

Section 1: Data Set Summary & Exploration


- The training set contains 34,799 traffic sign images
- The validation set contains 4410 images
- The test set contains 12,630 images
- The shape of a traffic sign image is 32x32x3
- The number of unique classes/labels in the data set is 43

Section 2: Exploratory visualization of the dataset

Below is a visualization of the training data set. It's a bar chart which shows how many example images there are per class and as you can see, the distribution is uneven. For example there are relatively fewer examples of signs that have labels 0,6,19,20,21 etc.



Below is a similar visualization of the validation data set. The distribution here is also uneven and is similar to the distribution of the training set.



I did not examine any characteristics of the test set. I thought that set should remain untouched and unknown because I wanted to keep the final test of the model as objective and unbiased as possible.

Section 3: Design and Test a Model Architecture

Preprocessing

I converted the images to grayscale in order to speed up their processing by the convolutional network. Initially I tried to get the model to train on the original unprocessed images (i.e. with all 3 colour channels intact) because I thought having colour information would help the network classify things, but doing this zeroed out the weights in the network for some reason. Why this happened is still unknown to me.


I then normalized the gray-scaled images so that pixel values range between 0 and 1. Below is a sequence of 3 images which shows an image before processing, after grayscaling and then after normalization has been performed.



In order to make sure the pre-processing didn't mangle the data, I printed out the ranges of some randomly chosen normalized images, the means of the pixel values and the standard deviations. The print-out below shows these statistics for the normalized image shown above.

```
The range of the randomly chosen normalized image is 0.098231827112 to 1.00196463654
The mean is 0.555860398248
The std is 0.298733666847
```





I also constructed histograms of the pixel values of randomly chosen normalized images. The histogram below shows the distribution of pixel values in the grayscale and normalized picture shown above:



Augmenting the training data

I decided to augment the training data because acquiring more data is often the best way to improve the performance of a model. I generated additional images by translating them in the x and y directions by random amounts ranging from 0 to 3 pixels. I also zoomed out of the images, making the signs smaller, in order to augment the data. (I experimented with zooming **in** on images by a factor of 1.3 as well but including these images in the training data significantly decreased the accuracy of the model. These zoomed in images were therefore not used to augment the training set).

The row of images below shows an image before and then after translation and the next row shows an image before and after zooming outwards:



Initially there were 34,799 training examples but augmenting the data tripled the size of the training dataset to 139,196 examples.

Model architecture

The model is almost identical to the LeNet model except that I have added another fully-connected layer. My model consists of two convolutional layers, 2 max pooling layers, and 4 fully connected layers. Also, I use the ReLU activation function rather than the sigmoid function used in the LeNet model because recent research suggests that the ReLU achieves better results.

Layer	Description
Input	32x32x1 grayscale and normalized images
Convolution 5x5	1x1 stride, padding='valid', outputs 28x28x24
Relu	
Max Pooling	2x2 stride, padding='valid', outputs 14x14x24
Convolution 5x5	1x1 stride, padding='valid', outputs 10x10x64
Relu	
Max Pooling	2x2 stride, padding='valid' outputs 5x5x64
Fully connected	500
Relu	
Fully connected	200
Relu	
Fully connected	100
Relu	
Fully connected	43

Model Training

I chose to implement a LeNet architecture. Accuracy of this model hovered around 91% when the system was trained with the given training set. I decided then to augment the training data set by translating and zooming out of the images. Augmenting the data in this way tripled the training set size from 34,799 to 139,196 examples. I also experimented with zooming **in** on images by a factor of 1.3 but including these images in the training data significantly decreased the accuracy of the model. These zoomed in images were therefore not used to augment the training set.

When I used the LeNet architecture with an Adam optimizer, a batch size of 128 and a learning rate of 0.0001 on the augmented data set, validation accuracy only barely exceeded 93% after 10 epochs and started fluctuating up and down at epoch 6. So I experimented by changing the activation function from the sigmoid function to ReLU. (Recent research suggests that the ReLU achieves better results than nearly all other activation functions).

I also added another fully connected layer because I thought that going from 1600 nodes (in the final convolution layer) to 43 classes in the space of just 3 fully connected layers was maybe too abrupt and might possibly throw away important information. This was an admittedly intuitive decision but given that there is no rigorous way right now of deciding beforehand on how many levels of architecture there should be for a given problem or what the

values of certain hyperparameters should be, I simply tried this out and the accuracy of the system improved. When this layer was added to the model, an accuracy of 94.5% was achieved on the validation set after 6 epochs of training.

The final results were therefore a training set accuracy of 97%, validation set accuracy of 94.5% and a test set accuracy of 93.6%

Section 4: Test a model on new images

I downloaded 8 new traffic sign images from the internet and they are shown below.



8 new images downloaded from the web. The model misclassified all of them.

The correct labels for these images are, starting from the top left and moving across each row from left to right: [29,22,18,28,25,27,23,14]. These images, which I'll refer to as the 'internet images', varied in size and so the program's first step was to resize them to have a shape of 32x32x3. They were then transformed into grayscale images and normalized in the same way the training set images were. However the classifier incorrectly classified all of these images.

Interestingly, most of these new internet images (seen above) have blue backgrounds whereas the training set images tend to have whiter or blacker backgrounds like this:



Some examples of training set images that have the same labels as the 'internet images'

In fact, after looking at randomly sampled examples of the training data I hardly ever came across an image that had a blue background. For example below is a random sampling of training set images. Perhaps this is what confused the network although you would think that transforming the images into grayscale would avoid or minimize this problem.



12 randomly chosen images from the training set

Another reason why the model performed so poorly on the internet images may be that the size of the signs within the training set images differ than those in the internet set. More experimentation with zooming out or in on images in order to make them appear more like the training data would address this.

The result of this experiment with the internet images is very interesting, but sobering, because it clearly shows that if a training set differs from a test set in ways that are significant, but hard for the human eye to discern, the performance of the neural network can be very poor. It appears that the training set and the set of internet images that I downloaded are drawn from different distributions that are somehow fundamentally different, even though both sets of data appear to be very similar. This really highlights how important it is that the validation set be drawn from the same distribution as the final test set. Andrew Ng makes this argument at point 41:45 in the following lecture: <https://www.youtube.com/watch?v=F1ka6a13S9I&t=65s>

I don't know how the German Traffic signs that I used for my training and validation sets were assembled but training a network using them but then testing the system on images that were not created or assembled in the same way shows how brittle a neural network solution can be and how important it is to train the network on the same kind of data with which you evaluate its final performance.

Model Certainty - Softmax probabilities

Let's look at the top 5 softmax probabilities for each of the 8 internet images to see if, at the very least, the correct label for a sign appears in the model's top 5 guesses for that sign. Here are the top 5 softmax probabilities that the neural network produced when choosing a label for each of the 8 internet images:

```
TopKV2 (values=array([[ 4.02324885e-01,   3.94627869e-01,   1.86988935e-01,
                        7.23987212e-03,   5.87017182e-03],
 [ 7.80312121e-01,   8.57489482e-02,   5.28570227e-02,
   4.08709645e-02,   3.81706990e-02],
 [ 3.58170956e-01,   2.78487951e-01,   2.01521531e-01,
   9.97604281e-02,   2.15502661e-02],
 [ 9.99967456e-01,   3.25112815e-05,   1.21842021e-08,
   1.51397506e-09,   1.27155421e-15],
 [ 5.15053153e-01,   3.92139792e-01,   5.97165935e-02,
   2.63524912e-02,   4.96213650e-03],
 [ 5.05172014e-01,   3.38022560e-01,   1.03151195e-01,
   1.95008721e-02,   1.26849907e-02],
 [ 8.39147210e-01,   1.55942380e-01,   3.03602335e-03,
   1.76994957e-03,   9.84320868e-05],
 [ 9.99927998e-01,   5.81595341e-05,   5.86297983e-06,
   3.64434595e-06,   3.54727058e-06]], dtype=float32), indices=array([[23, 30, 29,
 19, 20],
 [25, 24, 23, 29, 30],
 [16, 12,  9, 17, 10],
 [18, 27, 26, 11, 24],
 [23, 11, 28, 30, 20],
 [22, 26,  8, 29, 18],
 [27, 18, 24, 11, 30],
 [23, 20, 29, 31, 30]], dtype=int32))
```

The correct labels for the internet images are [29,22,18,28,25,27,23,14].

The correct label for internet image #1 is 29 and it appears 3rd in the top 5 probabilities.

The correct label for internet image #2 is 22 and it doesn't appear in the top 5 probabilities.
The correct label for internet image #3 is 18 and it doesn't appear in the top 5 probabilities.
The correct label for internet image #4 is 28 and it doesn't appear in the top 5 probabilities.
The correct label for internet image #5 is 25 and it doesn't appear in the top 5 probabilities.
The correct label for internet image #5 is 27 and it doesn't appear in the top 5 probabilities.
The correct label for internet image #6 is 23 and it doesn't appear in the top 5 probabilities.
The correct label for internet image #8 is 14 and it doesn't appear in the top 5 probabilities.

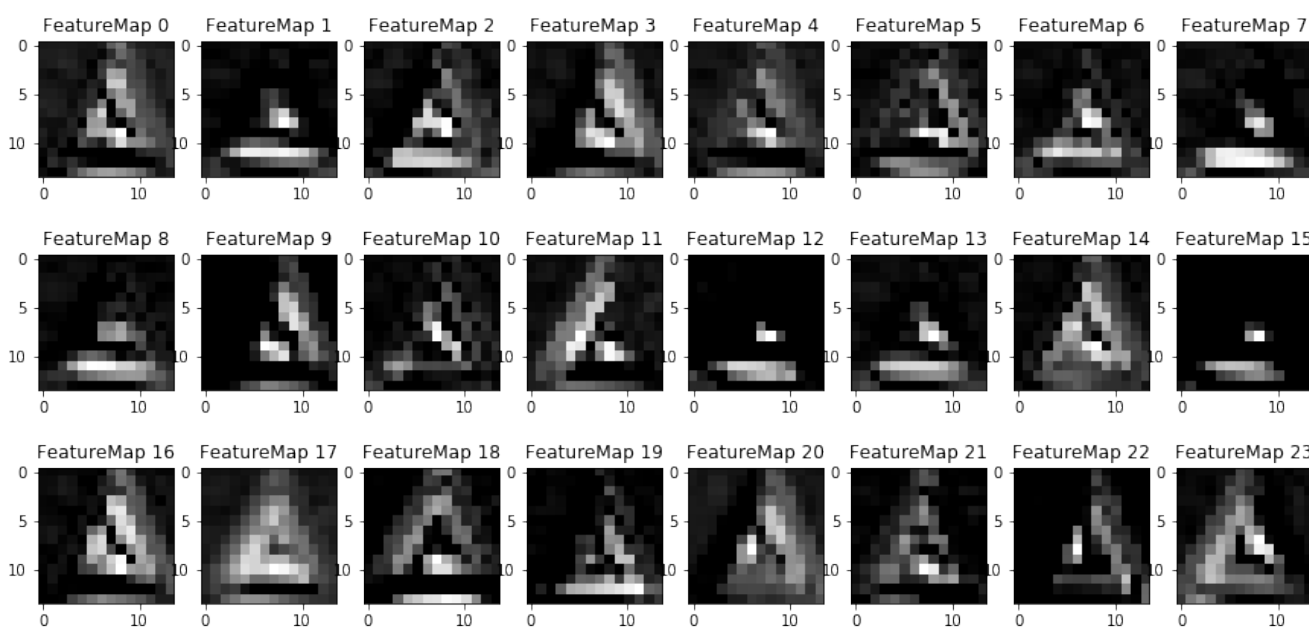
This analysis emphasizes how poorly the model performed in estimating the label for an internet image. Not only did the model not choose the correct label for any of the images, its top 5 guesses for a given image were also wrong (except for image #1).

Section 5: Visualize layers of the Neural Network

Below is the sign (Wild animals crossing) which I sent through the neural network in order to look at the salient features the network used to perform its classification task.




The following images show the features looked for by convolutional layer #1. For the most part it noticed large sections of triangles because the sign has a triangular shape.



Maps of the features looked at by convolutional layer #1 when the sign above is passed through the network.

The following images show the features looked for by convolutional layer #2 and these features appear to be more primitive kinds of shapes that the network looked for in order to classify the sign: horizontal lines in map 38, 41, 47 and diagonal lines in map 5, 7, 25 and 40.



Maps of the features looked at by convolutional layer #2 when a sign is passed through the network.