Image Classification using AlexNet on CIFAR-10

A Deep Learning Project Case Study

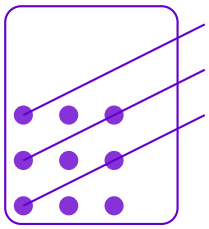

1. 1

Overview

This project implements a convolutional neural network based on the AlexNet architecture to classify images from the CIFAR-10 dataset. The dataset consists of 60,000 color images from 10 categories. The aim is to build a high-performance CNN with generalization and training efficiency.

2. 1

Objective

- Implement an AlexNet-style model using PyTorch.
- Achieve over 80% accuracy on the CIFAR-10 dataset.
- Apply data augmentation and regularization techniques.

3. 1

Tools and Technologies

- **Python** for scripting
- **PyTorch** for model building
- **Matplotlib & Sklearn** for evaluation
- **CUDA GPU** acceleration
- CIFAR-10 dataset

4. 1

Methodology Flow

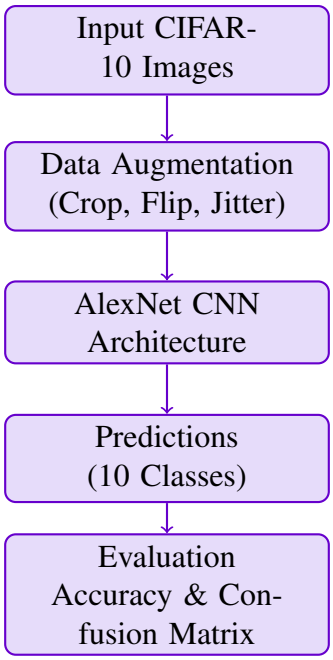

5. 1

Model Architecture

The network contains 3 convolutional blocks, each followed by ReLU, BatchNorm, and MaxPooling. This is followed by fully connected layers: 1024 → 512 → 10.

6. 1

Training Configuration

- **Loss Function:** CrossEntropyLoss
- **Optimizer:** Adam (LR = 0.001, weight decay = 1e-4)
- **Scheduler:** StepLR (step = 20, gamma = 0.5)
- **Epochs:** 50 **Batch Size:** 128

7. 1

Results

- **Final Training Accuracy:** 90%
- **Test Accuracy: 84.62%**
- Overfitting reduced via augmentation and dropout
- Confusion matrix showed balanced class-wise accuracy

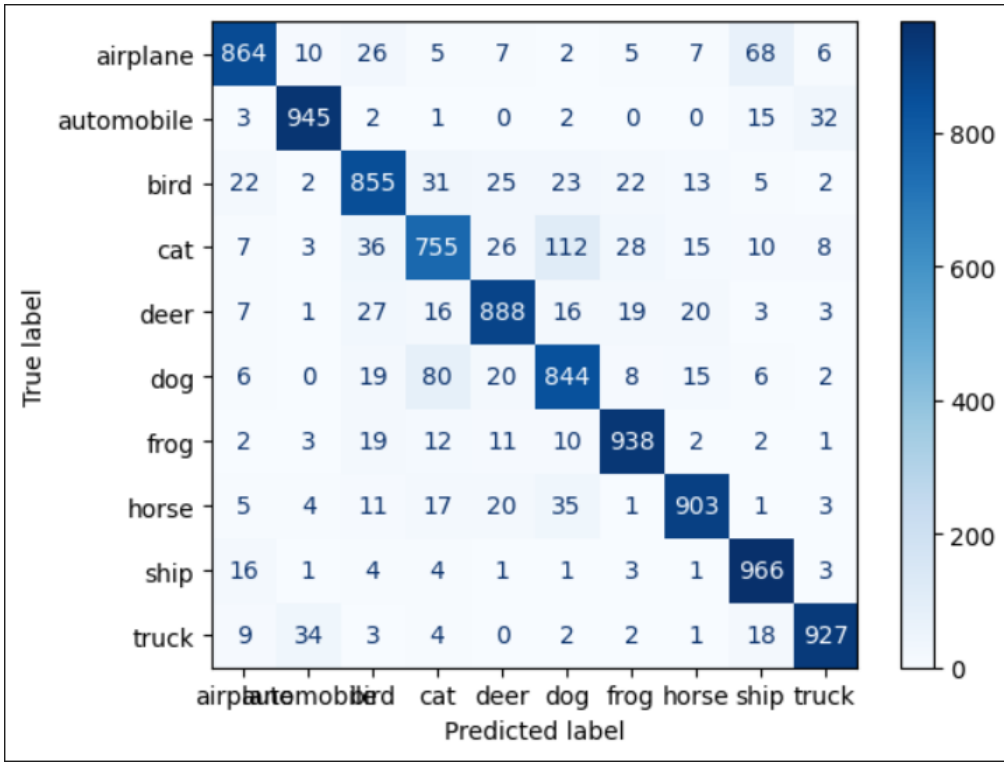

Figure 1: Confusion matrix of CIFAR-10 classification using AlexNet

8. 1

Key Learnings

- Data augmentation improves generalization significantly.
- Classic architectures like AlexNet still perform well.
- Visualization tools are vital for model diagnostics.

"This project enhanced my understanding of CNNs, especially in training stability, evaluation strategies, and practical application of data augmentation."