

Dog vs Cat Classification using Transfer Learning

Project Overview

This project implements an image classification system to distinguish between dog and cat images using transfer learning. A pre trained convolutional neural network is leveraged to extract meaningful visual features and build an efficient deep learning classifier.

Dataset Description

The dataset consists of labeled images of dogs and cats organized in separate directories. Images vary in size and orientation and represent real world visual variations that make classification a challenging task.

Data Preprocessing

Images are resized to a fixed resolution and normalized to ensure consistent input for the neural network. Data generators are used to efficiently load images and apply preprocessing during training.

Transfer Learning Approach

A pre trained convolutional neural network model trained on the ImageNet dataset is used as a feature extractor. The convolutional base is frozen to retain learned representations while custom classification layers are added on top.

Model Architecture

The model architecture includes the frozen convolutional base followed by fully connected dense layers. Nonlinear activation functions are applied to learn complex patterns that differentiate dogs and cats.

Model Compilation

The model is compiled using the Adam optimizer and binary crossentropy loss function which is suitable for two class classification problems. Accuracy is used as the primary evaluation metric.

Model Training

The model is trained on the processed image dataset for multiple epochs. Validation data is used to monitor performance and reduce overfitting during training.

Model Evaluation

After training the model is evaluated on a test dataset to measure classification accuracy on unseen images of dogs and cats.

Prediction System

The trained model can classify new images by predicting whether the input image represents a dog or a cat based on learned visual features.

Key Takeaways

This project demonstrates how transfer learning significantly reduces training time and improves performance in image classification tasks by leveraging pre trained deep learning models.

Author

Satyam Gajjar