Prof. Bernt Schiele, Dr. Mario Fritz
<{schiele,mfritz}@mpi-inf.mpg.de>

Rakshith Shetty, Yang He
<{rshetty, yang}@mpi-inf-mpg.de>

Exercise 2: Deep Neural Networks and Backpropagation

Deep neural networks have shown staggering performances in various learning tasks, including computer vision, natural language processing, and sound processing. They have made the model designing more flexible by enabling end-to-end training.

In this exercise, we get to have a first hands-on experience with neural network training. Many frameworks (*e.g.* PyTorch, Tensorflow, Caffe) allow easy usage of deep neural networks without precise knowledge on the inner workings of backpropagation and gradient descent algorithms. While these are very useful tools, it is important to get a good understanding of how to implement basic network training from scratch, before using this libraries to speed up the process. For this purpose we will implement a simple two-layer neural network and its training algorithm based on back-propagation using only basic matrix operations in questions 1 to 3. In question 4, we will use a popular deep learning library, PyTorch, to do the same and understand the advantanges offered in using such tools.

As a benchmark to test our models, we consider an image classification task using the widely used CIFAR-10 dataset. This dataset consists of 50000 training images of 32x32 resolution with 10 object classes, namely airplanes, cars, birds, cats, deer, dogs, frogs, horses, ships, and trucks. The task is to code and train a parametrised model for classifying those images. This involves

- Implementing the feedforward model (Question 1).

- Implementing the backpropagation algorithm (gradient computation) (Question 2).

- Training the model using stochastic gradient descent and improving the model training with better hyperparameters (Question 3).

- Using the PyTorch Library to implement the above and experiment with deeper networks (Question 4).

A note on notation: Throughout the exercise, notation v_i is used to denote the i^{th} element of vecotr v.

Please untar the handed out `ex2.tar.gz`. Questions 1-3 are based on the script `ex2_FCnet.py` and question 4 is based on the script `ex2_pytorch.py`. To download the CIFAR-10 dataset please execute the script `datasets/get_datasets.sh`.

Questions indicated by *report* should be answered in a separate report file - preferrably a pdf format file. Derivations are allowed to be handwritten (and scanned), but LaTeXgenerated documents would be preferred.

The completed exercise should be handed in in a `.tar.gz` format which compresses all the code + the report.

Question 1: Implementing the feedforward model (10 points)

In this question we will implement a two-layered a neural network architecture as well as the loss function to train it. Starting from the main file `ex2_FCnet.py`, complete the required code in the `two_layernet.py` to complete this question. Refer to the comments in the code to the exact places where you need to fill in the code.

Model architecture Our architecture is shown in Fig.1. It has an input layer, and two model layers – a hidden and an output layer. We start with randomly generated toy inputs of 4 dimensions and number of classes $K = 3$ to build our model in q1 and q2 and in q3 use images from CIFAR-10 dataset to test our model on a real-world task. Hence input layer is 4 dimensional for now.

In the hidden layer, there are 10 units. The input layer and the hidden layer are connected via linear weighting matrix $W^{(1)} \in \mathbb{R}^{10 \times 4}$ and the bias term $b^{(1)} \in \mathbb{R}^{10}$. The parameters $W^{(1)}$ and $b^{(1)}$ are to be learnt later on. A linear operation is performed, $W^{(1)}x + b^{(1)}$, resulting in a 10 dimensional vector $z^{(2)}$. It is then followed by a

Figure 1: Visualisation of the two layer fully connected network, used in Q1-Q3

relu non-linear activation ϕ, applied element-wise on each unit, resulting in the activations $a^{(2)} = \phi(z^{(2)})$. Relu function has the following form:

$$\phi(u) = \begin{cases} u, & \text{if } u \geq 0 \\ 0, & \text{if } u < 0 \end{cases} \tag{1}$$

A similar linear operation is performed on $a^{(2)}$, resulting in $z^{(3)} = W^{(2)}a^{(2)} + b^{(2)}$, where $W^{(2)} \in \mathbb{R}^{3 \times 10}$ and $b^{(2)} \in \mathbb{R}^3$; it is followed by the softmax activation to result in $a^{(3)} = \psi(z^{(3)})$. The softmax function is defined by:

$$\psi(u)_i = \frac{\exp^{u_i}}{\sum_j \exp^{u_j}} \tag{2}$$

The final functional form of our model is thus defined by

$$a^{(1)} = x \tag{3}$$

$$z^{(2)} = W^{(1)}a^{(1)} + b^{(1)} \tag{4}$$

$$a^{(2)} = \phi(z^{(2)}) \tag{5}$$

$$z^{(3)} = W^{(2)}a^{(2)} + b^{(2)} \tag{6}$$

$$f_\theta(x) := a^{(3)} = \psi(z^{(3)}) \tag{7}$$

which takes a flattened 4 dimensional vector as input and outputs a 3 dimensional vector, each entry in the output $f_k(x)$ representing the probability of image x corresponding to the class k. We summarily indicate all the network parameters by $\theta = (W^{(1)}, b^{(1)}, W^{(2)}, b^{(2)})$.

Implementation We are now ready to implement the feedforward neural network.

a) Implement the code in `two_layenet.py` for the feedforward model. You are required to implement Eq.3 to 7. Verify that the scores you generate for the toy inputs match the correct scores given in the `ex2_FCnet.py`. (4 points)

b) We later guide the neural network parameters $\theta = (W^{(1)}, b^{(1)}, W^{(2)}, b^{(2)})$ to fit to the given data and label pairs. We do so by minimising the loss function. A popular choice of the loss function for training neural network for multi-class classification is the cross-entropy loss. For a single input sample x_i, with label y_i, the loss function is defined as :

$$J(\theta, x_i, y_i) = -\log P(Y = y_i, X = x_i) \tag{8}$$

$$= -\log f_\theta(x_i)_{y_i} \tag{9}$$

$$= -\log \psi(z^{(3)})_{y_i} \tag{10}$$

$$J(\theta, x_i, y_i) = -\log \left[\frac{\exp^{z_{y_i}^{(3)}}}{\sum_j^K \exp^{z_j^{(3)}}} \right] \tag{11}$$

Averaging over the whole training set, we get

$$J(\theta, \{x_i, y_i\}_{i=1}^N) = \frac{1}{N} \sum_{i=1}^N -log \left[\frac{\exp^{z_{y_i}^{(3)}}}{\sum_j \exp^{z_j^{(3)}}} \right] \tag{12}$$

where K is the number of classes. Note that if the model has perfectly fitted to the data (*i.e.* $f_\theta^k(x_i) = 1$ whenever x_i belongs to class k and 0 otherwise), then J attains the minimum of 0.

Apart from trying to correctly predict the lable, we have to prevent overfitting the model to the current training data. This is done by encoding our prior belief that the correct model should be simple (Occam's razor); we add an L_2 regularisation term over the model parameters θ. Specifically, the loss function is defined by:

$$\tilde{J}(\theta) = \frac{1}{N} \sum_{i=1}^N -log \left[\frac{\exp^{z_{y_i}^{(3)}}}{\sum_j \exp^{z_j^{(3)}}} \right] + \lambda \left(||W^{(1)}||_2^2 + ||W^{(2)}||_2^2 \right) \tag{13}$$

where $|| \cdot ||_2^2$ is the squared L_2 norm. For example,

$$||W^{(1)}||_2^2 = \sum_{p=1}^{10} \sum_{q=1}^{4} W_{pq}^{(1)2} \tag{14}$$

By changing the value of λ it is possible to give weights to your prior belief on the degree of simplicity (regularity) of the true model.

Implement the final loss function in `two_layernet.py` and let it return the loss value. Verify the code by running and matching the output cost 1.30378789133. (4 points)

c) To be able to train the above model on large datsets, with larger layer widths, the code has to be very efficient. To do this you should avoid using any python *for* loops in the forward pass and instead use matrix/ vector multiplication routines in the numpy library. If you have written the code of parts (a) and (b) using loops, convert it to vectorized version using numpy operations (2 points).

Question 2: Backpropagation (15 points)

We train the model by solving

$$\min_\theta \tilde{J}(\theta) \tag{15}$$

via stochastic gradient descent. We therefore need an efficient computation of the gradients $\nabla_\theta \tilde{J}(\theta)$. We use backpropagation of top layer error signals to the parameters θ at different layers.

In this question, you will be required to implement the backpropagation algorithm yourself from a pseudocode. We will give a high-level description of what is happening at each line.

For those who are interested in the robust derivation of the algorithm, we include the optional exercise on the derivation of backpropagation algorithm. A prior knowledge on standard vector calculus including the chain rule would be helpful.

Backpropagation   Backpropagation algorithm is simply a sequential application of chain rule. It is applicable to any (sub-) differentiable model that is a composition of simple building blocks. In this exercise, we focus on the architecture with stacked layers of linear transformation + relu non-linear activation.

The intuition behind backpropagation algorithm is as follows. Given a training example (x, y), we first run the feedforward to compute all the activations throughout the network, including the output value of the model $f_\theta(x)$ and the loss J. Then, for each parameter in the model we want to compute the effect that parameter has on the loss. This is done by computing the derivatives of the loss w.r.t each model parameter.

Backpropagation algorithm is performed from the top of the network (loss layer) towards the bottom. It sequentially computes the gradient of the loss function with respect to each layer activations and parameters.

Let's start by deriving the gradients of the un-regularized loss function w.r.t final layer activations $z^{(3)}$. We will then use this in the chain rule to compute analytical experssions for gradients of all the model parameters.

a) Verify that the loss function defined in Eq.12 has the gradient w.r.t $z^{(3)}$ as below.

$$\frac{\partial J}{\partial z^{(3)}} \left(\{x_i, y_i\}_{i=1}^N \right) = \frac{1}{N} \left(\psi(z^{(3)}) - \Delta \right) \tag{16}$$

where Δ is a matrix of $N \times K$ dimensions with

$$
\Delta_{ij} = \begin{cases} 1, & \text{if } y_i = j \\ 0, & \text{otherwise} \end{cases} \tag{17}
$$

(*report*, 2 points)

b) To compute the effect of the weight matrix $W^{(2)}$ on the loss in Eq.12 incurred by the network, we compute the partial derivatives of the loss function with respect to W^2. This is done by applying the chain rule. Verify that the partial derivative of the loss w.r.t $W^{(2)}$ is

$$
\frac{\partial J}{\partial W^{(2)}} \left(\{x_i, y_i\}_{i=1}^N \right) = \frac{\partial J}{\partial z^{(3)}} \cdot \frac{\partial z^{(3)}}{\partial W^{(2)}} \tag{18}
$$

$$
= \frac{1}{N} (\psi(z^{(3)}) - \Delta) a^{(2)'} \tag{19}
$$

Similary, verify that the regularized loss in Eq.13 has the derivatives

$$
\frac{\partial \tilde{J}}{\partial W^{(2)}} = \frac{1}{N} (\psi(z^{(3)}) - \Delta) a^{(2)'} + 2\lambda W^{(2)} \tag{20}
$$

(*report*, 2 point)

c) We can repeatedly apply chain rule as discussed above to obtain the derivatives of the loss with respect to all the parameters of the model $\theta = (W^{(1)}, b^{(1)}, W^{(2)}, b^{(2)})$. Dervive the expressions for the derivatives of the regularized loss in Eq.13 w.r.t $W^{(1)}$, $b^{(1)}$, $b^{(2)}$ now. (*report*, 6 points)

d) Using the expressions you obtained for the derivatives of the loss w.r.t model parameters, implement the back-propogation algorithm in the file `two_layernet.py`. Run the `ex2_FCnet.py` and verify that the gradients you obtained are correct using numerical gradients (already implemented in the code). The maximum relative error between the gradients you compute and the numerical gradients should be less than 1e-8 for all parameters. (5 points)

Question 3: Stochastic gradient descent training (10 points)

We have implemented the backpropagation algorithm for computing the parameter gradients and have verified that it indeed gives the correct gradient. We are now ready to train the network. We solve Eq.15 with the stochastic gradient descent.

Stochastic gradient descent (SGD) Typically neural networks are large and are trained with millions of data points. It is thus often infeasible to compute the gradient $\nabla_\theta \tilde{J}(\theta)$ that requires the accumulation of the gradient over the entire training set. Stochastic gradient descent addresses this problem by simply accumulating the gradient over a small random subset of the training samples (minibatch) at each iteration. Specifically, the algorithm is as follows

---

Data: Training data $\{(x_i, y_i)\}_{i=1,\cdots,N}$, initial network parameter $\theta^{(0)}$, regularisation hyperparameter λ, learning rate α, batch size B, iteration limit T
Result: Trained parameter $\theta^{(T)}$
1 **for** $t = 1, \cdots, T$ **do**
2 $\quad \{(X_j', y_j')\}_{j=1}^B \leftarrow$ a random subset of the original training set $\{(X_i, y_i)\}_{i=1}^N$;
3 $\quad v \leftarrow -\alpha \nabla_\theta \tilde{J}(\theta^{(t-1)}, \{(X_j', y_j')\}_{j=1}^B)$;
4 $\quad \theta^{(t)} \leftarrow \theta^{(t-1)} + v$;
5 **end**

Algorithm 1: Stochastic gradient descent with momentum

where the gradient $\nabla_\theta \tilde{J}(\theta, \{(X_j', y_j')\}_{j=1}^B)$ is computed only on the current randomly sampled batch.

Intuitively, $v = -\nabla_\theta \tilde{J}(\theta^{(t-1)})$ gives the direction to which the loss $\tilde{J}$ decreases the most (locally), and therefore we follow that direction by updating the parameters towards that direction $\theta^{(t)} = \theta^{(t-1)} + v$.

a) Implement the stochastic gradient descent algorithm in `two_layernet.py` and run the training on the toy data. Your model should be able to obtain loss ¡= 0.02 on the training set and the training curve should look similar to the one shown in figure 2. (3 points)
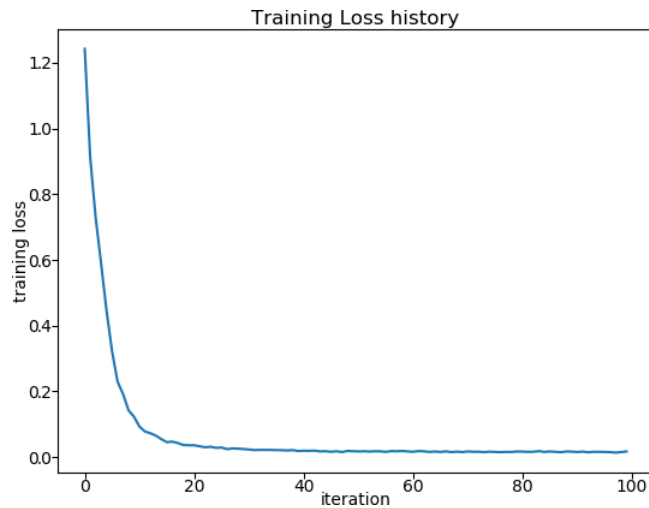
Figure 2: Example training curve on the toy dataset.

Figure 3: Example images from the CIFAR-10 dataset

b) We are now ready to train our model on real image dataset. For this we will use the CIFAR-10 dataset. Since the images are of size 32×32 pixels with 3 color channels, this gives us 3072 input layer units, represented by a vector $x \in \mathbb{R}^{3072}$. See figure 3 for example images from the dataset. The code to load the data and train the model is provided with some default hyper-parameters in `ex2_FCnet.py`. With default hyperparametres, if previous questions have been done correctly, you should get validation set accuracy of about 29%. This is very poor. Your task is to debug the model training and come up with beter hyper-parameters to improve the performance on the validation set. Visualize the training and validation performance curves to help with this analysis. There are several pointers provided in the comments in the `ex2_FCnet.py` to help you understand why the network might be underperforming (Line 224-250). Once you have tuned your hyper parameters, and get validation accuracy greater than 48% run your best model on the test set once and report the performance.

Question 4: Implement multi-layer perceptron using PyTorch library (10 points)

So far we have implemented a two-layer network by explicitly writing down the expressions for forward, backward computations and training algorithms using simple matrix multiplication primitives from numpy library.

However there are many libraries available designed make experimenting with neural networks faster, by abstracting away the details into re-usable modules. One such popular opensource library is PyTorch (https://pytorch.org/). In this final question we will use PyTorch library to implement the same two-layer network we did before and

train it on the Cifar-10 dataset. However, extending a two-layer network to a three or four layered one is a matter of changing two-three lines of code using PyTorch. We will take advantage of this to experiment with deeper networks to improve the performance on the CIFAR-10 classification. This question is based on the file `ex2_pytorch.py`

To install the pytorch library follow the instruction in https://pytorch.org/get-started/locally/ . If you have access to a Graphics Processing Unit (GPU), you can install the gpu verison and run the exercise on GPU for faster run times. If not, you can install the cpu version (select cuda version None) and run on the cpu. Having gpu access is not necessary to complete the exercise. There are good tutorials for getting started with pytorch on their website (https://pytorch.org/tutorials/).

a) Complete the code to implement a multi-layer perceptron network in the class *MultiLayerPerceptron* in `ex2_pytorch.py`. This includes instantiating the required layers from `torch.nn` and writing the code for forward pass. Intially you should write the code for the same two-layer network we have seen before. (3 points)

b) Complete the code to train the network. Make use of the loss function `torch.nn.CrossEntropyLoss` to compute the loss and `loss.backward()` to compute the gradients. Once gradients are computed, `optimizer.step()` can be invoked to update the model. Your should be able to achieve similar performance ($> 48\%$ accuracy on the validation set) as in Q3. Report the final validation accuracy you achieve with a two-layer network. (3 points)

c) Now that you can train the two layer network to achieve reasonable performance, try increasing the network depth to see if you can improve the performance. Experiment with networks of atleast 2, 3, 4, and 5 layers, of your chosen configuration. Report the training and validataion accuracies for these models and discuss your observations. Run the evaluation on the test set with your best model and report the test accuracy. (report, 4 points)