# CIFAR Object Recognition using ResNet50

## Project Overview

This project implements an object recognition system using deep learning and the ResNet50 architecture. The model is trained on the CIFAR image dataset to classify objects into predefined categories using transfer learning.

## Dataset Description

The CIFAR dataset consists of small color images representing multiple object categories. Each image contains three color channels and captures a wide variety of objects under different visual conditions.

## Data Preprocessing

Images are resized and normalized to match the input requirements of the ResNet50 model. Proper preprocessing ensures stable training and effective feature extraction.

## ResNet50 Architecture

ResNet50 is a deep convolutional neural network that uses residual connections to enable efficient training of very deep models. These connections help prevent vanishing gradient problems and improve learning performance.

## Transfer Learning Strategy

A pre trained ResNet50 model trained on ImageNet is used as a feature extractor. The convolutional layers are frozen and custom classification layers are added on top to adapt the model to the CIFAR dataset.

## Model Compilation

The model is compiled using the Adam optimizer and categorical crossentropy loss function. Accuracy is used as the evaluation metric to measure classification performance.

## Model Training

The model is trained for multiple epochs using training and validation datasets. Training progress is monitored to ensure effective learning and prevent overfitting.

## Model Evaluation

After training the model is evaluated on a test dataset to measure its performance on unseen images. The evaluation results demonstrate the effectiveness of transfer learning for object recognition tasks.

## Prediction System

The trained model can predict object categories for new images by analyzing learned visual features extracted by the ResNet50 network.

## Key Takeaways

This project demonstrates how deep residual networks combined with transfer learning can achieve strong performance in object recognition tasks using relatively small datasets.

## Author

Satyam Gajjar