Brock University
Faculty of Mathematics & Science
Department of Computer Science

COSC 4P80 Final Project

Prepared By:
Brett Terpstra #6920201
bt19ex@brocku.ca

Steve Mastrokalos #7276900
sm21ks@brocku.ca

Instructor:
Dave Bockus

January, 2025

Abstract

Deep Learning folds feature extraction into the traditional neural network architecture. This paper will use the MNIST handwritten digit dataset to show the benefit in using deep learning techniques.

Table of Contents

Chapter 1

Introduction

As previously mentioned, deep learning combines feature extraction through convolution and pooling with traditional neural networks, eliminating the need for humans to manually extract features from datasets. Convolution, in essence, is a filtering process where a trained filter slides over the input data to extract features and other useful information. Pooling is the subsequent process of taking local samples and selecting either the minimum, maximum, or average of those samples. This step helps identify feature locations and condenses the information produced by the convolution layer.

A typical deep learning pipeline consists of several convolution and pooling layers, followed by a few fully connected layers. In this work, we aim to demonstrate that using a deep learning configuration can reduce the size of the feed-forward section without compromising program performance, thereby highlighting the effectiveness of deep learning.

The MNIST database is a standard benchmark for image-processing neural networks. For our comparison, we will use a modified version of the DLIB deep learning example. This approach allows us to showcase the differences between standard feed-forward neural networks and deep learning networks without requiring expensive GPUs or AI accelerators. While the MNIST dataset is solvable using feed-forward neural networks, we intend to demonstrate that deep learning can achieve better classification performance, even on smaller networks.

Chapter 2

Experimental Setup

Our experiments are divided into two parts, each testing a deep learning network alongside its corresponding feed-forward network. For a fair comparison, the feed-forward test focuses explicitly on the feed-forward component of the deep learning network. This ensures that variables such as the number of layers or nodes in the feed-forward section remain consistent, minimizing potential biases and maintaining the integrity of our comparisons.

2.1   Experiment 1

2.2   Experiment 2