

Due: Wednesday, April 21 at 11:59 pm

Deliverables:

1. Submit your predictions for the test sets to Kaggle as early as possible. Include your Kaggle scores in your write-up (see below). The Kaggle competition for this assignment can be found at
 - <https://www.kaggle.com/c/spring21-cs189-hw6-cifar10>
2. The written portion:
 - Submit a PDF of your homework, **with an appendix listing all your code**, to the Gradescope assignment titled “Homework 6 Write-Up”. Please see section 3.3 for an easy way to gather all your code for the submission (you are **not** required to use it, but we would strongly recommend using it).
 - In addition, please include, as your solutions to each coding problem, the specific subset of code relevant to that part of the problem. Whenever we say “include code”, that means you can either include a screenshot of your code, or typeset your code in your submission (using markdown or \LaTeX).
 - You may typeset your homework in LaTeX or Word (submit PDF format, **not** .doc/.docx format) or submit neatly handwritten and scanned solutions. **Please start each question on a new page.**
 - If there are graphs, include those graphs in the correct sections. **Do not** put them in an appendix. We need each solution to be self-contained on pages of its own.
 - In your write-up, please state with whom you worked on the homework.
 - In your write-up, please copy the following statement and sign your signature next to it. (Mac Preview and FoxIt PDF Reader, among others, have tools to let you sign a PDF file.) We want to make it *extra* clear so that no one inadvertently cheats.

“I certify that all solutions are entirely in my own words and that I have not looked at another student’s solutions. I have given credit to all external sources I consulted.”
3. Submit all the code needed to reproduce your results to the Gradescope assignment entitled “Homework 6 Code”. Yes, you must submit your code twice: in your PDF write-up following the directions as described above so the readers can easily read it, and once in compilable/interpretable form so the readers can easily run it. Do **NOT** include any data files we provided. Please include a short file named README listing your name, student ID, and instructions on how to reproduce your results. Please take care that your code doesn’t take up inordinate amounts of time or memory. If your code cannot be executed, your solution cannot be verified.

1 Honor Code

Declare and sign the following statement:

“I certify that all solutions in this document are entirely my own and that I have not looked at anyone else’s solution. I have given credit to all external sources I consulted.”

Signature : _____

While discussions are encouraged, *everything* in your solution must be your (and only your) creation. Furthermore, all external material (i.e., *anything* outside lectures and assigned readings, including figures and pictures) should be cited properly. We wish to remind you that consequences of academic misconduct are *particularly severe*!

2 Background

This section will provide a background on neural-networks that is designed to help you complete the assignment. There are no questions in this part. The questions for the homework begin in section 4.

2.1 Neural Networks

Many of the most exciting recent breakthroughs in machine learning have come from “deep” (read: many-layer) neural networks, such as the [deep reinforcement learning algorithm](#) that learned to play Atari from pixels, or the [GPT-2](#) model, which generates text that is nearly indistinguishable from human-generated text.

Neural network libraries such as Tensorflow and PyTorch have made training complicated neural network architectures very easy. However, we want to emphasize that neural networks begin with fundamentally simple models that are just a few steps removed from basic logistic regression. In this assignment, you will build two fundamental types of neural network models, all in plain numpy: a **feed-forward fully-connected network**, and a **convolutional neural network**. We will start with the essential elements and then build up in complexity.

A neural network model is defined by the following.

- An **architecture** defining the flow of information between computational layers. This defines the composition of functions that the network performs from input to output.
- A **cost function** (e.g. cross-entropy or mean squared error).
- An **optimization algorithm** (e.g. stochastic gradient descent with backpropagation).
- A set of **hyperparameters** (e.g. learning rate, batch size, etc.).

Each *layer* is defined by the following components.

- A **parameterized function** that defines the layer’s map from input to output (e.g. $f(x) = \sigma(Wx + b)$).
- An **activation function** σ (e.g. ReLU, sigmoid, etc.).
- A set of **parameters** (e.g. weights and biases).

Neural networks are commonly used for supervised learning problems, where we have a set of inputs and a set of labels, and we want to learn the function that maps inputs to labels. To learn this function, we need

to update the parameters of the network (the weights and biases). We do this using **mini-batch gradient descent**. In order to compute the gradients for gradient descent, we will use a dynamic programming algorithm called **backpropagation**.

In the backpropagation algorithm, we first compute what is called a “forward pass” of the network. In the forward pass, we send a mini-batch of input data (e.g. 50 datapoints) through the network. The result is a set of outputs, which we use to compute our loss function. We then take the derivatives of this loss with respect to the parameters of each layer, starting with the output of the network and using the chain rule to propagate backwards through the layers. This is called the “backward pass.” During the backward pass we compute the derivatives of the loss function with respect to each of the model parameters, starting from the last layer and “propagating” the information from the loss function backwards through the network. This lets us calculate derivatives with respect to all the parameters of our network while letting us avoid computing the same derivatives multiple times.

To summarize, training a neural network involves three steps.

1. Forward propagation of inputs.
2. Computing the cost.
3. Backpropagation and parameter updates.

2.2 Batching

When Building neural networks, we have to carefully consider the data. In homework 4, you coded both batch gradient descent and stochastic gradient descent for logistic regression. For the stochastic version, where only a single data point was used, the form of derivatives used in gradient descent were different than those of batch gradient descent. Neural-networks always operate on mini-batches, or subsets of the data matrix. This is because operating on all the data at once would be impossible given current compute and actually bad for optimization, while operating on a single data point would be far too noisy. Thus, every step of your neural network must be defined to operate on batches of data. For example, the input to a fully connected neural network would be a matrix of shape (B, d) where B is the batch size and d is the number of features. The input to a convolutional neural network would be a tensor of shape (B, H, W, C) where B is the batch size, H is the height of the image, W is the width of the image, and C is the number of channels in the image (3 for RGB).

Because you are writing the gradient descent algorithm to work on mini-batches, all of your derivations must work for mini-batches as well. This is important to keep in mind as you complete this assignment, and many of the derivations may be different than those you have seen in class. Thinking in terms of mini-batches often changes the shapes and operations you do. **Your derivations must be for mini-batches and cannot use for loops to iterate over individual data points.** This is one aspect of the assignment that students often find to be very difficult.

2.3 Feed-Forward, Fully-Connected Neural Networks

A feed-forward, fully-connected neural network layer performs an [affine transformation](#) of an input, followed by a nonlinear activation function. We will use the following notation when defining fully-connected layers, with superscripts surrounded by brackets indexing layers and subscripts indexing the vector/matrix elements.

- x : A single data vector, of shape $1 \times d$, where d is the number of features.




Figure 1: A 3-layer fully-connected neural network.




Figure 2: A single fully-connected neuron.

- y : A single label vector, of shape $1 \times k$, where k is the number of classes (for a classification problem), or the number of output features (for a regression problem).
- $n^{[l]}$: The number of neurons in layer l .
- $W^{[l]}$: A matrix of weights connecting layer $l - 1$ with layer l , of shape $n^{[l-1]} \times n^{[l]}$. At layer 0, it is shape $d \times n^{[l]}$.
- $b^{[l]}$: The bias vector for layer l , of shape $1 \times n^{[l]}$.
- $h^{[l]}$: The output of layer l . This is a vector of shape $1 \times n^{[l]}$.
- $\sigma^{[l]}(\cdot)$: The nonlinear “activation function” applied at layer l .

A fully-connected layer l is a function

$$\phi^{[l]}(h^{[l-1]}) = \sigma^{[l]}(h^{[l-1]}W^{[l]} + b^{[l]}) = h^{[l]}.$$

At layer 0, $h^{[l-1]}$ is simply the data vector x . We will use the term $z^{[l]} = h^{[l-1]}W^{[l]} + b^{[l]}$ as shorthand for the intermediate result within layer l before applying the activation function σ . Each layer is computed sequentially and the output of one layer is used as the input to the next. A neural network is thus a *composition of functions*. We want to find the parameters of the function that takes us from our input examples x to our labels y .

When used for classification, feed forward networks accept inputs as d -dimensional vectors, and the labels will be k -dimensional “one-hot” vectors, where k is the number of classes. A one-hot vector is a binary vector whose elements are computed according to the following function:

$$y_i = \begin{cases} 1 & x \in \text{class } i, \\ 0 & \text{otherwise.} \end{cases}$$

For example, for a classification problem with 3 classes, the label encoding an example from class 2 (zero-indexing) would be: $[0, 0, 1]$.

2.4 Convolutional Neural Networks

We will use the following notation when defining a convolutional neural network layer.

- X : A single image tensor, of shape $1 \times d_1 \times d_2 \times c$, where d_1 and d_2 are the spatial dimensions, and c is the number of channels.
- y : A single label vector, of shape $1 \times k$, where k is the number of classes.
- $n^{[l]}$: The number of neurons in layer l .
- $(k_1, k_2)^{[l]}$: The size of the spatial dimensions of the filters in layer l . Also referred to as the kernel¹ size.
- $W^{[l]}$: The tensor of filters convolved at layer l . This tensor has shape $k_1 \times k_2 \times n^{[l-1]} \times n^{[l]}$.
- $b^{[l]}$: The bias vector for layer l , of shape $1 \times n^{[l]}$.
- $H^{[l]}$: The output of layer l . This is a tensor of shape $1 \times r_1 \times r_2 \times n^{[l]}$, where (r_1, r_2) is the shape of output of the convolution operation. Below we will discuss how to calculate this.
- $\sigma^{[l]}(\cdot)$: The nonlinear “activation function” applied at layer l .

In a convolutional layer, each filter is convolved with the input image, across every image channel. This operation is, essentially, a sliding sum of element-wise products. Figure 3 gives a visual example. To compute a single element in the intermediate output Z , for a single neuron n and a single channel c in the input X , we compute

$$Z[d_1, d_2, n] = (X * W)[d_1, d_2, n] = \sum_i \sum_j \sum_c W[i, j, c, n] X[d_1 + i, d_2 + j, c] + b[n].$$

Please note that the formula above is the **cross-correlation** formula from signal processing and NOT the convolution formula. Nevertheless this is what ML people call convolution and so will we. It actually makes sense to use cross-correlation instead of using convolution because the former can be interpreted as producing an output which is higher at locations where the image has the pattern in the filter and low elsewhere. Furthermore, convolution is the same as cross-correlation with a flipped filter, and we learn the filters, so it should not make any difference operationally whether you implement convolution or cross-correlation. However, in order to pass the tests, you must implement **cross-correlation** and call that convolution because that’s how we do it in ML-land.

In this equation, we drop the layer superscripts for clarity, and index elements of the matrices in brackets. The output of this operation is what we call a “feature map,” which essentially captures the strength of each filter at every region in the image. In the equation above, we slide the filter over the image in increments of one pixel. We can choose to take a larger steps instead. The size of the step taken in the convolution operation is referred to as the *stride*.

The output of the convolutional layer is

$$H^{[l]} = \sigma^{[l]}(Z^{[l]}).$$

¹“Kernel” is an overloaded word. In the context of convolutional networks, the convolutional filter is also called the “convolutional kernel”. This has no relationship with the “kernel” in the kernel trick and kernel methods. The convolutional kernel is also referred to as a “mask” in lecture.




Figure 3: Figure showing an example of one convolution.




Figure 4: Figure showing an example of a max pooling layer with a kernel size of 2 and stride of 2.

A pooling layer is used to downsample the input feature maps. It takes an input array of shape $1 \times d_1 \times d_2 \times c$ and outputs an array of shape $1 \times r_1 \times r_2 \times c$. Note that it does **not** change the number of channels, but typically reduces the number of spatial dimensions, i.e., $r_1 < d_1$ and $r_2 < d_2$. In order to do this, we have a **kernel** of shape $k_1 \times k_2$ and a stride s . For each channel, we take either the max or the average of all the points in the window of size $k_1 \times k_2$. Then we slide the window by s pixels and repeat until we have performed this operation over the entire input image. This is illustrated in Figure 4. When the operation performed over each sliding window is max, it is called **max pooling**, whereas when the operation is averaging, then it is called **average pooling**. Using similar notation as above, the function computed by a max pooling layer is:

$$Z[r_1, r_2, c] = \text{MaxPool}(X)[r_1, r_2, c] = \max(X[r_1 s : r_1 s + k_1, r_2 s : r_2 s + k_2, c]).$$

Replacing the max function with the average function, we get an average pooling layer. This function is abstracted away as `pool_fn` in the code. Note that `pool_fn` takes an array as input and outputs a single float. The notation on the right hand side should be read as array slicing as in `numpy`.

Traditional CNNs operate on images by combining convolutional layers with pooling layers to progressively "shrink" the spatial size of the input until it is small enough to be fed to fully-connected network layers for classification.

3 The Neural Nets Package

3.1 Structure

We have provided a modularized codebase for constructing neural networks. The codebase has the following structure.

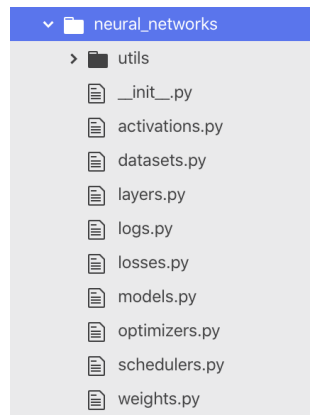


Figure 5: The structure of the starter codebase.

As you can see, the modules in the codebase reflect the structure outlined above. Different losses, activations, layers, optimizers, hyperparameters, and neural network architectures can be combined to yield different architectures.

In the codebase we have provided, each layer is an object with a few relevant attributes.

- **parameters:** An `OrderedDict` containing the weights and biases of the layer.
- **gradients:** An `OrderedDict` containing the derivatives of the loss with respect to the weights and biases of the layer, with the same keys as **parameters**.

- **cache:** An `OrderedDict` containing intermediate quantities calculated in the forward pass that are useful for the backward pass.
- **activation:** An `Activation` instance that is the activation function applied by this layer.
- **n_in:** The number of input units (input channels in CNN).
- **n_out:** The number of output units (output channels in CNN).

You will pass the layer a parameter that selects an activation function from those defined in `activations.py`. This will be stored as an attribute of the layer, which can be called as `layer.activation()`. The forward and backward passes of the layer are defined by the following methods.

- **forward** This method takes as input the output X from the previous layer (or input data). This method computes the function $\phi(\cdot)$ from above, combining the input with the weights W and bias b that are stored as attributes. It returns an output `out` and saves the intermediate value Z to the `cache` attribute, as it is needed to compute gradients in the backward pass.

```
def forward(self, X: np.ndarray) -> np.ndarray:
    """Forward pass: multiply by a weight matrix, add a bias, apply activation.
    Also, store all necessary intermediate results in the `cache` dictionary
    to be able to compute the backward pass.
    """
    ...
```

- **backward** This method takes the gradient of the downstream loss as input and uses the cached values to compute gradients with respect to its inputs and weights. It returns the gradient of the loss with respect to the input of the layer.

```
def backward(self, dLdY: np.ndarray) -> np.ndarray:
    """Backward pass for fully connected layer.
    Compute the gradients of the loss with respect to:
    1. the weights of this layer (mutate the `gradients` dictionary)
    2. the bias of this layer (mutate the `gradients` dictionary)
    3. the input of this layer (return this)
    """
    ...
```

Each activation function has a similar (but simpler) structure:

```
class Linear(Activation):
    def __init__(self):
        super().__init__()

    def forward(self, Z: np.ndarray) -> np.ndarray:
        """Forward pass for  $f(z) = z$ ."""
        return Z

    def backward(self, Z: np.ndarray, dY: np.ndarray) -> np.ndarray:
        """Backward pass for  $f(z) = z$ ."""
        return dY
```


3.2 Testing

We are using the `unittest` module built into Python for testing. Note that **the tests are not comprehensive**, so it is possible that your code passes all the tests but is still broken and gives you poor performance. We test the `forward` and `backward` methods of all the layers that you implement, along with a few that you don't have to implement (which should already be passing their tests). We also test the weight initialization of the `FullyConnected` layer so that you have some guidance when you implement a layer for the first time. When you are done with the homework, **you must make sure that your code passes all the tests** except the tests for the L2 loss function and the Sigmoid activation function. You should have 21/25 tests passing at least. (You can implement the L2 loss and Sigmoid activation function if you want to make those tests pass.)

To run all the tests, make sure that you are in the root directory of the project (and **not** in either of the `neural_networks/` and `tests/` directories). Then run the following:

```
python -m unittest -v
```

25 tests will be run. Before you implement anything, 17 tests should fail and 8 should pass.

The tests are located in the `tests/` directory. **Please do not modify the pre-existing tests** because they rely on data generated by a correct implementation, and any modifications can cause the tests to fail. The data is located in the `tests/data` directory. However, if you wish to write your own tests in addition to the ones that are already provided, then please refer to [this tutorial](#) and [the official documentation](#). Note that we don't require you to write your own tests as part of this homework.

The tests are organized into separate classes that mirror the structure of the classes that they test. For example, the `forward` method of the `Conv2D` class located in `neural_networks/layers.py` is tested by the `test_forward` method of the `TestConv2D` class located in `tests/test_layers.py`. If you want to run that specific test, you can use the following command **root directory of the project**:

```
python -m unittest -v tests.test_layers.TestConv2D.test_forward
```

If you instead want to run all the tests defined in the `TestConv2D` class, you should run:

```
python -m unittest -v tests.test_layers.TestConv2D
```

And if you want to run all the tests in `test_layers.py`, then you should run:

```
python -m unittest -v tests.test_layers
```

We have also included a Jupyter Notebook named `check_gradients.ipynb`, which only contains gradient checks for your layers, which you can use for debugging your layers' gradients. This can provide more debuggability than simply running the provided unit tests. Note that we do **not** require you to provide us the output of those tests, since gradients are implicitly tested in the `test_backward` methods of the testing classes.

3.3 Submission

Please run

```
python3 generate_submission.py --help
```

for instructions on how to use the script to extract your implementations for the submission.

The script `generate_submission.py` extracts all the code from your implementations and produces either a \LaTeX or markdown file containing your code implementations, when given the flag `--format latex` and `--format markdown` respectively. The generated document contains your functions in separate sections for activation functions, layers, losses, and the model. These can be sections, subsections, or subsubsections depending on whether you supply the flag `--heading_level 1`, `--heading_level 2`, or `--heading_level 3`.

For example, if you want \LaTeX output with each part (activations, layers, losses, and the model) in a different subsection, with the output saved to `submission.tex`, you would run the following:

```
python3 generate_submission.py --format latex --heading_level 2 --output submission.tex
```

whereas if you want markdown output with each part (activations, layers, losses, and the model) in a different subsubsection, with the output saved to `submission.md`, you would run the following:

```
python3 generate_submission.py --format markdown --heading_level 3 --output submission.md
```

We would suggest running these commands to see exactly what they do.

The markdown document will compile by itself, but you would most likely want to create a markdown cell in a Jupyter Notebook and copy-paste the generated markdown into that cell. That should work seamlessly provided you can already compile Jupyter Notebooks into PDFs.

Note that the \LaTeX document will **not** compile by itself. It is meant to generate code that you can then `\input{}` into your \LaTeX document.

Feel free to play around with the script if you want to. Notably, if you change some function which is not in the `student_implementations` list, then you could add that function to the `student_implementations` list to have the script automatically gather your code from that function (but we think that most students will **not** have to do this).

Questions

The background section of the homework has ended. The following sections contain the questions you must complete.

4 Basic Network Layers

In this question you will implement the layers needed for basic classification neural networks. For each part, you will be asked to 1) derive the gradients, 2) write the matching code, 3) pass the tests.

Keep in mind that your solutions to all layers **must operate on mini-batches of data** and should not use for-loops to iterate over the data points in a batch.

4.1 ReLU

First, you will implement the ReLU activation function in `activations.py`.

$$\sigma_{\text{ReLU}}(\gamma) = \begin{cases} 0 & \gamma < 0, \\ \gamma & \text{otherwise.} \end{cases}$$

Note that the activation function is applied element-wise to a vector input.

Instructions

1. First, derive the gradient of the downstream loss with respect to the input of the ReLU activation function, Z . **You must arrive at a solution for batched input X .**
2. Next, implement the forward and backward passes of the ReLU activation in the script `activations.py`. Include your code in your writeup (either a screenshot or typesetting is fine). **Do not iterate over training examples, use batched operations.**
3. Include the output of running `python -m unittest -v tests.test_activations.TestReLU`. Make sure that the tests are passing.

4.2 Fully-Connected Layer

Now you will implement the forward and backward passes for the fully-connected layer in the `layers.py` script. The code is marked with `YOUR CODE HERE` statements indicating what to implement and where. Please read the docstrings and the function signatures too. Write the fully-connected layer for a general input h that contains a mini-batch of m examples with d features.

When implementing a new layer, it is important to manually verify correctness of the forward and backward passes. We have provided a Jupyter notebook `check_gradients.ipynb` for you to use to numerically check the gradients of your layer implementations. Simply run the cell corresponding to each layer. The printed errors should be very small, usually on the order of 10^{-8} or smaller.

Instructions

1. First, derive the gradients of the downstream loss L with respect to W and b in the fully-connected layer, $\frac{\partial L}{\partial W}$ and $\frac{\partial L}{\partial b}$. You will also need to take the derivative of the loss with respect to the input of the layer $\frac{\partial L}{\partial X}$, which will be passed to lower layers. **Again, you must assume must arrive at a solution for batched X and Z**

2. Implement the forward and backward passes of the fully-connected layer in `layers.py`. First, initialize the weights of the model using `_init_parameters`, which takes the shape of the design matrix as input and initializes the parameters, cache, and gradients of the layer. The `backward` method takes in an argument `dLdY`, the derivative of the loss with respect to the output of the layer, which is computed by higher layers and backpropagated. This should be incorporated into your gradient calculation. In your writeup, include the code you have implemented. **Do not iterate over training examples, use batched operations.**
3. Include the output of `python -m unittest -v tests.test_layers.TestFullyConnected`. Make sure that the tests are passing.

4.3 Softmax Activation

Next, we need to define an activation function for the output layer. The ReLU activation function returns continuous values that are (potentially) unbounded to the right. Since we are building a classifier, we want to return *probabilities* over classes. The softmax function has the desirable property that it outputs a probability distribution. That is, the softmax function squashes continuous values into the range $[0, 1]$ and normalizes the outputs so that they add up to 1. For this reason, many classification neural networks use the softmax activation. The softmax activation takes in a vector s of k un-normalized values s_1, \dots, s_k and outputs a probability distribution over the k possible classes. The forward pass of the softmax activation on input s_i is

$$\sigma_i = \frac{e^{s_i}}{\sum_{j=1}^k e^{s_j}},$$

where k ranges over all elements in s . Due to issues of numerical stability, the following modified version of this function is commonly used.

$$\sigma_i = \frac{e^{s_i - m}}{\sum_{j=1}^k e^{s_j - m}},$$

where $m = \max_{j=0}^k s_j$. We recommend implementing this method.

Instructions

1. Derive the Jacobian of the softmax activation function. **You do not need to use batched inputs for this question, an answer for a single training point is acceptable.**
2. Implement the forward and backward passes of the softmax activation in `activations.py`. We recommend vectorizing the backward pass for efficiency. **For this question only may you use a for loop over training examples in the batch.**
3. Include the output of `python -m unittest -v tests.test_activations.TestSoftMax`. Make sure that the tests are passing.

4.4 Cross-Entropy Loss

For this classification network, we will be using the multi-class cross-entropy loss function

$$L = -y \ln(\hat{y}),$$

where y is the binary one-hot vector encoding the ground truth labels and \hat{y} is the network's output, a vector of probabilities over classes. The cross-entropy cost calculated for a mini-batch of m samples is

$$J = -\frac{1}{m} \left(\sum_{i=1}^m Y_i \ln(\hat{Y}_i) \right).$$

Instructions

1. Derive the gradient of the cross-entropy cost with respect to the network's predictions, \hat{Y} . **You must use batched inputs.**
2. Implement the forward and backward passes of the cross-entropy cost. Note that in the codebase we have provided, we use the words “loss” and “cost” interchangeably. This is consistent with most large neural network libraries, though technically “loss” denotes the function computed for a single datapoint whereas “cost” is computed for a batch. You will be computing over batches. **Do not iterate over training examples, use batched operations.**
3. Include the output of `python -m unittest -v tests.test_losses.TestCrossEntropy`. Make sure that the tests are passing.

5 Two-Layer Fully Connected Networks

Now, you will use the methods you've written to train a two-layer network (also referred to as a one *hidden* layer network). You will use the **Iris Dataset**, which contains 4 features for 3 different classes of irises.

Instructions

1. Fill in the `forward`, `backward`, and `predict` methods for the `NeuralNetwork` class in `models.py`. Define the parameters of your network in `train_ffnn.py`. We have provided you with several other classes that are critical for the training process.
 - The data loader (in `datasets.py`), which is responsible for loading batches of data that will be fed to your model during training. You may wish to alter the data loader to handle data pre-processing. Note that all datasets you are given have not been normalized or standardized.
 - The stochastic gradient descent optimizer (in `optimizers.py`), which performs the gradient updates and optionally incorporates a [momentum term](#).
 - The learning rate scheduler (in `schedulers.py`), which handles the optional learning rate decay. You may choose to use either a constant or exponentially decaying learning rate.
 - Weight initializers (in `weights.py`). We provide you with many options to explore, but we recommend using `xavier_uniform` as a default.
 - A logger (in `logs.py`), which saves hyperparameters and learned parameters and plots the loss as your model trains.

Outputs will be saved to the folder `experiments/`. You can change the name of the folder a given run saves to by changing the parameter called `model_name`. Be careful about overwriting folders; if you forget to change the name and perform a run with identical hyperparameters, your previous run will be overwritten!

2. Train a 2-layer neural network on the Iris Dataset while varying the following hyperparameters.

- Learning rate
- Hidden layer size

You must try at least 3 different *combinations* of these hyperparameters. Report the results of your exploration, including the values of the parameters you explored and which set of parameters gave the best test error. Provide plots showing the loss versus iterations for your best model and report your final test error.

6 CNN Layers

In this problem, you will write the forward and backward passes of a convolutional neural network layer. Convolutional neural networks take considerably longer to train than the feedforward layers we have built, and the numpy implementation you will write here will be impractically slow. So, you will not be asked to train your network. Instead, we will simply run tests on the forward and backward passes of your network.

Again, remember that all your implementations and derived gradients must be for mini-batches, and not single examples. For CNN's, the input tensor will be of shape (m, H, W, C) .

6.1 Convolutional Layer

1. Derive the gradient of the loss with respect to the input and parameters (weights and biases) of a convolutional layer. **For this question your answer may be in the form of individual component partial derivatives.** You may also ignore stride and assume both the filter and image are infinitely zero-padded outside of their bounds.
2. Fill in the forward and backward passes of the Conv2D layer in `layers.py`. Include your code in your submission. **Do not iterate over training examples, use batched operations.**
3. Run `python -m unittest -v tests.test_layers.TestConv2D`. Make sure the tests are passing and include the output in your submission.

6.2 Pooling Layers

1. Explain how we can use the backprop algorithm to compute derivatives through the max pooling and average pooling operations. (A plain English answer will suffice; equations are optional.) Carefully consider the values you must keep track of!
2. Fill in the forward and backward passes of the Pool2D layer in `layers.py`. Include your code in your submission. **Do not iterate over training examples, use batched operations.**
3. Run `python -m unittest -v tests.test_layers.TestPool2D`. Make sure the tests are passing and include the output in your submission.

7 Kaggle

Please see the Google Colab Notebook [here](#) that contains the questions.

As with every homework, you are allowed to use any setup you wish. However, we highly recommend you to use [Google Colab](#) for free access to GPUs, which will significantly improve the speed of neural network

training. Instructions are provided in the notebook itself. Note that if you have access to GPUs locally, then feel free to run the notebook on your computer.

The following sections mirror the Colab Notebook and provide the deliverables for each question.

7.1 MLP for Fashion MNIST

Deliverables:

1. Provide code for training an MLP on the fashion MNIST dataset (can be tagged from colab notebook and in the code appendix)
2. A plot of the training and validation loss for each epoch of training for at least 8 epochs.
3. A plot of the training and validation accuracy for each epoch, achieving a final validation accuracy of at least 82%.

7.2 CNNs for CIFAR-10

Deliverables:

1. Provide the code for training your CNN model (can be in appendix).
2. Submit to Kaggle and include your test accuracy in your report. Our simple reference solution gets a test accuracy of 74.8% after 10 epochs.
3. Provide at least 1 training curve for your model, depicting loss per epoch or step after training for at least 8 epochs.
4. Explain the components of your final model, and how you think your design choices contributed to it's performance.