

Homework 4: Basic CNN Construction - PyTorch

Objective:

The objective of this homework was to introduce or refamiliarize us with basic CNN construction, specifically in PyTorch using the COCO dataset as training and validation data. The project exercises our ability to read in a dataset, construct a CNN, build a training and validation wrapper, and evaluate the results.

Tasks:

Task 1: Prepare the dataset

The first step in this exercise was the preparation of the dataset for ingestion by a CNN. The homework refers to this as dataset creation, however, dataset creation typically refers to the painstaking process of label assignment to each image. In this particular case, we are using the COCO dataset which comes pre-labeled. Fortunately, the COCO dataset also comes with a python specific library that greatly helps in terms of parsing the data. We were asked to select 8000 training images and 2000 validation images in the following classes: boat, couch, dog, cake, motorcycle. The output of my code (Appendix A – Code Snippet 1) for this section balanced the classes by randomly retrieving 2000 of each class in the training set and 400 of each in the validation set, however, I belatedly realized that each image had multiple images so the result was 8000 and 2000 annotations (respective to training and validation sets) but only 7862 and 1992 images respectively. It also resized the images according to the requirements set out in the homework. Lastly, it produced to csv log files, one for training and one for validation, that logged image names, assigned class, and class number. The first 3 images of each class are presented in Figure 1.




Figure 1: Examples of COCO images by required classes

Task 2: CNN Construction

CNN construction was relatively simple given the example provided in the homework. Establishing the network architecture is very much streamlined compared to the last homework through the use of Pytorch, however, it is vital to understand how information is being passed between layers and how each layer interacts and changes the data before it is passed to the next layer. To calculate the layer output size and total number of layer parameters, I used the following two equations modified with variables for my own understanding taken from [here](#) and [here](#).

<u>Calculate layer output size</u> $\frac{W - F + 2P}{S} + 1$ Where W is input size, F is filter size (user input), P is padding, and S is stride	<u>Calculate layer parameters</u> $OC * (KS^2 * IC + 1)$ Where KS is kernel size, IC is input channels, and OC is output channels
---	---

Figure 2 below highlights the application of these equations to the three networks we were required to build, specifically as they relate to the homework requirements. The output feature map of the first fully connected layer were of the following sizes [6272, 8192, 6272] with respect to Net1, Net2, and Net3. For Net3, we were asked to ensure that the added convolutional layers had padding, however, we were not asked to add padding to the first two convolutional layers as compared to Net1. The total number of parameters were [406885, 535178, 499365] respectively. The final build of these networks in code form can be found in Appendix A – Code Snippet 2.




Figure 2: Calculation of layer inputs, outputs, and parameters for each required network

Task 3: Data Loading, Model Training, and Validation Inference

After the networks were built, the next step involved developing a modified data loader subclass that modifies the imagery in preparation for passing to the networks for training or evaluation. This subclass reads in the csv data saved previously when the training and validation data was appropriated and separated into new folders as well as modifying images as they are read in. There is a transform that is applied converting the image to a tensor while normalizing between the values of -1 to 1. I also noticed that there were some grayscale images (1xchannel) that needed to be converted to RGB when appropriate. This data loader subclass then passes the modified images to the DataLoader class during training. Code for the subclass can be found in Appendix A – Code Snippet 3.

The other section of this code necessitated the development of a model training function that accomplished several functions. This is the section that caused me the most trouble, so it is heavily commented, but the first step was importing the data by first assigning the data loader subclass to the correct location and then passing to the data loader base class. Next, I assigned the GPU to conduct the work before loading the network, assigning the optimizer, conducting model training, and logging the training loss. Upon completion of training, this function also switches to evaluation mode and inferences all the validation data, producing a dataframe with predictions and “ground truth” labels. I would normally split these two steps so that the model weights were saved to memory and later loaded in for validation (really testing) and application, however, that sort of workflow is unnecessary for the homework requirements so this condensed version works better. The output of this function is the loss log for every epoch and the prediction/reference log. The function itself can be found in Appendix A – Code Snippet 4.

Task 4: Model Evaluation

The last coding requirement was the evaluation of the model through traditional means. The functions I used to plot the results and the __main__ call function tying all the previous code together in one bundle can be found in Appendix A – Code Snippet 5 and 6 respectively. The plotting function produces both a loss comparison between networks, as well as confusion matrices for each network. Lastly, it produces one csv per network that provides overall accuracy per class. All tables and figures are provided below.

Network	COCO Class					Overall
	Boat	Couch	Dog	Cake	Motorcycle	
Net 1	0.613	0.540	0.343	0.558	0.555	0.522
Net 2	0.650	0.480	0.265	0.485	0.633	0.525
Net 3	0.600	0.498	0.335	0.508	0.568	0.515

Table 1: Model accuracies by COCO class and CNN




Figure 3: Training Loss Comparison between the three networks.




Figure 4: Confusion matrix for each model.

Discussion Questions:

1. My results suggest that adding padding didn't make much of a difference in terms of classification performance, however, I understand that this is not always the case. Padding ensures minimal loss of data when passing through convolutional layers which may be important depending on your application. My intuition is that padding becomes more important either when image resolution is poor or when object relationships within the images are complex. In the first case, each pixel gains in importance since there are so few to spare, while in the second case, loss of data may blur the relationship between objects.
2. Net 3 does have a vanishing gradient problem as highlighted by the resultant loss curve which is not a smooth downward slope toward convergence. I was actually surprised that Net 3 didn't result in over training, although it is likely if the epoch count was increased beyond the 60 I ran for my models. Net 3 just has too many layers for this application, and it highlights that when it comes to CNNs, increased complexity isn't always better.
3. Results indicate similar performance between all three networks. Network 1 is the best performing, followed by network 2 and then network 3, but the differences in accuracy are negligible. Network 1 performed best in three of the five classes (couch, dog, cake) while network 2 performed best in the other two of five classes (boat, motorcycle) but worst in the other three.
4. The easiest class to identify across all three networks was boat which I would expect would relate to image context since most boat pictures are going to include water which is a rather discernable feature as compared to the environments in which the other classes could be found. For a similar reason, motorcycles are next which are always going to be found in areas with roads, asphalt, or concrete. Conversely, couches, dogs, and cakes are found in environments that are highly varied (wall color, rug color, furniture, people, etc.).
5. For application, I would suggest moving forward with network 1 as it is the simplest of the networks that also happens to produce the best results if we continued to work with 64x64 images. To improve performance, I would first add in all training data that contained labels of the classes of interest from the COCO dataset. Assuming we are continuing to work with 64x64 images, I would next focus on hyperparameter tuning. Lastly, I would apply cross validation to ensure the model had an opportunity to train on all available images given the limited number available. This would require the separation of a true test set but might be worth it in the end.

Appendix A – Code

```
1 #Import Statements
2 import os
3 from PIL import Image
4 import matplotlib.pyplot as plt
5 import pandas as pd
6 from pycocotools.coco import COCO
7 import random
8 import torchvision.transforms as tvf
9 import torch
10 #####
11 #Set initial variables
12 annotations_path=r"C:/BME_646/coco_dataset/annotations_trainval2017/annotations/instances_train2017.json"
13 images_path=r"C:/BME_646/coco_dataset/train2017/"
14 hw_train_path=r"C:/BME_646/coco_dataset/hw_subset/train/"
15 hw_val_path=r"C:/BME_646/coco_dataset/hw_subset/val/"
16 csv_out_path=r"C:/BME_646/coco_dataset/hw_subset/"
17 categories=['boat','couch','dog','cake','motorcycle']
18 num_train=8000
19 num_val=2000
20 num_figs=15
21 #####
22 #Load the Annotation file using COCO
23 annotations=COCO(annotations_path)
24 #####
25 #Load images, resize and save them
26 i=0
27 train_temp=[]
28 val_temp=[]
29 for category in categories:
30     #reassign class numbers
31     cls_num=i
32     i+=1
33     #Create a list of image names based on category
34     img_ids=annotations.getImgIds(catIds=annotations.getCatIds(catNms=[category]))
35     #creating balanced respresentation of classes in train and val datasets
36     print("Working on {}".format(category))
37     img_ids=random.sample(img_ids, int((num_train-num_val)/len(categories)))
38     j=0
39     for img_id in img_ids:
40         j+=1
41         #Load image
42         img_meta=annotations.loadImgs(img_id)[0]
43         temp_img=Image.open(images_path+img_meta['file_name'])
44         #resize
45         temp_img=temp_img.resize((64,64))
46         if j<int(num_train/len(categories)):
47             train_temp.append({'img_num':img_meta['file_name'],'class':category,'class_num':cls_num})
48             temp_img.save(hw_train_path+img_meta['file_name'])
49         else:
50             val_temp.append({'img_num':img_meta['file_name'],'class':category,'class_num':cls_num})
51             temp_img.save(hw_val_path+img_meta['file_name'])
52 #Convert training and val lists to dataframes and write to csv
53 #to parse during data loading
54 train_df=pd.DataFrame(train_temp)
55 val_df=pd.DataFrame(val_temp)
56 train_df.to_csv(csv_out_path+"training_log.csv", index=False)
57 val_df.to_csv(csv_out_path+"validation_log.csv", index=False)
58 #####
59 #Plot the first 3 images per class
60 three_instances=train_df.groupby('class').head(3)
61 fig, axes = plt.subplots(nrows=5, ncols=3, figsize=(10,15))
62 axes=axes.flatten()
63 #Loop through the DataFrame and plot each image
64 for i, (fig, (_, row)) in enumerate(zip(axes, three_instances.iterrows())):
65     img_path = hw_train_path+row['img_num']
66     img = Image.open(img_path)
67     fig.imshow(img)
68     fig.set_xlabel(row['class'])
69     fig.xaxis.set_label_position('bottom')
70     fig.tick_params(axis='both', which='both', length=0) # Hide tick marks
71     fig.xaxis.set_ticks_position('none')
72     fig.yaxis.set_ticks_position('none')
73     fig.set_yticklabels([])
74     fig.set_xticklabels([])
75 #Adjust layout
76 plt.tight_layout()
77 plt.show()
```

Code Snippet 1: Dataset preparation code - Selection and delineation of training and validation sets from the COCO dataset as well as plotting of three examples from each class.

```
24 ##### BUILD NETWORKS
25 #Establish Network 1 - From Homework
26 class Net1(nn.Module):
27     def __init__(self):
28         super(Net1,self).__init__()
29         self.conv1 = nn.Conv2d(3, 16, 3)
30         self.pool = nn.MaxPool2d(2, 2)
31         self.conv2 = nn.Conv2d(16, 32, 3)
32         self.fc1 = nn.Linear(6272, 64)
33         self.fc2 = nn.Linear(64,5)
34         #How data is fed through the network based on attributes of class
35     def forward(self, x):
36         x = self.pool(F.relu(self.conv1(x)))
37         x = self.pool(F.relu(self.conv2(x)))
38         x = x.view(x.shape[0], -1)
39         x = F.relu(self.fc1(x))
40         x = self.fc2(x)
41         return x
42
43 #Establish Network 2 - From Homework
44 class Net2(nn.Module):
45     def __init__(self):
46         super(Net2, self).__init__()
47         self.conv1 = nn.Conv2d(3, 16, 3, padding=1)
48         self.pool = nn.MaxPool2d(2, 2)
49         self.conv2 = nn.Conv2d(16, 32, 3, padding=1)
50         self.fc1 = nn.Linear(8192, 64)
51         self.fc2 = nn.Linear(64,5)
52         #How data is fed through the network based on attributes of class
53     def forward(self, x):
54         x = self.pool(F.relu(self.conv1(x)))
55         x = self.pool(F.relu(self.conv2(x)))
56         x = x.view(x.shape[0], -1)
57         x = F.relu(self.fc1(x))
58         x = self.fc2(x)
59         return x
60
61 #Establish Network 3 - Modification from Net2
62 class Net3(nn.Module):
63     def __init__(self):
64         super(Net3,self).__init__()
65         self.conv1 = nn.Conv2d(3, 16, 3)
66         self.pool = nn.MaxPool2d(2, 2)
67         self.conv2 = nn.Conv2d(16, 32, 3)
68         self.extra_conv1 = nn.Conv2d(32, 32, 3, padding=1)
69         self.extra_conv2 = nn.Conv2d(32, 32, 3, padding=1)
70         self.extra_conv3 = nn.Conv2d(32, 32, 3, padding=1)
71         self.extra_conv4 = nn.Conv2d(32, 32, 3, padding=1)
72         self.extra_conv5 = nn.Conv2d(32, 32, 3, padding=1)
73         self.extra_conv6 = nn.Conv2d(32, 32, 3, padding=1)
74         self.extra_conv7 = nn.Conv2d(32, 32, 3, padding=1)
75         self.extra_conv8 = nn.Conv2d(32, 32, 3, padding=1)
76         self.extra_conv9 = nn.Conv2d(32, 32, 3, padding=1)
77         self.extra_conv10 = nn.Conv2d(32, 32, 3, padding=1)
78         self.fc1 = nn.Linear(6272, 64)
79         self.fc2 = nn.Linear(64,5)
80         #How data is fed through the network based on attributes of class
81     def forward(self, x):
82         x = self.pool(F.relu(self.conv1(x)))
83         x = self.pool(F.relu(self.conv2(x)))
84         x = F.relu(self.extra_conv1(x))
85         x = F.relu(self.extra_conv2(x))
86         x = F.relu(self.extra_conv3(x))
87         x = F.relu(self.extra_conv4(x))
88         x = F.relu(self.extra_conv5(x))
89         x = F.relu(self.extra_conv6(x))
90         x = F.relu(self.extra_conv7(x))
91         x = F.relu(self.extra_conv8(x))
92         x = F.relu(self.extra_conv9(x))
93         x = F.relu(self.extra_conv10(x))
94         x = x.view(x.shape[0], -1)
95         x = F.relu(self.fc1(x))
96         x = self.fc2(x)
97         return x
```

Code Snippet 2: Network architectures built Pytorch

```
100 ##### DATA LOADING
101 #Establish the image to tensor transformation
102 transform_comp=tv.t.Compose([tv.t.ToTensor(),tv.t.Normalize(mean =[0.5], std =[0.5])])
103
104 #Establish Dataset Class
105 class cocoData_loader(Dataset):
106     def __init__(self, csv_path, img_dir, transform=transform_comp):
107         #Passed Attributes
108         self.csv_data=pd.read_csv(csv_path)
109         self.img_dir=img_dir
110         self.transform=transform
111         #Derived attributes
112         self.img_num=self.csv_data['img_num']
113         self.class_name=self.csv_data['class']
114         self.class_num=self.csv_data['class_num']
115
116     def __len__(self):
117         return (len(self.csv_data))
118
119     def __getitem__(self, row):
120         temp_img=Image.open(self.img_dir+self.img_num[row])
121         label=torch.tensor(self.class_num[row], dtype=torch.int64)
122         #I came across some images that were gray, so these need
123         #to be converted to RGB (e.g., 4971, 5324, etc.)
124         if temp_img.mode=='L':
125             temp_img=temp_img.convert('RGB')
126         #Apply transform
127         temp_img=self.transform(temp_img)
128         return temp_img, label
```

Code Snippet 3: Data Loader subclass designed to read in csv data and modify images before passing to DataLoader base class.

```
131 ##### MODEL TRAINING
132 def train_model(csv_dir, img_dir, network, epochs, batch_s=32, learning_rate=1e-3):
133     ### MODEL TRAINING
134     #Import the data
135     train_set = cocoData_loader(csv_path = os.path.join(csv_dir, 'training_log.csv'),
136                                img_dir = os.path.join(img_dir, 'train/'))
137     train_loader = DataLoader(train_set, batch_size = batch_s, shuffle = True, num_workers = 4)
138
139     val_set = cocoData_loader(csv_path = os.path.join(csv_dir, 'validation_log.csv'),
140                               img_dir = os.path.join(img_dir, 'val/'))
141     val_loader = DataLoader(val_set, batch_size = batch_s, shuffle = False, num_workers = 4)
142     #Assign gpu if available
143     device = torch.device("cuda:0" if torch.cuda.is_available() else "cpu")
144     #Pass the model to the gpu
145     model = network.to(device)
146     criterion = nn.CrossEntropyLoss()
147     #Assign Torch's built in Adam optimizer
148     optimizer = torch.optim.Adam(model.parameters(), lr = learning_rate, betas=(0.9, 0.99))
149     #Begin model training
150     model.train()
151     #Keep a record of the loss
152     loss_log = []
153     for epoch in range(epochs):
154         running_loss = 0.0
155         epoch_loss = 0.0
156         for data in train_loader:
157             #Pass the inputs
158             inputs, labels = data
159             inputs, labels = inputs.to(device), labels.to(device)
160             #Clear old gradients from last epoch
161             optimizer.zero_grad()
162             #conduct the forward pass
163             outputs = model(inputs)
164             #Calculate the batch loss
165             loss = criterion(outputs, labels.long())
166             #Calculate the loss with respect to the parameters
167             loss.backward()
168             #update the parameters accordingly
169             optimizer.step()
170             #record the loss for posterity
171             running_loss += loss.item()
172             epoch_loss += loss.item()
173             #Update the loss list for later plotting
174             loss_log.append(epoch_loss / len(train_loader))
175             print('Epoch {} Complete - Epoch Loss= {:.2f} , Total Loss={:.2f}'.format(str(epoch), epoch_loss, running_loss))
176         print('Finished Training')
177
178     ##### VALIDATION DATASET
179     #Convert model to evaluation mode
180     model.eval()
181     #Initialize empty results
182     total_correct = 0
183     total_samples = 0
184     pred_list = []
185     ref_list = []
186
187     with torch.no_grad():
188         for images, labels in val_loader:
189             #Pass the inputs
190             images, labels = images.to(device), labels.to(device)
191             outputs = model(images)
192             #Apply softmax and find the maximum probability class
193             _, predicted = torch.max(F.softmax(outputs, dim=1), 1)
194             #update the result parameters
195             total_correct += (predicted == labels).sum().item()
196             total_samples += labels.size(0)
197             #Append predictions and references to the lists
198             pred_list.extend(predicted.cpu().numpy())
199             ref_list.extend(labels.cpu().numpy())
200     #Convert lists to DataFrame
201     pred_ref = pd.DataFrame({'pred': pred_list, 'ref': ref_list})
202     accuracy = 100 * total_correct / total_samples
203     print(f'Accuracy on {total_samples} validation images: {accuracy:.2f} %')
204
205     return loss_log, pred_ref
206
```

Code Snippet 4: Data Loader and Validation Inference

```
131 ##### MODEL TRAINING
132 def train_model(csv_dir, img_dir, network, epochs, batch_s=32, learning_rate=1e-3):
133     ### MODEL TRAINING
134     #Assign gpu if available
135     device = torch.device("cuda:0" if torch.cuda.is_available() else "cpu")
136     #Import the data
137     train_set = cocoData_loader(csv_path = os.path.join(csv_dir, 'training_Log.csv'),
138                                img_dir = os.path.join(img_dir, 'train/'))
139     train_loader = DataLoader(train_set, batch_size = batch_s, shuffle = True, num_workers = 4)
140
141     val_set = cocoData_loader(csv_path = os.path.join(csv_dir, 'validation_Log.csv'),
142                              img_dir = os.path.join(img_dir, 'val/'))
143     val_loader = DataLoader(val_set, batch_size = batch_s, shuffle = False, num_workers = 4)
144     #Pass the model to the gpu
145     model = network.to(device)
146     criterion = nn.CrossEntropyLoss()
147     #Assign Torch's built in Adam optimizer
148     optimizer = torch.optim.Adam(model.parameters(), lr = learning_rate, betas=(0.9, 0.99))
149     #Begin model training
150     model.train()
151     #Keep a record of the loss
152     loss_log = []
153     for epoch in range(epochs):
154         running_loss = 0.0
155         epoch_loss = 0.0
156         for data in train_loader:
157             #Pass the inputs
158             inputs, labels = data
159             inputs, labels = inputs.to(device), labels.to(device)
160             #Clear old gradients from last epoch
161             optimizer.zero_grad()
162             #conduct the forward pass
163             outputs = model(inputs)
164             #Calculate the batch loss
165             loss = criterion(outputs, labels.long())
166             #Calculate the loss with respect to the parameters
167             loss.backward()
168             #update the parameters accordingly
169             optimizer.step()
170             #record the loss for posterity
171             running_loss += loss.item()
172             epoch_loss += loss.item()
173             #Update the loss list for later plotting
174             loss_log.append(epoch_loss / len(train_loader))
175             print('Epoch {} Complete - Epoch Loss= {:.2f} , Total Loss= {:.2f}'.format(str(epoch), epoch_loss, running_loss))
176         print('Finished Training')
177
178     ### VALIDATION DATASET
179     #Convert model to evaluation mode
180     model.eval()
181     #Initialize empty results
182     total_correct = 0
183     total_samples = 0
184     pred_list = []
185     ref_list = []
186
187     with torch.no_grad():
188         for images, labels in val_loader:
189             #Pass the inputs
190             images, labels = images.to(device), labels.to(device)
191             outputs = model(images)
192             #Apply softmax and find the maximum probability class
193             _, predicted = torch.max(F.softmax(outputs, dim=1), 1)
194             #update the result parameters
195             total_correct += (predicted == labels).sum().item()
196             total_samples += labels.size(0)
197             #Append predictions and references to the lists
198             pred_list.extend(predicted.cpu().numpy())
199             ref_list.extend(labels.cpu().numpy())
200     #Convert lists to DataFrame
201     pred_ref = pd.DataFrame({'pred': pred_list, 'ref': ref_list})
202     accuracy = 100 * total_correct / total_samples
203     print(f'Accuracy on {total_samples} validation images: {accuracy:.2f} %')
204
205     return loss_log, pred_ref
```

Code Snippet 5: Model training and accuracy calculations against the validation set

```
207 ##### MODEL EVALUATION
208 def plot_results(loss_logs, preds_refs, epoch_num):
209     model_names = ['Net1', 'Net2', 'Net3']
210     # Plotting Loss
211     plt.figure(figsize=(12, 10))
212     for i, loss_log in enumerate(loss_logs):
213         plt.plot(range(epoch_num), loss_log, label=model_names[i])
214     plt.xlabel('Epochs')
215     plt.ylabel('Loss')
216     plt.legend()
217     plt.savefig('C:/BME_646/coco_dataset/output/Loss_Comparison.png')
218     #plt.show()
219     #Plotting Confusion Matrix
220     categories = ['boat', 'couch', 'dog', 'cake', 'motorcycle']
221     for preds_ref, net_name in zip(preds_refs, model_names):
222         cm = confusion_matrix(preds_ref['ref'].values, preds_ref['pred'].values)
223         class_acc = pd.DataFrame({'class': categories, 'accuracy': [cm[i][i] / sum(cm[i]) for i in range(len(categories))])})
224         class_acc.to_csv('C:/BME_646/coco_dataset/output/class_accuracy_{}.csv'.format(net_name), index=False)
225         cm_plot = ConfusionMatrixDisplay(cm, display_labels=categories).plot()
226         cm_plot.figure_.savefig('C:/BME_646/coco_dataset/output/confusion_matrix_{}.png'.format(net_name))
```

Code Snippet 6: Plotting Function for both loss comparison and confusion matrices.

```
229 ##### RUN ALL THIS STUFF
230 #Do all the above when this script is called
231 if __name__ == '__main__':
232     #Initialize some hyperparameters
233     epoch_num=60
234     batch_s=32
235     learning_rate=1e-3
236     csv_log_root='C:/BME_646/coco_dataset/hw_subset/csv_logs'
237     data_root='C:/BME_646/coco_dataset/hw_subset'
238     # Run All three models
239     results = [train_model(csv_log_root, data_root, net, epoch_num) for net in [Net1(), Net2(), Net3()]]
240     # Splitting the results into separate lists
241     loss_lists, preds_refs = zip(*results)
242     # Saving predictions references
243     for preds_ref, i in zip(preds_refs, range(1, 4)):
244         preds_ref.to_csv('C:/BME_646/coco_dataset/output/pred_ref{}.csv'.format(i), index=False)
245     # Plotting loss and confusion matrix for all networks
246     plot_results(loss_lists, preds_refs, epoch_num)
```

Code Snippet 7: __main__ run function