

Intelligent World Informatics Lecture V

Final Project

Paul Lödige
Student ID: 37-229753
paul.loedige@student.kit.edu

February 3, 2023

1	Assignment	2
1.1	Additional Details	2
2	Project Goal	3
2.1	The Dataset	3
3	Methods	3
3.1	Dense Neural Network	3
3.2	Convolutional Neural Network (CNN)	3
4	Comparison	4
4.1	VGG16	4
5	Summary	4
A	Appendix	5
A.1	Dense Neural Network Comparison	5
A.2	CNN Comparison	6
A.3	Transfer Learning VGG16	7
A.4	Code	8

1 Assignment

- Define your own project
 - Describe the task and the goal you would like to achieve
 - Experiment on any public dataset provided by Tensorflow:
 - Full list: <https://www.tensorflow.org/datasets/catalog/overview>
 - Simple datasets: https://www.tensorflow.org/api_docs/python/tf/keras/datasets
 - Do one of the followings:
 1. Apply two different deep learning techniques you learned from this class
 - e.g.: Comparing Dense DNN vs RNN method on the same dataset (what pros and cons?)
 - e.g.: Combining CNN and RNN on a dataset (why the combination is preferable?)
 2. Apply transfer learning with two different base models
 - Explain the reasons why the base models are reasonable choices.
 - Survey on the internet with keywords such as “best model for imagenet/cifar10/mnist”, “tensorflow pre-trained models”, etc, to find the name of the model
 - Download the base model through tensorflow
 - Just like in the code, simply change the name of the base model. List of available base models in Tensorflow:
 - https://www.tensorflow.org/api_docs/python/tf/keras/applications
 - Add layers, dropouts, use different learning rate, epoch, batch_size, etc.

1.1 Additional Details

- Write 1 ~ 2 page(s) (excluding the code) describing your project:
 - Imagine the project as a kind of “mini research paper”
 - What is the problem? What is your motivation?
 - How are you solving them? What deep learning algorithms?
 - Outline the experiment and test on the dataset.
 - What are the results?
 - Add at least 1 figure and insert your analysis
 - Attach the code at the end.
 - Include your name, affiliation, student number, as well as your university e-mail address in case if we need to contact you for clarification
 - Put all the above in a single PDF and upload to ITC-LMS

2 Project Goal

This project aims to compare different deep learning techniques. The goal is to discern the pros and cons of the different techniques when it comes to their use on character recognitions in images. This approach was chosen because it doesn't focus on the learning of object concepts but rather on pattern recognition. Since pattern recognition and character recognition in particular is a less complex topic than concept classification it allows for the use of less complex models which are easier to understand.

2.1 The Dataset

The project uses the "svhn_cropped" dataset (https://www.tensorflow.org/datasets/catalog/svhn_cropped). The Street View House Numbers (SVHN) dataset was created at Stanford University and includes "73 257 digits for training, 26 032 digits for testing, and 531 131 additional, somewhat less difficult samples, to use as extra training data" (<http://ufldl.stanford.edu/housenumbers/>). The dataset provided by TensorFlow uses the "MNIST-like 32-by-32 images centered around a single character" version.

3 Methods

3.1 Dense Neural Network

The first method tried to find a good classification model was to use regular Dense Neural Networks. Through experimentation it was determined that the classification accuracy only varied slightly when using a different amount of hidden layers (see appendix A.1).

3.2 Convolutional Neural Network (CNN)

When it comes to image classification CNNs are often used as they reduce the dimensionality of the input. This in turn simplifies the problem and makes it easier to classify. Multiple configurations for CNNs were tried (see appendix A.2) The best CNN found during experimentation had the following configuration:

```
0 configs['convolutional neural network (3 convolutions)'] = {
1     'model': models.Sequential([
2         layers.Conv2D(32, (3,3), activation='relu', input_shape=(32,32,3)),
3         layers.MaxPooling2D((2,2)),
4         layers.Conv2D(64, (3,3), activation='relu'),
5         layers.MaxPooling2D((2,2)),
6         layers.Conv2D(128, (3,3), activation='relu'),
7         layers.MaxPooling2D((2,2)),
8         layers.Flatten(),
9         layers.Dense(64, activation='relu'),
10        layers.Dense(64, activation='relu'),
11        layers.Dense(10, activation='softmax')
12    ]),
13    'optimizer': 'adam',
14    'loss': 'sparse_categorical_crossentropy',
15    'metrics': ['accuracy'],
16    'epochs': 10,
17    'batch_size': 128
18 }
```

4 Comparison

While the Dense Neural Networks can be trained efficiently (a few seconds per epoch), they are also less accurate than the CNNs. As can be seen in fig. 1 the Dense Neural Networks achieved a validation accuracy of around 70 %. The CNNs on the other hand achieved around 80 – 85 %.

It can also be seen that the type of model has a bigger impact on the accuracy than its complexity. Even the CNN with just one convolutional layer performs better than the Dense neural network with 5 hidden dense layers.

4.1 VGG16

As can be seen from the comparison of the two previous methods CNNs hold a clear advantage when it comes to image classification. In order to find an even better model for classification it therefore makes sense to look at CNNs developed by other researchers. VGG16 is one CNN that is commonly used for these kind of tasks and is also available as part of the `tensorflow.keras` library. While the attempt at transfer learning was unsuccessful (see appendix A.3), training the whole model pre-trained on the "Imagenet" dataset with a learning rate of 0.001 produced a model with even higher accuracy than the CNNs mentioned before. While the training of this model is very time intensive it gave a validation accuracy of over 92 % (fig. 1).

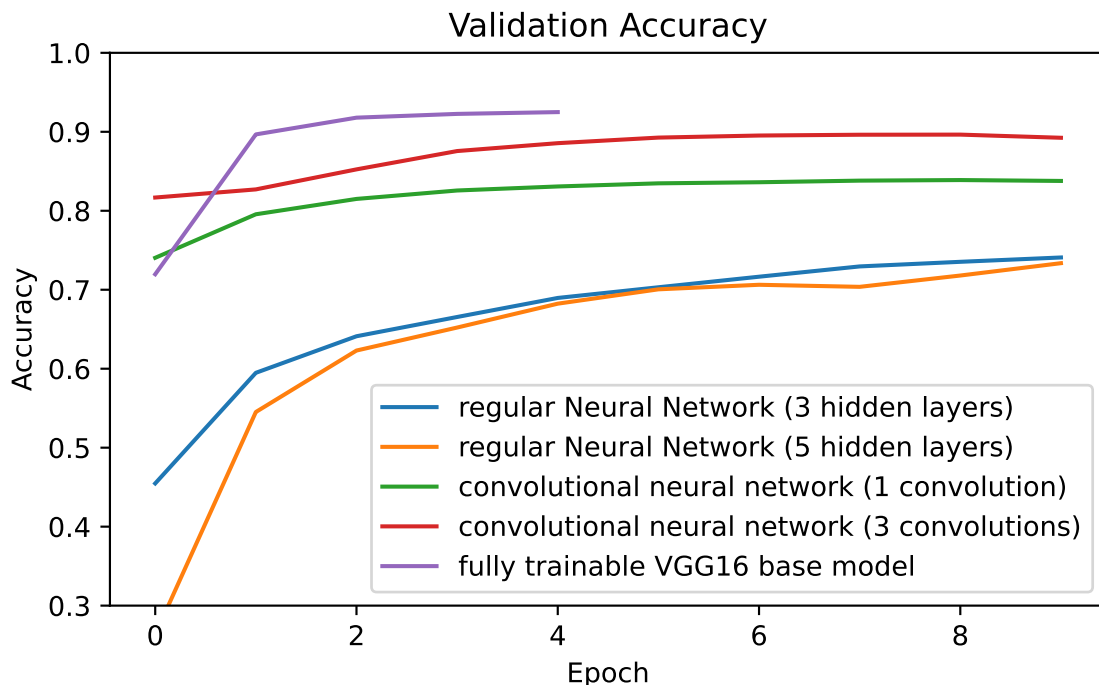



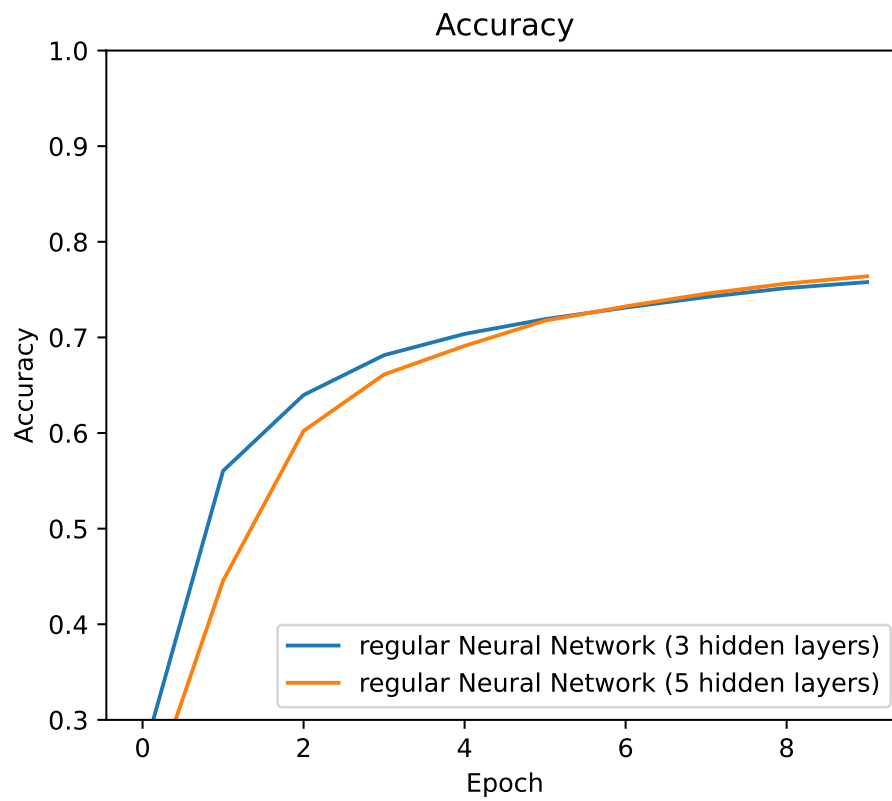
Figure 1: Validation Accuracy of the different models

5 Summary

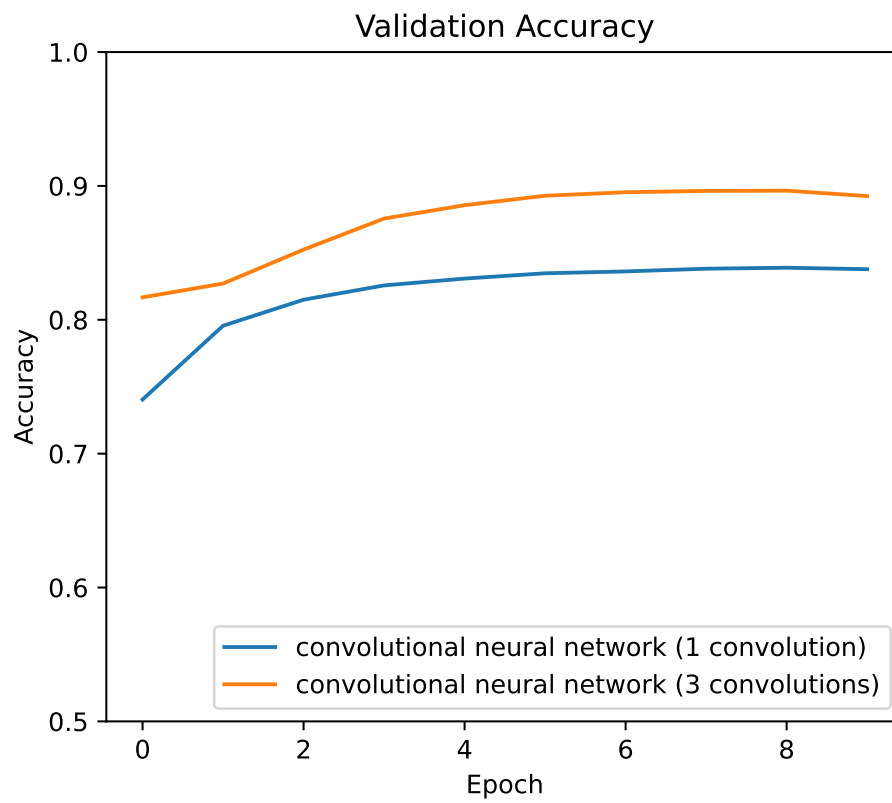

In this project it has been shown that Convolutional Neural Networks outperform Dense Neural Networks when it comes to classifying the SVHN dataset. Additionally, if the goal is to achieve the highest accuracy possible one should use a pre-existing model designed for solving problems in a similar context.

A Appendix


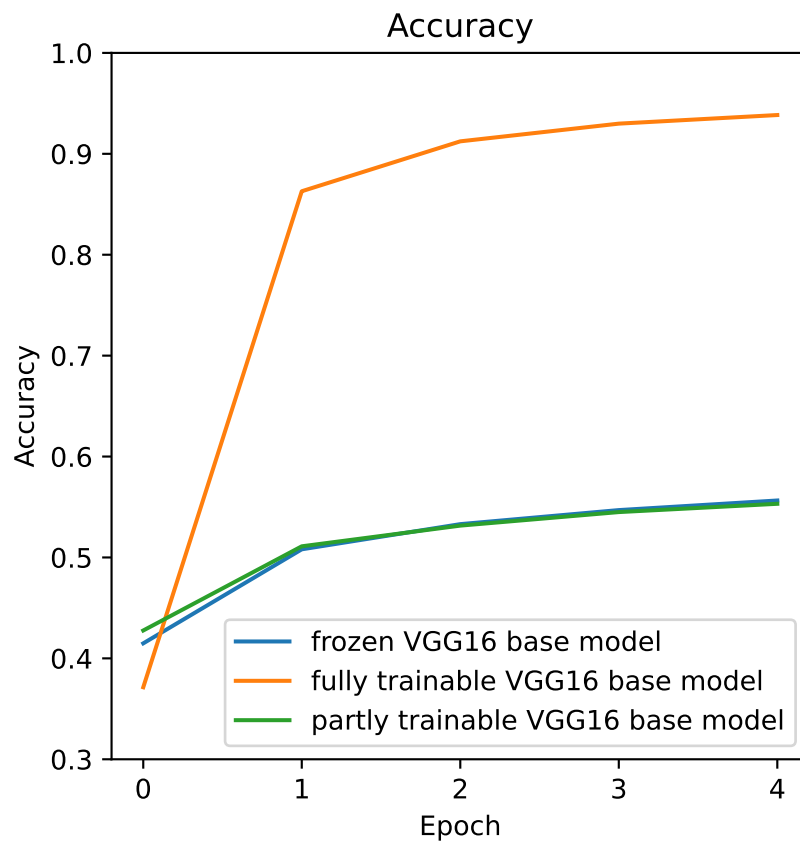
A.1 Dense Neural Network Comparison



A.2 CNN Comparison



A.3 Transfer Learning VGG16



A.4 Code

The following code is an exported and slightly formatted version of a Jupyter Notebook that is also available at https://git.ploedige.com/Intelligent_World_Informatics_V/Final_Project.

```
0 # Import libraries
1 import matplotlib.pyplot as plt
2 import tensorflow as tf
3 from tensorflow.keras import layers, models, applications, optimizers
4 import tensorflow_datasets as tfds
5 import tkinter.messagebox # for notifications
6
7 # Load Dataset
8 # https://www.tensorflow.org/datasets/keras_example
9 # split into training and test data
10 ds_train, ds_test = tfds.load(
11     'svhn_cropped',
12     split=['train', 'test'],
13     shuffle_files=True,
14     as_supervised=True
15 )
16
17 # Confirm Data
18 for image, label in ds_train.take(1):
19     plt.imshow(image)
20     plt.show()
21
22 # Define normalization function
23 def normalize_img(image, label):
24     return tf.cast(image, tf.float32) / 255.0, label
25
26 # Build a training pipeline
27 ds_train = ds_train.map(normalize_img, num_parallel_calls=tf.data.AUTOTUNE)
28 ds_train = ds_train.cache()
29 ds_train = ds_train.batch(128)
30 ds_train = ds_train.prefetch(tf.data.AUTOTUNE)
31
32 # Build an evaluation pipeline
33 ds_test = ds_test.map(normalize_img, num_parallel_calls=tf.data.AUTOTUNE)
34 ds_test = ds_test.cache()
35 ds_test = ds_test.batch(128)
36 ds_test = ds_test.prefetch(tf.data.AUTOTUNE)
37
38 # Configurations
39 configs = {}
40
41 configs['regular Neural Network (3 hidden layers)'] = {
42     'model': models.Sequential([
43         layers.Flatten(input_shape=(32, 32, 3)),
44         layers.Dense(256, activation='relu'),
45         layers.Dense(128, activation='relu'),
46         layers.Dense(64, activation='relu'),
47         layers.Dense(10, activation='softmax')
48     ]),
49     'optimizer': 'adam',
50     'loss': 'sparse_categorical_crossentropy',
51     'metrics': ['accuracy'],
52     'epochs': 30,
53     'batch_size': 64
54 }
55
56 configs['regular Neural Network (5 hidden layers)'] = {
57     'model': models.Sequential([
58         layers.Flatten(input_shape=(32, 32, 3)),
59         layers.Dense(1024, activation='relu'),
60         layers.Dense(512, activation='relu'),
61         layers.Dense(256, activation='relu'),
62         layers.Dense(128, activation='relu'),
63         layers.Dense(64, activation='relu'),
64         layers.Dense(10, activation='softmax')
65     ]),
66     'optimizer': 'adam',
67     'loss': 'sparse_categorical_crossentropy',
68     'metrics': ['accuracy'],
```



```

69         'epochs': 30,
70         'batch_size': 64
71     }
72
73 configs['convolutional neural network (1 convolution)'] = {
74     'model': models.Sequential([
75         layers.Conv2D(32, (3,3), activation='relu', input_shape=(32,32,3)),
76         layers.MaxPooling2D((2,2)),
77         layers.Flatten(),
78         layers.Dense(64, activation='relu'),
79         layers.Dense(64, activation='relu'),
80         layers.Dense(10, activation='softmax')
81     ]),
82     'optimizer': 'adam',
83     'loss': 'sparse_categorical_crossentropy',
84     'metrics': ['accuracy'],
85     'epochs': 10,
86     'batch_size': 64
87 }
88
89 configs['convolutional neural network (3 convolutions)'] = {
90     'model': models.Sequential([
91         layers.Conv2D(32, (3,3), activation='relu', input_shape=(32,32,3)),
92         layers.MaxPooling2D((2,2)),
93         layers.Conv2D(64, (3,3), activation='relu'),
94         layers.MaxPooling2D((2,2)),
95         layers.Conv2D(128, (3,3), activation='relu'),
96         layers.MaxPooling2D((2,2)),
97         layers.Flatten(),
98         layers.Dense(64, activation='relu'),
99         layers.Dense(64, activation='relu'),
100        layers.Dense(10, activation='softmax')
101    ]),
102    'optimizer': 'adam',
103    'loss': 'sparse_categorical_crossentropy',
104    'metrics': ['accuracy'],
105    'epochs': 10,
106    'batch_size': 128
107 }
108
109 vgg16_base_model = applications.VGG16(input_shape=(32,32,3), include_top= False)
110 # freeze the VGG16 model
111 vgg16_base_model.trainable = False
112 configs['frozen VGG16 base model'] = {
113     'model': models.Sequential([
114         vgg16_base_model,
115         layers.Flatten(),
116         layers.Dense(64, activation='relu'),
117         layers.Dense(32, activation='relu'),
118         layers.Dense(10, activation='softmax'),
119     ]),
120     'optimizer': 'adam',
121     'loss': 'sparse_categorical_crossentropy',
122     'metrics': ['accuracy'],
123     'epochs': 5,
124     'batch_size': 64
125 }
126
127 vgg16_base_model = applications.VGG16(input_shape=(32,32,3), include_top= False)
128 configs['fully trainable VGG16 base model'] = {
129     'model': models.Sequential([
130         vgg16_base_model,
131         layers.Flatten(),
132         layers.Dense(64, activation='relu'),
133         layers.Dense(32, activation='relu'),
134         layers.Dense(10, activation='softmax'),
135     ]),
136     'optimizer': optimizers.Adam(learning_rate=0.001),
137     'loss': 'sparse_categorical_crossentropy',
138     'metrics': ['accuracy'],
139     'epochs': 5,
140     'batch_size': 64
141 }
142
143 vgg16_base_model = applications.VGG16(input_shape=(32,32,3), include_top= False)
144 # freeze the VGG16 model

```

```

145 vgg16_base_model.trainable = False
146 for layer in vgg16_base_model.layers[-6:]:
147     layer.trainable=True
148
149 configs['partly trainable VGG16 base model'] = {
150     'model': models.Sequential([
151         vgg16_base_model,
152         layers.Flatten(),
153         layers.Dense(64, activation='relu'),
154         layers.Dense(32, activation='relu'),
155         layers.Dense(10, activation='softmax'),
156     ]),
157     'optimizer': optimizers.Adam(learning_rate=0.001),
158     'loss': 'sparse_categorical_crossentropy',
159     'metrics': ['accuracy'],
160     'epochs': 5,
161     'batch_size': 64
162 }
163
164 # Compile the Models
165 for config in configs.values():
166     if 'history' not in config.keys():
167         config['model'].compile(
168             optimizer=config['optimizer'],
169             loss=config['loss'],
170             metrics=config['metrics']
171         )
172
173 # Train and Evaluate the Models
174 try:
175     for config_name, config in configs.items():
176         if 'history' in config.keys():
177             print(f'Already trained model "{config_name}"')
178         else:
179             print(f'Now training model "{config_name}"')
180             config['history'] = config['model'].fit(ds_train, epochs=config['epochs'],
181 validation_data=ds_test, batch_size=config['batch_size'])#, verbose=0)
182 except Exception as e:
183     print(e)
184     tkinter.messagebox.showerror("ERROR", f"ERROR: {e}")
185
186 # Plot the accuracy
187 plt.figure(figsize=(10,5))
188 for config_name, config in configs.items():
189     plt.plot(config['history'].history['accuracy'], label=config_name)
190 plt.xlabel('Epoch')
191 plt.ylabel('Accuracy')
192 # plt.ylim([0.75, 1])
193 plt.legend(loc='lower right')
194 plt.title("Accuracy")
195 plt.show()
196
197 # Plot the validation accuracy
198 plt.figure(figsize=(10,5))
199 for config_name, config in configs.items():
200     plt.plot(config['history'].history['val_accuracy'], label=config_name)
201 plt.xlabel('Epoch')
202 plt.ylabel('Accuracy')
203 # plt.ylim([0.8, 1])
204 plt.legend(loc='lower right')
205 plt.title("Validation Accuracy")
206 plt.show()
207
208 # Notify when done
209 tkinter.messagebox.showinfo("DONE", "DONE")

```