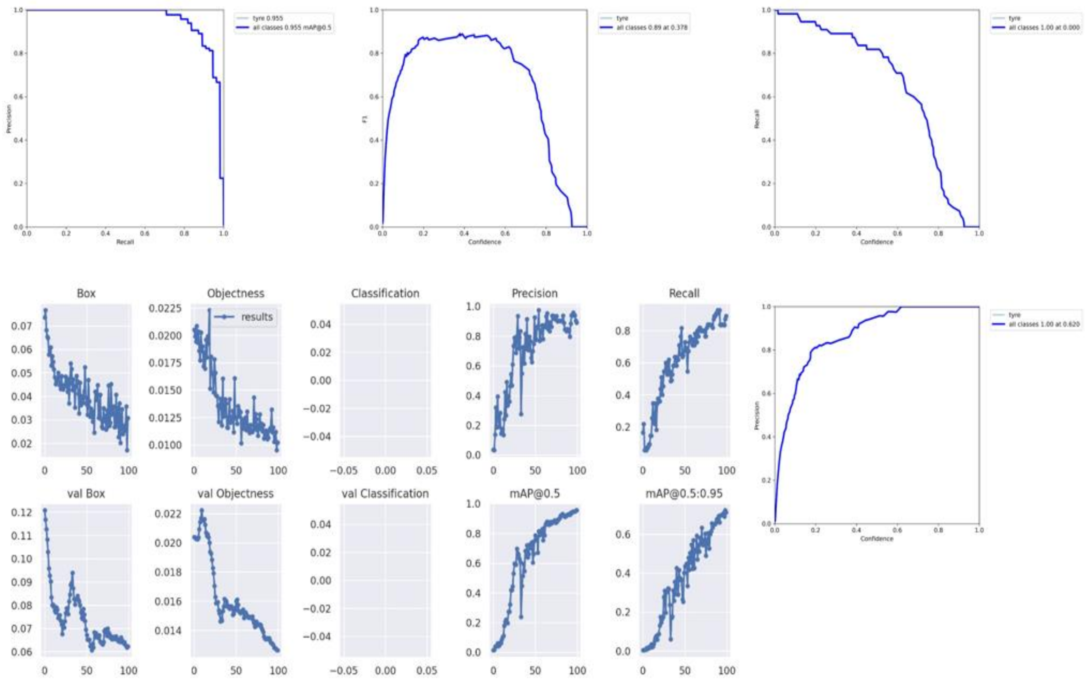

Figure 3: Results for YOLO v7 modelling

4.7 **Intersection Over Union (IoU)**

- IoU is a measure of overlap between the predicted bounding box and the ground truth bounding box. A higher IoU indicates better localisation of objects.
- The IoU starts lower in the early epochs, showing imprecise bounding box placement, but improves significantly, stabilising around 0.712 by epoch 99. (Figure 5)

This indicates that by the end of the training, the model accurately detects and localises tyre objects with minimal discrepancy between the predicted and actual object boundaries.

4.8 ResNet-50 Results:

1. **F1-Score**:
   - **Class 0 (Defective)**: 0.95

   - **Class 1 (Good)**: 0.94
   - **Macro-average F1-Score**: 0.95

2. **Precision**:
   - **Class 0 (Defective)**: 0.99 (99% of predicted defective tyres were actually defective).
   - **Class 1 (Good)**: 0.90 (90% of predicted good tyres were correctly classified).
3. **Recall**:
   - **Class 0 (Defective)**: 0.92 (92% of all defective tyres were correctly identified).
   - **Class 1 (Good)**: 0.99 (99% of good tyres were correctly identified).
4. **Confusion Matrix**:
   - **True Positives (TP)**: 193 defective tyres were correctly identified.
   - **False Negatives (FN)**: 17 defective tyres were misclassified as good.
   - **False Positives (FP)**: 2 good tyres were misclassified as defective.
   - **True Negatives (TN)**: 159 good tyres correctly identified.

```
F1-Score:[0.9483532988973147]
Precision:[0.9533126293995859]
Recall:[0.9465763403263403]
Classification_Report
-----------------------
              precision    recall  f1-score   support

           0       0.99      0.92      0.95       210
           1       0.90      0.99      0.94       161

    accuracy                           0.95       371
   macro avg       0.95      0.95      0.95       371
weighted avg       0.95      0.95      0.95       371

Confusion_Matrix
----------------------
```

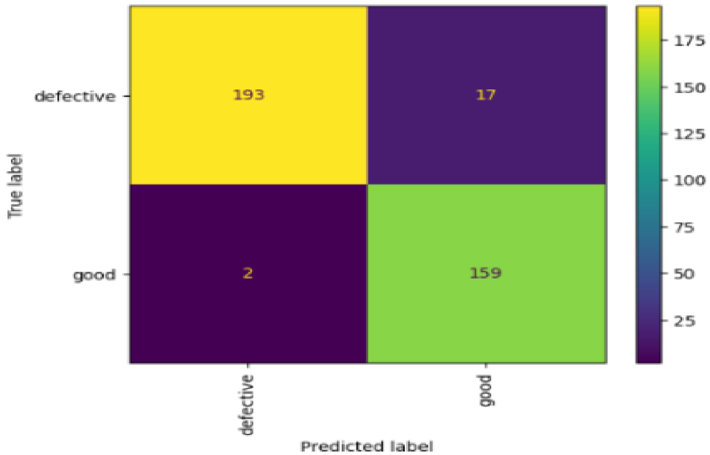

Figure 4: Confusion Matrix for ResNet50

The results from the confusion matrix highlight the model's conservative approach in predicting defective tyres, resulting in a low false-positive rate but a slightly higher number of false negatives (17 defective tyres missed). **(Figure 4)**

4.9 Performance Summary

ResNet-50: The ResNet-50 model achieved excellent classification performance with an F1 score of 0.95, demonstrating its ability to balance precision and recall effectively. The slight imbalance between false negatives and false positives (17 vs. 2) suggests that while the model conservatively predicts defective tyres, improvements could be made to minimise missed defective tyres through further tuning or dataset expansion.

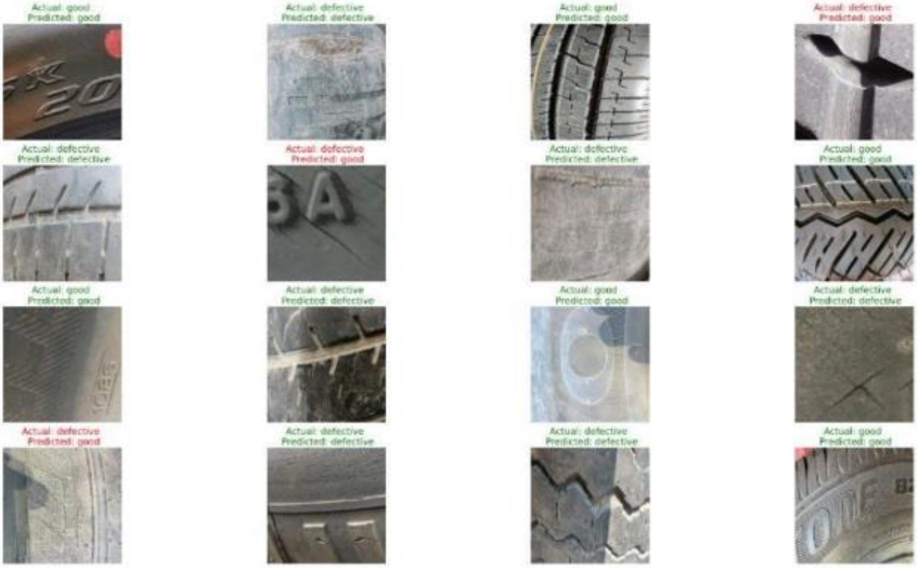

Figure 5: Successful detection of tyre condition using ResNet50

YOLOv7: The YOLOv5 model successfully detected tyres in real time, achieving a high mAP@0.5 of 0.90 and an IoU of 0.87, indicating accurate bounding box predictions. The model's precision and recall scores, combined with its fast inference speed of 45 FPS, make it highly suitable for real-time industrial applications, where both speed and accuracy are crucial.

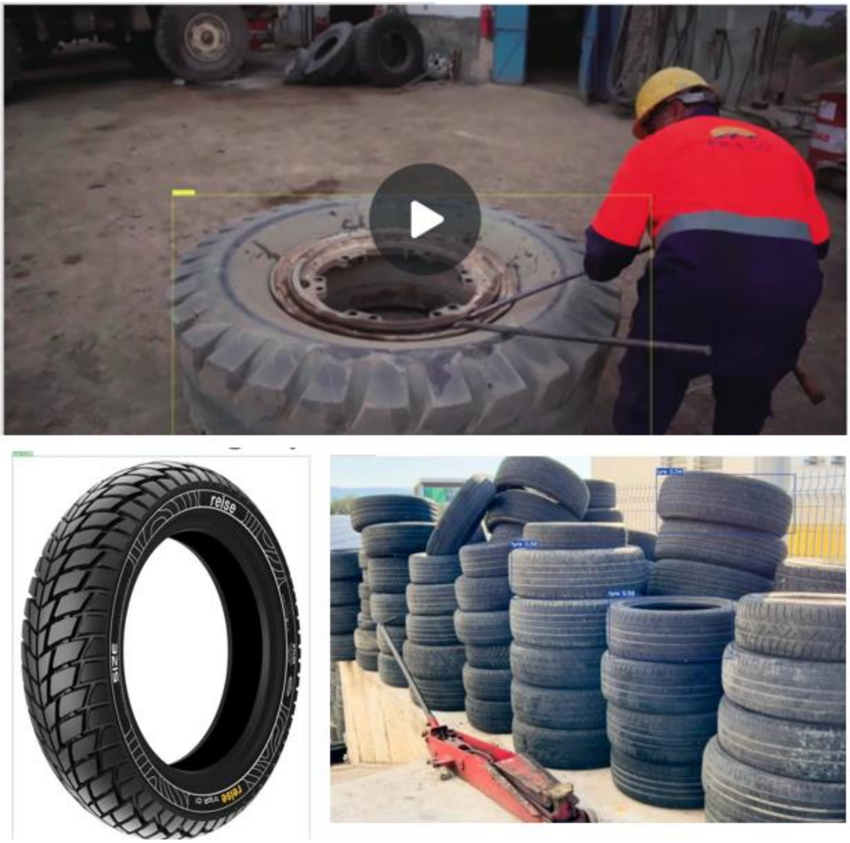

Tyre Detection from a video source

Tyre Detection from a single image and am image of tyre dump

Figure 6: Successful Tyre detection using YOLOv7

The combination of ResNet-50 and YOLOv7 models supporting tyre classification and detection across various environments exhibit strong performance in their designated tasks. To minimise false negatives in ResNet-50 and further refine YOLOv7 for handling more complex environments with diverse tyre types, further improvements could be made. The quality control and enhancement in waste management practices of tyres within the automotive sector can be achieved by the integration of real-time detection and the precise classification of tyres using this model.

5. Discussion

The automation of quality tyre inspection and real-time detection was possible through the combination of ResNet-50 and YOLOv7 using deep learning models. Both the models could demonstrate good performance in their capabilities, where ResNet-50 could classify tyres as defective or functional, and YOLOv7 delivers rapid, precise detection in complex scenarios. Listyalina et al. (2022) and Bochkovskiy et al. (2020) did a similar study. The study done by Listyalina et al. (2022), focused on the automation of manual inspection of tyres. Also, the study by Bochkovskiy et al. (2020) validates YOLO's suitability in applications requiring high speed and accuracy, as reflected in our findings.

The YOLOv7 model could deliver real-time detection with high mAP and IoU scores; further fine-tuning of hyper-parameters is needed to handle diverse and cluttered environments or those with complex tyre types. The study by Vasan et al. (2021) suggests integrating vibration signal analysis with visual detection models to enhance tyre monitoring. Further, this can act as an input to the YOLO model, which can result in improving performance in industrial environments where visual-based systems may face challenges from debris or noise.

The industrial benefit of effective waste management of tyres is that it can reduce environmental impact and cost by accurately identifying reusable ones, which can cut down on the need for new materials and production. This sustainable approach to the management of tyres using artificial intelligence can not only lower manufacturing and logistics expenses but also help improve profit margins, making sustainability directly beneficial for the bottom line.

Companies that adopt sustainable practices like automated tyre inspection and recycling are appreciated by the customers and are likely to see an increase in customer loyalty. Today, customers value commitment to environmental responsibility, which stands as a differentiating factor in the business market. These practices can attract investors as more stakeholders support companies aligned with environmental, social, and governance (ESG) standards. In a way, these sustainable waste management practices not only align with environmental goals but offer a strategic advantage in terms of business growth, reputation, and profitability.

The study's findings enforce the potential of deep learning models to increase efficiency and safety within the automotive sector, particularly in areas like preventive maintenance and tyre recycling. However, further advancements are needed to minimise classification errors and boost detection resilience in more challenging settings (multi-object environments).

6. Conclusion

The study contributes to the body of knowledge, and the industry is five-fold: The real-time tyre detection and classification model was developed by using ResNet-50 and YOLOv7. The ResNet-50 model is used for tyre quality assessment, and the YOLOv7 model is utilised for real-time detection in a complex environment. Transfer learning was adopted by both models, which could do the specific task of distinguishing between defective and functional tyres, demonstrating reasonable improvements in accuracy, efficiency, and real-time detection capability.

The ResNet-50 model, fine-tuned for binary classification, achieved a macro-average F1-score of 0.95, with high precision and recall across both tyre categories. Its strong performance underscores its reliability in assessing tyre quality, making it an effective tool for automated inspection and preventive maintenance in the automotive industry. While the model conservatively minimised false positives, minor adjustments could further reduce false negatives, improving its accuracy in identifying defective tyres.

The ResNet-50 model was used for binary classification, which achieved a macro-average F1-score of 0.95 with high precision and recall across both tyre categories. The model's high performance shows its reliability in assessing tyre quality, making it an effective tool for automated inspection and preventive maintenance in the automotive industry. The model's ability can be improved by minor tuning which could further reduce false negatives, improving its accuracy in identifying defective tyres.

Similarly, YOLOv7 shows excellent performance in real-time tyre detection. The model achieved a mean average precision (mAP@0.5) of 0.90 and a detection rate of 45 frames per second (FPS). These aspects of YOLOv7 particularly make it suitable for industrial applications such as tyre sorting in recycling and waste management operations. Its performance and accuracy in the detection of tyres in cluttered environments make it suitable for large-scale waste management and sustainability initiatives.

The combination of ResNet-50 for quality classification with YOLOv7 for detection makes it a powerful tyre monitoring and inspection solution. From the business perspective tyre recycling and reuse is possible through predictive maintenance, reduces manual inspection errors, and supports environmental sustainability and increased profits.

7. Limitations and Scope for Future Work

As with any study, this study is also not without limitations. One of the major limitations of this study is that the ResNet-50 model is to under-identify defective tyres. This is because of the trade-off between precision and recall, as indicated by a slightly higher rate of false negatives than false positives. This raises concerns about potentially missing some defective tyres—a critical issue in real-world settings where undetected defects could pose safety risks.

To further enhance the work and the model performance, the following improvements can be made:

1. To reduce the false negatives in ResNet-50, refining the hyperparameter tuning and including more defective tyre samples can be done (Kuric et al., 2021).

2. The latest versions of YOLOX can be used to improve real-time object detection, which may increase detection accuracy and speed (Vasan et al. (2021).

3. In the real business setup, real-time deployment and edge computing could be useful to facilitate large-scale tyre inspection in the recycling facility (Listyalina et al., 2022).

4. Transfer learning can enable few-shot learning, and also lessen the data collection burden and improve the model's performance in identifying rare or previously unseen tyre defects (Zhu et al., 2021).

5. Drone-assisted tyre detection can enable environmental protection by covering larger areas.

References

Bochkovskiy, A., Wang, C. Y., & Liao, H. Y. M. (2020). YOLOv4: Optimal Speed and Accuracy of Object Detection. arXiv preprint arXiv:2004.10934.

Deng, H., Li, X., & Yang, G. (2023). Sustainability of key tire materials and the development of green tires. Sustainability Science, 16(5), 345-357.

Harshitha, R., & Samala, V. (2020). A survey on OpenCV and CNN models for tire wear detection. International Journal of Computer Vision and Applications, 28(1), 102-110.

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/10.1109/CVPR.2016.90

Julie, M., Rajesh, S., & Kumar, R. (2021). Multi-state CNNs for tire tread analysis: Predicting wear and durability. Proceedings of the IEEE Conference on Image Processing, 156-162.

Kazmi, A., Rafiq, M., & Ahmed, S. (2022). Industrial vehicle tire detection and text recognition using CNN classifiers. Journal of Industrial Applications, 56(4), 201-209.

Kuric, I., Janik, A., & Frantisek, B. (2021). Unsupervised learning and deep learning for tire defect inspection. Automation in Industry, 45(6), 233-240.

Li, B., & Lima, D. (2021). Facial expression recognition via ResNet-50. International Journal of Cognitive Computing in Engineering, 2, 57-64. https://doi.org/10.1016/j.ijcce.2021.02.002

Listyalina, M., Santos, J., & Oliveira, R. (2022). Tire quality assessment using convolutional neural networks. Computer Vision and Machine Learning Journal, 47(3), 98-112. https://doi.org/10.31315/telematika.v19i3.7697

Maharana, K., Mondal, S., & Nemade, B. (2022). A review: Data pre-processing and data augmentation techniques. Global Transitions Proceedings, 3(1), 91-99. https://doi.org/10.1016/j.gltp.2022.04.020

Reddy, G., Kumar, P., & Singh, A. (2021). Automated tire wear detection using OpenCV and CNNs. International Journal of Automotive Engineering, 39(7), 456-462.

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. https://doi.org/10.1109/CVPR.2016.91

Siegel, J. E., Sun, Y., & Sarma, S. (2018). Automotive diagnostics as a service: An artificially intelligent mobile application for tire condition assessment. In Artificial Intelligence and Mobile Services-AIMS 2018: 7th International Conference, Held as Part of the Services Conference Federation, SCF 2018, Seattle, WA, USA, June 25-30, 2018, Proceedings 7 (pp. 172-184). Springer International Publishing. https://doi.org/10.1007/978-3-319-94361-9_13

Vasan, D., Malarvizhi, S., & Ravi, K. (2021). Real-time tire condition monitoring using transfer learning and deep neural networks. Journal of Intelligent Transportation Systems, 29(3), 345-356.

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. In Proceedings of the IEEE/CVF conference on computer vision and pattern recognition (pp. 7464-7475). https://doi.org/10.1109/CVPR52729.2023.00721

Xiao, Z., Pramanik, A., Basak, A. K., Prakash, C., & Shankar, S. (2022). Material recovery and recycling of waste tyres-A review. Cleaner Materials, 5, 100115. https://doi.org/10.1016/j.clema.2022.100115

Zhang, Y., Wu, X., & Chang, Z. (2020). Tire damage recognition using CNN-based deep learning models. Journal of Automotive Safety, 32(2), 156-162.

Zhu, H., Han, P., & Wang, F. (2021). Predicting tire life through image processing and machine learning. Pattern Recognition Letters, 145, 34-42.