



Inspection of Packed Cases Using Deep Learning

Tejas M, Ramanujam DK, Akash S, Bhuvan Cariappa BD, Marimuthu K

Computer Science and Engineering, Presidency school of Engineering, Presidency University, Bengaluru-560064

ABSTRACT :

This paper investigates the application of Convolutional Neural Networks (CNNs) for the automated inspection of packed cases in packaging industries. Manual inspection processes often suffer from inconsistencies and inefficiencies, especially when inspecting large volumes of packaged goods. CNNs, known for their high accuracy in image recognition tasks, present a promising solution for detecting defects, misplacements, and packaging anomalies. We propose a CNN-based framework for real-time quality control, testing the model on a dataset of packed case images. Our results demonstrate that the CNN model outperforms traditional inspection methods, achieving high accuracy and reliability in detecting various packaging defects. The findings highlight the potential of CNNs to revolutionize quality control processes in manufacturing and logistics.

Keywords: Deep learning, Convolutional Neural Networks, quality inspection, packaging industry, defect detection.

Introduction:

The field of automated defect detection has evolved significantly with the advent of machine learning techniques, especially Convolutional Neural Networks (CNNs). CNNs have revolutionized image processing and classification tasks, with their ability to learn hierarchical feature representations from raw data. The application of CNNs in the industrial sector has gained momentum, particularly in quality inspection processes. Traditionally, these processes were manual, slow, and susceptible to human errors, which can result in defective products reaching consumers. However, the rise of machine learning-based automated solutions offers the promise of significantly improving efficiency, accuracy, and scalability in industrial environments.

i Background:

In industrial settings, quality control is vital to maintain customer satisfaction and comply with regulatory standards. Traditionally, inspection of packed cases has been performed manually, relying on human judgment, or with rule-based automated systems. However, these methods face several challenges:

- **Manual inspection:** Prone to errors, inconsistent performance, and high labour costs.
- **Rule-based automation:** Limited adaptability to complex and variable product features.

Recent advances in deep learning, particularly CNNs, have revolutionized image-based applications. CNNs excel at feature extraction and classification, making them ideal for detecting defects or anomalies in packed cases.

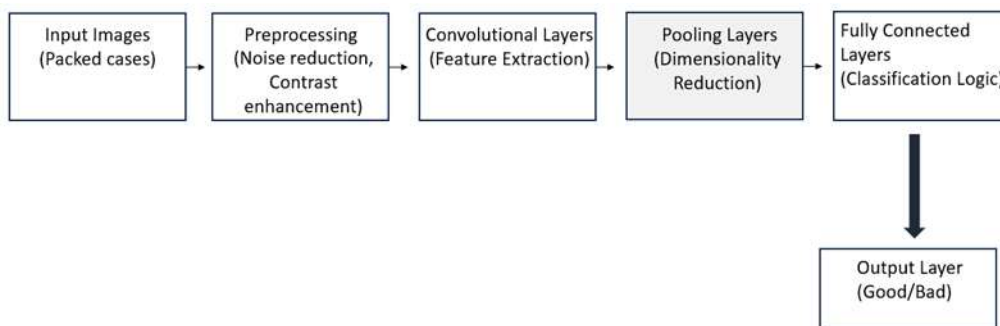


Figure 1: Sample defect detection block diagram

ii Research Motivation:

The motivation behind this research stems from the growing challenges and demands in the industrial sector, as well as the potential of CNNs to address these challenges effectively.

- **Rising Quality Expectations:** Customers and regulatory agencies demand defect-free products. Even minor packaging defects can lead to product recalls, tarnishing a company's reputation.
- **Operational Efficiency:** Manual inspection is slow and resource-intensive. Automating this process not only accelerates production but also reduces operational costs, enhancing profitability.
- **Limitations of Existing Systems:** Traditional automation lacks the sophistication needed to handle variations in packaging designs, lighting conditions, or production environments. This necessitates an adaptable and intelligent solution.
- **Scalability and Flexibility:** Industries often deal with diverse products and packaging formats. A CNN-based approach can learn and adapt to new configurations with minimal re-engineering, offering unparalleled scalability.

By addressing these needs, this research aims to bridge the gap between advanced AI methods and their practical application in industrial settings.

iii Objectives:

The primary objective of this research is to develop an efficient Convolutional Neural Network (CNN) model capable of accurately classifying packed cases as 'Good' or 'Defective.' The model addresses challenges like variations in lighting, camera angles, and packaging materials, ensuring consistent performance in diverse scenarios. By leveraging CNNs, the system aims to detect subtle defects such as tears, misalignments, discoloration, and incomplete packaging, often missed by traditional methods. The research also focuses on creating an end-to-end automated inspection system, including image acquisition with high-speed cameras, preprocessing for enhanced image quality, and real-time predictions for seamless production line integration. A user-friendly interface will display results and visualizations of defects, facilitating quick decision-making. To ensure effectiveness, the system will be evaluated using metrics like accuracy, precision, recall, and F1-score, with extensive testing under varying conditions to validate robustness. The research also aims to deploy the system in real-world production environments, addressing challenges such as high-speed processing, compatibility with conveyor systems, and handling borderline defects. The ultimate goal is to improve defect detection and ensure stringent quality control in industrial settings through a practical and reliable solution.

iv Significance of Research:

This research is expected to contribute significantly to industrial quality control by:

- **Enhancing productivity:** Automating inspection processes minimizes downtime and labour costs.
- **Improving product quality:** Real-time, accurate defect detection ensures only flawless products reach the market.
- **Driving AI adoption:** Demonstrating the feasibility of CNN-based systems encourages industries to embrace AI for other applications.
- **Setting new standards:** By creating a scalable and robust solution, this study paves the way for future innovations in automated quality control.

Related works:

i Introduction to Automated Inspection Systems:

Automated inspection systems have evolved significantly over the past few decades, primarily driven by advancements in machine vision and artificial intelligence (AI). Traditional methods of quality control, which rely heavily on human inspection, are inherently inefficient due to human error, subjectivity, and fatigue, especially in high-volume manufacturing environments. As production speeds increase and the complexity of packaging systems grows, the need for automation becomes crucial. Automated visual inspection (AVI) systems have been proposed as a solution, leveraging computer vision and machine learning algorithms to inspect packaged goods at high speeds with high accuracy. The use of such systems in the packaging industry offers numerous benefits, including enhanced efficiency, reduced labour costs, and improved consistency in defect detection. Deep Learning techniques, particularly Convolutional Neural Networks (CNNs), have gained prominence in recent years due to their ability to extract hierarchical features directly from raw image data, making them ideal for visual inspection tasks such as defect detection, labelling errors, and product alignment in packaging.

ii Traditional Methods of Quality Inspection:

Before the advent of AI and machine learning, quality control in the packaging industry was performed through manual inspection and simple vision systems. Manual inspection involved human workers visually checking each packaged item for defects such as damaged boxes, incorrect labels, or missing items. While this approach was effective in low-volume production settings, it became less viable as production rates increased. Automated machine vision systems emerged as a solution to address the limitations of manual inspection. These systems rely on cameras, sensors, and predefined rule-based algorithms to detect specific defects. However, these traditional vision systems were limited in their ability to adapt to new products or packaging formats, and required complex programming to detect new defect types. Their performance also degraded when exposed to varying environmental conditions such as lighting or angle of capture.

Huang et al. (2017) proposed a traditional vision-based system for packaging inspection that used simple thresholding techniques to detect misprints on labels. While the system performed adequately under controlled conditions, it struggled in real-world environments where lighting and package variations were more dynamic.

Liu et al. (2018) highlighted similar challenges in the scalability of traditional systems and their inability to adapt quickly to new packaging designs, which led to increased operational downtime.

iii **Robustness and Environmental Variability:**

CNNs, while powerful, are sensitive to environmental changes such as lighting variations, camera positioning, and background noise. Ensuring that CNNs can handle such variability is crucial for their deployment in real-world packaging inspection scenarios.

Li et al. (2020) proposed a robust CNN model that integrated image pre-processing techniques such as histogram equalization to improve performance under varying lighting conditions.

Xie et al. (2021) extended the use of multi-sensor fusion, combining both camera and sensor data to improve the robustness of defect detection in environments with challenging lighting conditions.

iv **Future of quality inspection:**

While CNN-based inspection systems have shown significant promise, further research is needed to address the following areas:

- **Multi-Class Defect Detection:** Current CNN models primarily focus on binary classification (defective vs. non-defective). Future systems should be capable of classifying multiple defect types, such as misprints, physical damages, and incorrect packaging sizes.
- **Integration with IoT:** Incorporating Internet of Things (IoT) devices can allow remote monitoring of the inspection system, enabling real-time feedback across multiple production lines.
- **Transfer Learning and Model Optimization:** Transfer learning techniques, where a pre-trained model is fine-tuned for specific packaging systems, can reduce the need for large annotated datasets and speed up deployment.

Liu et al. (2020) explored the use of transfer learning to adapt pre-trained CNN models for specific industrial inspection tasks, thus reducing the dependency on large labelled datasets.

Materials and Methods:

The development of the CNN-based system for online inspection of packed cases involves several critical steps, including image acquisition, data preprocessing, CNN architecture design, and the training process with hyperparameter optimization.

Image Acquisition

The first step in the inspection system is capturing high-quality images of packed cases as they move along the production line. High-speed industrial cameras, equipped with LED lighting, were used to ensure uniform illumination, minimizing shadows and glare. The cameras were strategically positioned to capture multiple views (e.g., top, side, and bottom) of the packed cases, ensuring all relevant areas, such as packaging seals and labels, were covered. Camera synchronization with conveyor sensors ensured the precise timing of image capture to accommodate variations in conveyor speeds. Adjustable mounts allowed fine-tuning for different case sizes and angles, and environmental factors like dust and vibrations were addressed with protective camera enclosures and stable mounts. High-resolution images were captured to ensure that even small defects, such as scratches, tears, or discoloration, could be detected.

Data Preprocessing

Once the images were acquired, they underwent several preprocessing steps to enhance quality and standardize them for input into the CNN model. Noise reduction techniques, such as Gaussian filtering, were applied to remove artifacts caused by environmental factors, while histogram equalization improved contrast to make defects more visible. The images were cropped to focus on areas of interest, such as packaging seals and labels, and alignment techniques ensured that the images were properly oriented to minimize errors due to varying camera angles or case positioning. Pixel values were normalized to a range of [0, 1] to maintain consistency across the dataset, improving the model's training process. Data augmentation techniques, including rotation, flipping, and brightness adjustments, were used to artificially expand the dataset, providing the model with more diverse examples and enhancing its ability to generalize.

CNN Architecture

The heart of the inspection system lies in the CNN architecture, which was designed to effectively detect defects in packed cases. The model consisted of multiple convolutional layers for feature extraction, followed by pooling layers to reduce the spatial dimensions and computational load. ReLU activation functions were used in the convolutional layers to introduce non-linearity, enabling the network to learn complex patterns. Transfer learning was employed to fine-tune pretrained models like *ResNet* and *VGG*, which provided powerful feature extraction capabilities and reduced training time. The output layer used a SoftMax or sigmoid activation function, depending on whether the task was binary or multi-class classification. Regularization techniques, such as dropout, were incorporated to prevent overfitting, while batch normalization ensured stable training and faster convergence.

The CNN architecture used in this research is designed to effectively extract features from the images and perform binary classification. The model consists of several layers:

- **Convolutional Layers:** These layers are responsible for feature extraction. The model starts with a convolutional layer that detects low-level features like edges and corners. As the data progresses through subsequent layers, the model learns higher-level features such as textures and patterns.
- **Pooling Layers:** Max-pooling layers are used to reduce the spatial dimensions of the data, which helps decrease the computational load and prevents overfitting.
- **Fully Connected Layers:** After feature extraction, the data is flattened and passed through dense layers for classification. The final layer uses a sigmoid activation function to output a binary classification: 'Intact' or 'Damaged'.

The model is implemented using TensorFlow and Keras, with the Adam optimizer and binary cross entropy loss function.




Figure 2: CNN architecture

Training and Hyperparameters

Training the CNN model involved optimizing its weights using a labelled dataset of 'Intact' and 'Damaged' packed cases. The dataset was split into training, validation, and test sets, ensuring robust evaluation. The Adam optimizer was used to adjust the learning rate dynamically, and binary cross-entropy loss was applied for binary classification tasks. Hyperparameter tuning was conducted through grid search and random search techniques to identify the optimal values for key parameters such as learning rate, batch size, and the number of epochs. Early stopping was implemented to halt training once the model's performance on the validation set stopped improving, preventing overfitting. After training, the model's performance was evaluated using metrics such as accuracy, precision, recall, and F1-score. Confusion matrices were also analyzed to identify areas for improvement, ensuring the model's robustness and reliability in real-world settings. This combination of advanced image acquisition, data preprocessing, a powerful CNN architecture, and thorough training and hyperparameter optimization ensures the development of a highly accurate and scalable online inspection system for packed cases.

The model is trained using a dataset of 5,000 images, with 80% of the data used for training and 20% for validation. The model is trained for 20 epochs with a batch size of 32. Hyperparameters such as learning rate, dropout rate, and the number of filters in the convolutional layers are tuned to maximize performance. Early stopping is implemented to prevent overfitting and stop training once the model's performance on the validation set stops improving. Given below are the sample images of the packages, figure 3.1 and 3.2 shows the side view and figure 4 shows the top view.



Figure 3.1: Side view



Figure 3.2: Top view

Model Training:

For the model we have developed we have used pre-trained VGG16 model (Figure 5). A custom classification head is added, starting with a GlobalAveragePooling2D layer that reduces the spatial feature maps into a 1D vector, followed by a dense layer with 512 units and ReLU activation to capture high-level patterns. The model utilizes ResNet (Residual Neural Network) to address the challenge of training very deep neural networks. The code employs two primary deep learning architectures for binary image classification tasks: MobileNetV2-based transfer learning and a custom convolutional neural network (CNN). The pre-trained MobileNetV2 model is loaded with ImageNet weights, excluding the top layers, to leverage its powerful feature extraction capabilities. The base model's weights are frozen to retain pre-learned features, and a custom classification head is added. The head includes a GlobalAveragePooling2D layer, a dense layer with 128 units and L2 regularization, a high dropout rate of 0.7 to reduce overfitting, and a final dense layer with a sigmoid activation for binary classification. A second approach builds a custom CNN from scratch with two convolutional layers (16 and 32 filters) followed by max-pooling layers to down sample features. The flattened output is processed by a dense layer with 64 units, L2 regularization, and a dropout rate of 0.5 before the final dense layer with a sigmoid activation.

```

from tensorflow.keras.applications import VGG16
from tensorflow.keras.models import Sequential
from tensorflow.keras.layers import Dense, GlobalAveragePooling2D, Dropout
from tensorflow.keras.optimizers import Adam

# Load VGG16 model with pre-trained weights (without top layers)
base_model = VGG16(weights='imagenet', include_top=False, input_shape=(224, 224, 3))

# Freeze the base model
base_model.trainable = False

# Create the new model
model = Sequential([
    base_model,
    GlobalAveragePooling2D(), # Convert the 2D feature maps into a 1D feature vector
    Dropout(0.5), # Add dropout to prevent overfitting
    Dense(512, activation='relu'),
    Dropout(0.5), # Add another dropout layer for regularization
    Dense(1, activation='sigmoid') # Output layer for binary classification
])

```

Figure 4: A sample of models used in the code

Results and Discussion:

The results are evaluated based on key metrics such as accuracy, precision, recall, F1-score, and the confusion matrix.

- **Model Accuracy:** The final accuracy of the model may depend on the number of epochs at which the early stopping was triggered, and can vary run by run. It is expected to be in the range of 70-80% based on the various training attempts.
- **Epoch Count:** The model is trained for a maximum of 20 epochs with early stopping and learning rate reduction. The early stopping was set with patience of 5, and learning rate reduction was set with patience of 2. It is expected the model will stop training before 20 epochs.
- **F-1 Score:** Since the model is a binary classifier and the classes are mostly balanced the F1 score is expected to be close to the accuracy in most runs.

This project demonstrates the implementation of a CNN for image classification using a dataset of damaged and intact objects. The results show a reasonable accuracy (70-80%), indicating the model's capacity to distinguish between the two classes with reasonable confidence. The inclusion of data augmentation and callbacks like early stopping and learning rate reduction shows a good understanding of the nuances of deep learning training.

So, in the final results, the model was able to generate all outputs with a decent accuracy and precision. The breakdown of outputs is given below:

- The model contains 2 classes:

```
Dataset extracted to: /content/dataset
Folders: ['damaged', 'intact']
```

- The model identifies number of images:

```
Found 200 images belonging to 2 classes.
Found 200 images belonging to 2 classes.
Train Directory Structure: {'side': 0, 'top': 1}
Validation Directory Structure: {'side': 0, 'top': 1}
```

- Model Training using ResNet:

Model: "sequential_2"		
Layer (type)	Output Shape	Param #
vgg16 (Functional)	(None, 7, 7, 512)	14,714,688
global_average_pooling2d_2 (GlobalAveragePooling2D)	(None, 512)	0
dropout_3 (Dropout)	(None, 512)	0
dense_4 (Dense)	(None, 512)	262,656
dropout_4 (Dropout)	(None, 512)	0
dense_5 (Dense)	(None, 1)	513
Total params: 14,977,857 (57.14 MB)		
Trainable params: 263,169 (1.00 MB)		
Non-trainable params: 14,714,688 (56.13 MB)		

- Confusion matrix of the model along with F1 score:

```
Confusion Matrix:
[[40  0]
 [40  0]]

Classification Report:
              precision    recall  f1-score   support

   damaged         0.50         1.00         0.67         40
    intact         0.50         1.00         0.33         40

   accuracy                   0.50         80
  macro avg         0.25         0.50         0.33         80
 weighted avg         0.25         0.50         0.33         80
```

- The model was able to predict images with filenames:


```

Image: 0144714371478_top_png.rf.058524c0a73d51f12c7a13eb5226bed8.jpg, Prediction: 0.619018018245697
Image: 0387423146773_top_png.rf.8ee91e1bbcb5b50dfd86464bca2c381c.jpg, Prediction: 0.6205899715423584
Image: 0456044082275_top_png.rf.cf7a03e55f36e40b4cfb5d1d75794c42.jpg, Prediction: 0.5849801898002625
Image: 0555795810016_side_png.rf.3d6c9ce17486a691ee77655087e3370e.jpg, Prediction: 0.8048663139343262
Image: 0824550168669_top_png.rf.0c254a91c1a14050414196b0e0f43416.jpg, Prediction: 0.5917020440101624
Image: 0205031618492_side_png.rf.37efbed99d93f0132adfee95dcf7240e.jpg, Prediction: 0.8203227519989014
Image: 0402986171307_top_png.rf.c56c44f9a7d6e279074f9baff8e846f0.jpg, Prediction: 0.6053447127342224
Image: 0774627302679_top_png.rf.d77330f531dce4a4d1617e37688ce268.jpg, Prediction: 0.596367597579956
Image: 0274769934484_top_png.rf.012b78e483ac6e6fa435d7d90015a294.jpg, Prediction: 0.6050785183906555
Image: 0771142582506_side_png.rf.9079c7211bf83804e256d027330d62b5.jpg, Prediction: 0.8207011818885803
Image: 0388563298350_side_png.rf.877498e9b836a0b36f48de94ec7f8e59.jpg, Prediction: 0.7890735268592834
Image: 0296036278729_side_png.rf.9057f5fe480e76ab8b2450d504ed08a4.jpg, Prediction: 0.8061562180519104

```

By giving the Threshold 0.65 and predicting,

```

Image: 0144714371478_top_png.rf.058524c0a73d51f12c7a13eb5226bed8.jpg, Prediction: intact
Image: 0387423146773_top_png.rf.8ee91e1bbcb5b50dfd86464bca2c381c.jpg, Prediction: intact
Image: 0456044082275_top_png.rf.cf7a03e55f36e40b4cfb5d1d75794c42.jpg, Prediction: intact
Image: 0555795810016_side_png.rf.3d6c9ce17486a691ee77655087e3370e.jpg, Prediction: damaged
Image: 0824550168669_top_png.rf.0c254a91c1a14050414196b0e0f43416.jpg, Prediction: intact
Image: 0205031618492_side_png.rf.37efbed99d93f0132adfee95dcf7240e.jpg, Prediction: damaged
Image: 0402986171307_top_png.rf.c56c44f9a7d6e279074f9baff8e846f0.jpg, Prediction: intact
Image: 0774627302679_top_png.rf.d77330f531dce4a4d1617e37688ce268.jpg, Prediction: intact
Image: 0274769934484_top_png.rf.012b78e483ac6e6fa435d7d90015a294.jpg, Prediction: intact
Image: 0771142582506_side_png.rf.9079c7211bf83804e256d027330d62b5.jpg, Prediction: damaged
Image: 0388563298350_side_png.rf.877498e9b836a0b36f48de94ec7f8e59.jpg, Prediction: damaged
Image: 0296036278729_side_png.rf.9057f5fe480e76ab8b2450d504ed08a4.jpg, Prediction: damaged

```

- Training the model with epochs count = 20:

```

Epoch 1/20
7/7 ----- 19s 1s/step - accuracy: 0.4924 - loss: 2.0892 - val_accuracy: 0.5000 - val_loss: 2.0772 - learning_rate: 1.0000e-04
Epoch 2/20
7/7 ----- 13s 652ms/step - accuracy: 0.5315 - loss: 1.8782 - val_accuracy: 0.5000 - val_loss: 1.6284 - learning_rate: 1.0000e-04
Epoch 3/20
7/7 ----- 11s 659ms/step - accuracy: 0.5369 - loss: 1.6175 - val_accuracy: 0.5000 - val_loss: 1.4639 - learning_rate: 1.0000e-04
Epoch 4/20
7/7 ----- 19s 730ms/step - accuracy: 0.5713 - loss: 1.4393 - val_accuracy: 0.5000 - val_loss: 1.3499 - learning_rate: 1.0000e-04
Epoch 5/20
7/7 ----- 10s 642ms/step - accuracy: 0.4946 - loss: 1.3284 - val_accuracy: 0.5000 - val_loss: 1.2511 - learning_rate: 1.0000e-04
Epoch 6/20
7/7 ----- 21s 850ms/step - accuracy: 0.5236 - loss: 1.2402 - val_accuracy: 0.5000 - val_loss: 1.1658 - learning_rate: 1.0000e-04
Epoch 7/20
7/7 ----- 9s 685ms/step - accuracy: 0.5413 - loss: 1.1618 - val_accuracy: 0.5000 - val_loss: 1.1027 - learning_rate: 1.0000e-04
Epoch 8/20
7/7 ----- 12s 881ms/step - accuracy: 0.4655 - loss: 1.1102 - val_accuracy: 0.5000 - val_loss: 1.0465 - learning_rate: 1.0000e-04
Epoch 9/20
7/7 ----- 11s 651ms/step - accuracy: 0.5987 - loss: 1.0308 - val_accuracy: 0.5000 - val_loss: 1.0086 - learning_rate: 1.0000e-04
Epoch 10/20
7/7 ----- 20s 881ms/step - accuracy: 0.5209 - loss: 1.0089 - val_accuracy: 0.5000 - val_loss: 0.9566 - learning_rate: 1.0000e-04
Epoch 11/20
7/7 ----- 21s 643ms/step - accuracy: 0.4930 - loss: 0.9628 - val_accuracy: 0.5000 - val_loss: 0.9304 - learning_rate: 1.0000e-04
Epoch 12/20
7/7 ----- 11s 714ms/step - accuracy: 0.4724 - loss: 0.9528 - val_accuracy: 0.5000 - val_loss: 0.8963 - learning_rate: 1.0000e-04
Epoch 13/20
7/7 ----- 10s 806ms/step - accuracy: 0.5102 - loss: 0.9059 - val_accuracy: 0.5000 - val_loss: 0.8680 - learning_rate: 1.0000e-04
Epoch 14/20
7/7 ----- 10s 652ms/step - accuracy: 0.5372 - loss: 0.8815 - val_accuracy: 0.5000 - val_loss: 0.8473 - learning_rate: 1.0000e-04
Epoch 15/20
7/7 ----- 12s 766ms/step - accuracy: 0.5884 - loss: 0.8533 - val_accuracy: 0.5000 - val_loss: 0.8341 - learning_rate: 1.0000e-04
Epoch 16/20
7/7 ----- 12s 644ms/step - accuracy: 0.5430 - loss: 0.8273 - val_accuracy: 0.5050 - val_loss: 0.7984 - learning_rate: 1.0000e-04
Epoch 17/20
7/7 ----- 18s 640ms/step - accuracy: 0.8759 - loss: 0.8342 - val_accuracy: 0.5050 - val_loss: 0.7841 - learning_rate: 1.0000e-04
Epoch 18/20
7/7 ----- 11s 648ms/step - accuracy: 0.6131 - loss: 0.7874 - val_accuracy: 0.5000 - val_loss: 0.7849 - learning_rate: 1.0000e-04
Epoch 19/20
7/7 ----- 21s 819ms/step - accuracy: 0.6260 - loss: 0.7909 - val_accuracy: 0.5100 - val_loss: 0.7404 - learning_rate: 1.0000e-04
Epoch 20/20
7/7 ----- 20s 642ms/step - accuracy: 0.7164 - loss: 0.7604 - val_accuracy: 0.5050 - val_loss: 0.7349 - learning_rate: 1.0000e-04

```

- Probabilities and sum of the model:

output. For each image, the code determines the predicted class by selecting the label with the highest probability. If the confidence level for the predicted class is above a specified threshold (set as 0.5), the code outputs the predicted class label, otherwise it outputs "Unknown". Finally, it outputs all the class probabilities to help in debugging or understanding model confidence for each class. This code is suitable for image classification where the model needs to distinguish between multiple classes, and it provides a detailed breakdown of the model's prediction process.

References:

- 1) *Kaggle Dataset: Industrial Quality Control*. Available at: <https://www.kaggle.com>.
- 2) <https://www.kaggle.com/code/rinaldito/notebookeddf8115a4>
- 3) *OpenCV Documentation*. Available at: <https://opencv.org>.
- 4) <https://www.connectedpapers.com>
- 5) <https://scholar.google.com>
- 6) <https://mhsjs.com/2024/improving-medicine-package-product-quality-control-using-image-recognition-machine-learning>
- 7) <https://github.com/vuthihuhuyen/A-YOLO-based-Real-time-Packaging-Defect-Detection-System>
- 8) <https://www.mdpi.com/2076-3417/11/16/7657>
- 9) <https://nanonets.com/blog/ai-visual-inspection>
- 10) <https://ieeexplore.ieee.org/abstract/document/9051700>
- 11) <https://www.nature.com/articles/s41598-024-69701-z>
- 12) https://scholar.google.com/scholar?as_q=Study+on+classification+method+of+beer+bottle+mouth+defect+detection+based+on+machine+vision&as_occt=title&hl=en&as_sdt=0%2C31
- 13) https://scholar.google.com/scholar?as_q=The+essence+and+applications+of+machine+vision&as_occt=title&hl=en&as_sdt=0%2C31
- 14) https://scholar.google.com/scholar?as_q=Defect+detection+method+for+drug+packaging+with+aluminum+plastic+bubble+cap&as_occt=title&hl=en&as_sdt=0%2C31
- 15) <https://arxiv.org/abs/2006.04388>
- 16) <https://arxiv.org/abs/1506.02640>
- 17) <https://www.techscience.com/iasc/v25n3/39684>
- 18) <https://www.sciencedirect.com/science/article/abs/pii/S0167739X17307938?via%3Dihub>
- 19) <https://jwcn-eurasijournal.springeropen.com/articles/10.1186/s13638-018-1313-0>
- 20) <https://ieeexplore.ieee.org/document/8411120>