Fish Species Classification Report

Author: Enes Yamac

Model: Custom CNN (trained from scratch, no transfer learning)

Dataset: 9-class fish species dataset

This report summarises the training setup and results of a convolutional neural network for classifying nine fish species. All characters are ASCII to avoid rendering issues.

Date: 02.12.2025

1. Methods

The model is a sequential convolutional neural network trained from scratch (no pre-trained backbone).
Input images are resized to 224 by 224 pixels and normalised by dividing pixel values by 255.
The training and validation generators use the same image size and normalisation settings.

Main architecture:
- Input: 224 x 224 x 3 RGB image.
- Four convolutional blocks with 3 x 3 filters and ReLU activation.
- Batch normalisation after each convolutional layer.
- Max pooling after every two convolutional layers.
- Feature flattening layer.
- Dense layer with 256 units and ReLU activation.
- Dropout layer with rate 0.5 for regularisation.
- Output dense layer with 9 units and softmax activation.

The model is trained with the Adam optimiser and sparse categorical cross entropy loss.
We track two metrics during training: accuracy and top 3 accuracy.
Learning rate scheduling and early stopping are used to stabilise training and avoid overfitting.

2. Results

Final validation scores of the best saved model:
- Validation accuracy: about 84 percent.
- Top 3 accuracy: about 95 percent.
- Validation loss: about 0.76.

Class level performance summary:
- Excellent: Black Sea Sprat and Shrimp (precision and recall close to 1.00).

- Very good: Red Sea Bream (F1 score around 0.94).
- Good: Horse Mackerel, Trout and Striped Red Mullet (F1 score around 0.80 to 0.84).
- More difficult classes: Gilt Head Bream (about 0.69), Red Mullet (about 0.75) and Sea Bass (about 0.57).
These species are visually similar, which explains most confusion.

The confusion matrix confirms that most mistakes happen between species that look similar.
For a model trained from scratch, the overall performance is strong and consistent.

3. Reflection

The project shows that a well designed CNN trained from scratch can reach good performance on a multi class fish dataset.
Transfer learning with a model such as ResNet50 might still give higher accuracy, especially for visually similar species, but the custom CNN offers a good balance between complexity and performance.

Possible future improvements include hyper parameter tuning, stronger data augmentation and class specific analysis
to reduce confusion between similar fish types. The trained model has also been deployed as a Streamlit web application
that allows users to upload an image and see the predicted fish species and class probabilities.