SUMMARY

The primary objective of this thesis was to evaluate the effectiveness of the YOLOv8 algorithm for object detection in industrial quality control. The analysis focused on efficiency, accuracy, and computational resource requirements. The process was divided into several stages: data collection, experimental setup, algorithm implementation, and comparative analysis.

The YOLOv8 algorithm, known for its real-time object detection capabilities, was chosen due to its balance between speed and accuracy. The training process utilised a custom dataset provided by RAIKU Packaging OÜ, consisting of images of wooden springs used in packaging, annotated for defects. The training was conducted using Python within the Jupyter Notebook environments.

The training of the YOLOv8n model, a nano version optimised for being lightweight, was carried out over 100 epochs using a robust hardware setup. The hardware included an NVIDIA GeForce RTX 3060 GPU, AMD Ryzen 7 5800H CPU, 32GB RAM, and 1TB NVMe SSD. The training and validation processes were monitored through various metrics, including training loss, validation loss, mean average precision (mAP), precision, recall, and F1 score.

Key results included the following: both box loss and classification loss decreased consistently, indicating effective learning. The validation loss mirrored the training loss, showing good generalisation. The final mAP on the validation set was 16.6% for mAP@0.5, suggesting moderate accuracy. Precision peaked at 1.00 at a confidence level of 0.425. Recall showed higher values at lower confidence levels, indicating the model's ability to detect most objects. The F1 score balanced precision and recall, with a peak F1 score of 0.29 at a confidence level of 0.089. The confusion matrix highlighted areas of strength and weakness in classifying defects, with some difficulty noted in detecting specific defect classes.

The results of this thesis demonstrated the YOLOv8n model's capability to perform object detection tasks with reasonable accuracy and efficiency. The training process showed effective learning patterns, and the evaluation metrics provided a comprehensive understanding of the model's performance.

The main hypothesis of this study was that YOLOv8, being a newer algorithm, is expected to exhibit better performance in terms of training efficiency and accuracy compared to widely used algorithms like R-CNN and its variants. This hypothesis was supported by findings from related literature. Studies indicated that YOLOv8 demonstrated faster training times and required fewer computational resources than Faster R-CNN, making it more efficient for real-time applications. Literature evidence showed that YOLOv8 achieved competitive accuracy levels while maintaining high speed, which is crucial for industrial quality control.

Moreover, the trained model successfully detected cracks in the wooden springs provided by RAIKU Packaging OÜ. For quality control purposes, any detection of cracks, regardless of the number, would classify the product as defective. The results prove that the practical object detection problem from the company was solved effectively.

However, several limitations were identified. The moderate mAP and F1 score indicated room for improvement in detection accuracy. The confusion matrix revealed specific challenges in classifying certain defect types, suggesting that the model could benefit from further tuning and additional training data. In terms of computational efficiency, the YOLOv8n model performed well given the hardware constraints, but better hardware specifications could potentially enhance training speed and model performance.

Future work should focus on addressing these limitations by exploring more advanced data augmentation techniques, fine-tuning hyperparameters, and increasing the dataset size. Additionally, comparing the YOLOv8n model with other emerging object detection algorithms could provide further insights into its relative strengths and areas for improvement.

Overall, this thesis confirmed the critical role of advanced object detection algorithms in industrial quality control. The findings contribute valuable knowledge to the field and highlight the potential for ongoing improvements and innovations in automated quality control systems. The proven hypothesis underscores YOLOv8's superiority in efficiency and accuracy, making it a more viable solution for real-time object detection challenges in industrial environments compared to traditional algorithms like Faster R-CNN, as evidenced by related literature.