

ECE 60146

## Homework 5

March 6, 2023

Alex Rogers

roger299@purdue.edu

---

### 1 Introduction

This assignment is an introduction into using skip connections and creating a network for single object detection as well as creating and utilizing a custom dataset for training. It also gives experience with a well known and used dataset such as COCO. It also gives a brief introduction into performance metrics for classifiers and bounding box applications.

### 2 Custom COCO Image Classification Dataset

The COCO dataset was downloaded locally and then was processed using the code in section 5.1. As described by the assignment, only the three classes *bus*, *cat*, *pizza* were kept with the condition that a single dominant object of the class greater than 40,000 pixels was present. The code was designed to ensure there was only one of these dominant objects in the image and to ensure there were no duplicate images. Separate datasets were created for the train and validation sets using the corresponding COCO versions. In total, there were 3788 training images and 1971 validation images. A figure showing examples of the training set with the dominant image highlighted via the bounding box can be found in 1. Note that the pizza image in the bottom right of the figure has a bounding box which takes up the entire image and thus it can't be seen in the outputted image.

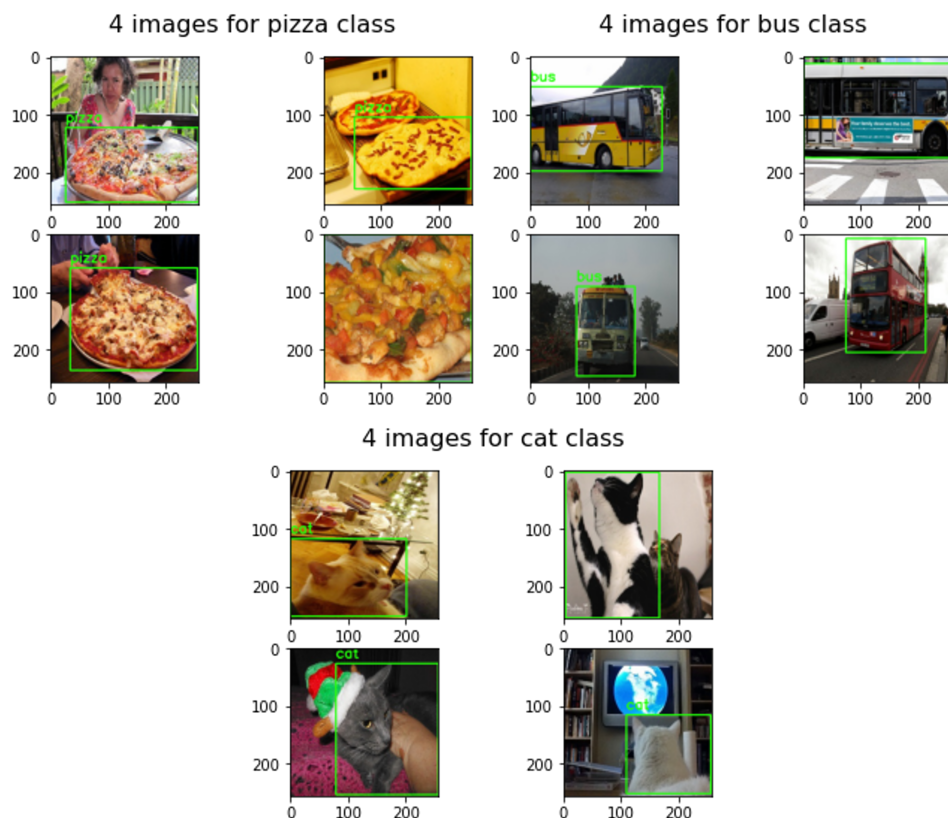


Figure 1: Four Images for Each Class in Custom Dataset

### 3 Image Classification Using CNNs

A custom dataloader was first created in order to easily load in the custom dataset. This along with the code for the network architecture, training, testing, and data visualization can be found in section 5.2. The `SkipBlock` class is in this section as well which outlines the details of the created skip connection.

#### 3.1 Mean Squared Error Regression Network

The first network created and trained was one using mean squared error for the bounding box regression task. This network had a total of **144 layers** as determined by the `net.parameters()` logic we were given in the assignment. The loss of both the regression and classification tasks can be found in figure 2 where 20 epochs with a learning rate of  $1e-4$  was used. The confusion matrix can be found in figure 3. The average classification

accuracy for all three classes is **86.10%** and the mean IOU is **0.576**. An example of the ground truth and predicted bounding boxes for each image can be found in figure 4.

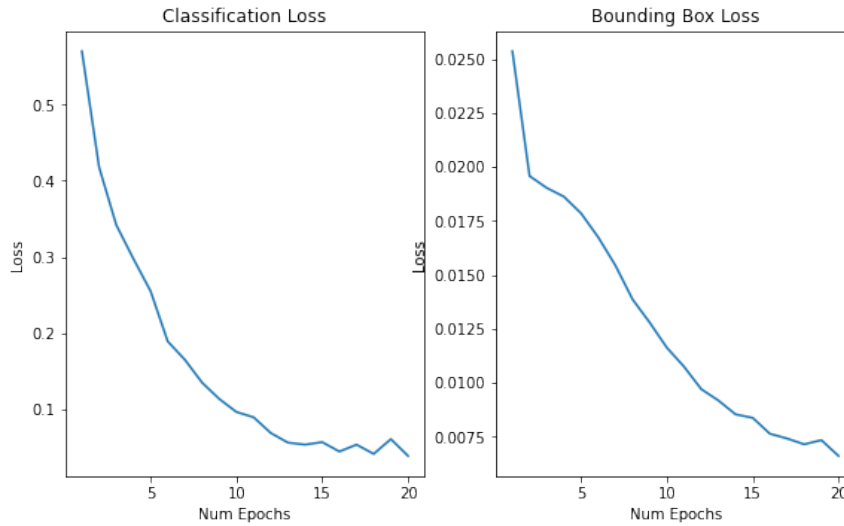


Figure 2: Average Epoch Loss For MSE Network

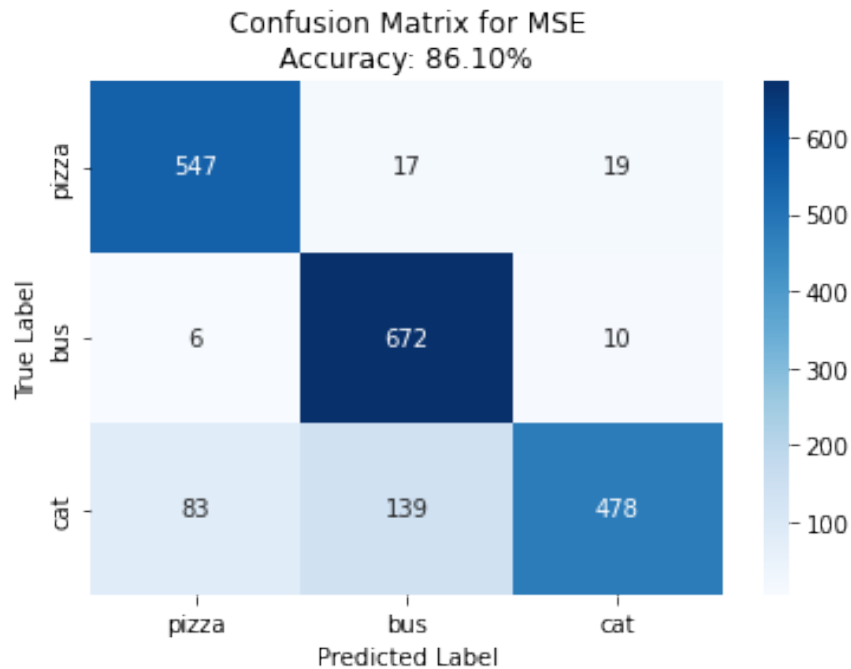


Figure 3: Confusion Matrix for MSE Network

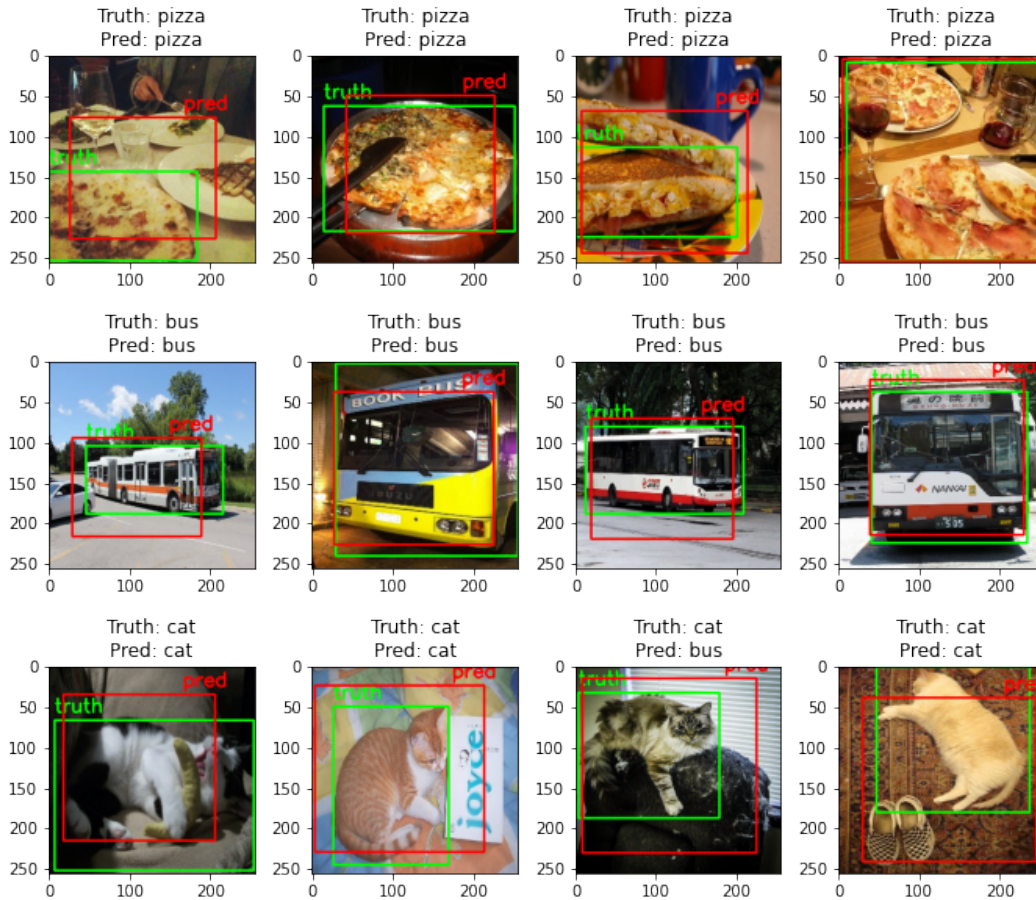


Figure 4: Comparison of Ground Truth and Predicted BBox for MSE Network

### 3.2 Complete IoU Regression Network

The second network created and trained was one using complete IoU (CIoU) for the bounding box regression task. This network also had a total of **144 layers** as determined by the `net.parameters()` logic we were given in the assignment. The loss of both the regression and classification tasks can be found in figure 5 where 20 epochs with a learning rate of  $1e - 4$  was used. The confusion matrix can be found in figure 6. The average classification accuracy for all three classes is **90.26%** and the mean IOU is **0.645**. An example of the ground truth and predicted bounding boxes for each image can be found in figure 7.

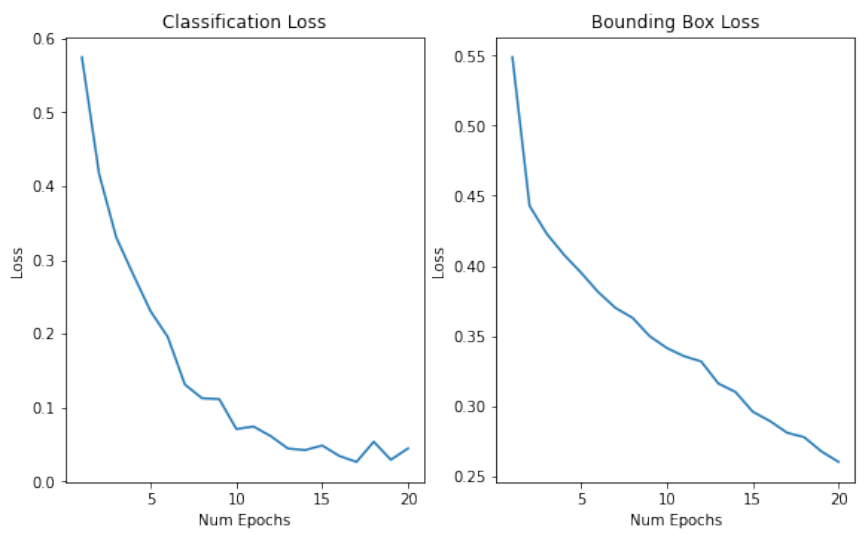


Figure 5: Average Epoch Loss For CIOU Network

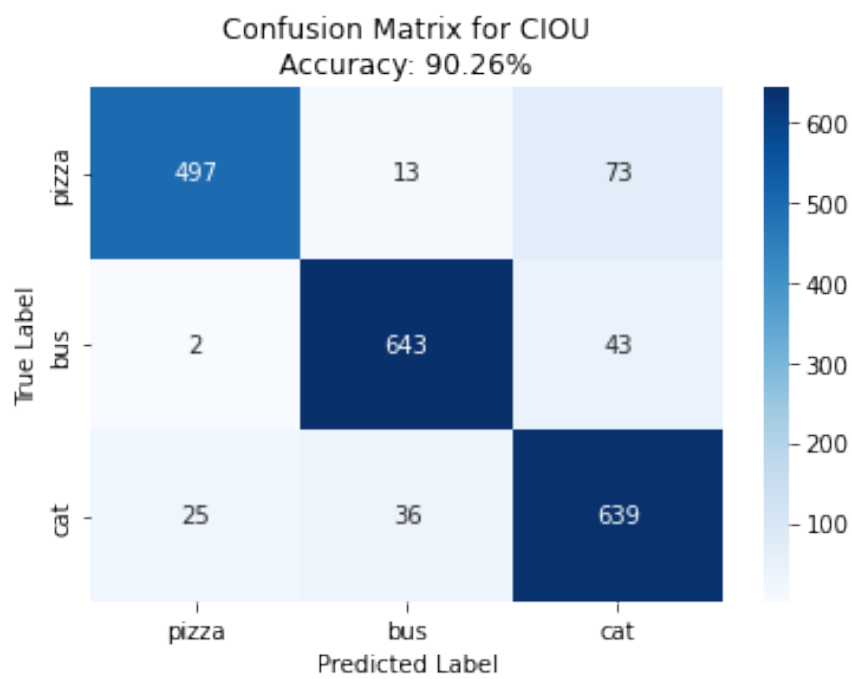


Figure 6: Confusion Matrix for CIOU Network

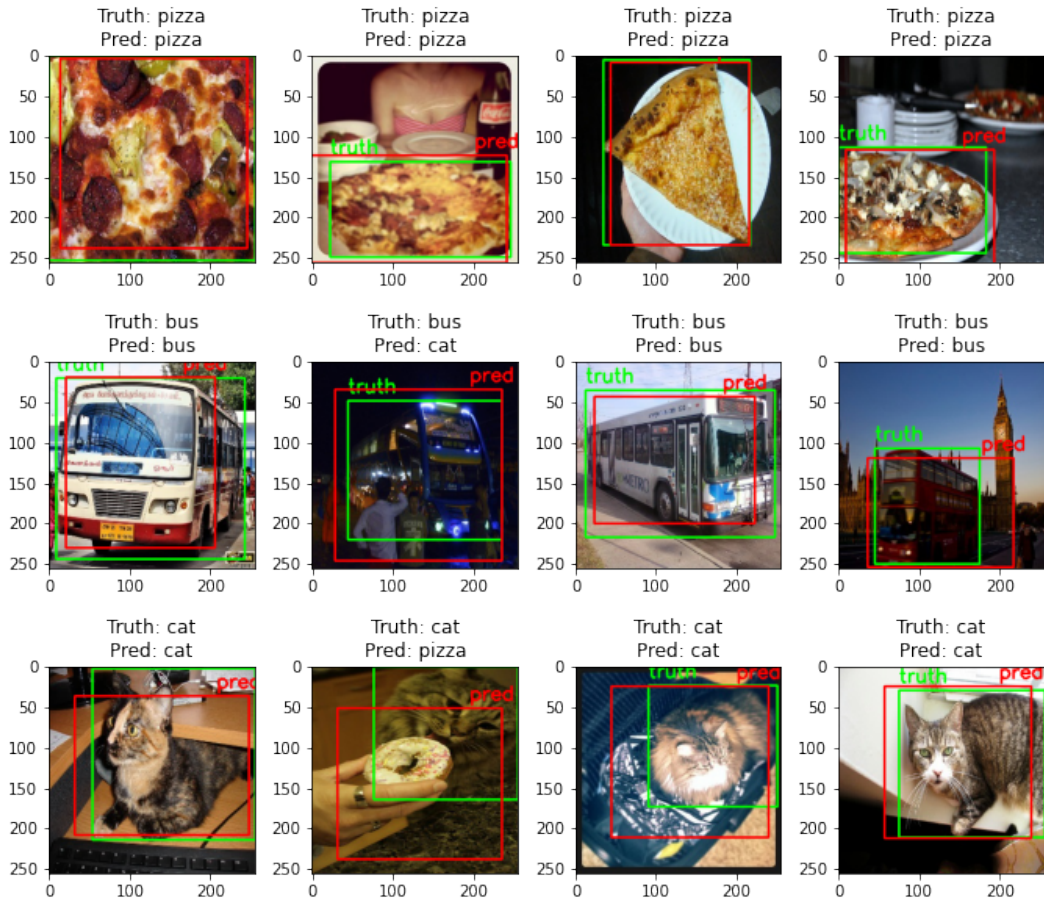


Figure 7: Comparison of Ground Truth and Predicted BBox for CIOU Network

## 4 Concluding Discussion

In all, the classifier seems to work relatively well on the given dataset. All of the losses decrease during training and appear to approach some form of a minimum which is a good sign that the network is learning. The classification of the classes performs at a pretty high rate only missing a few of the weirder images. Looking at some of the images, it is obvious why the classifier performs poorly since it only has a portion of the object in view or there are multiple objects in the image. The multiple objects seemed to mess up the bounding box the most as it would identify a box covering most of the multiple objects image instead of just the dominant one. Despite this, it seems to draw the bounding boxes well for both loss functions. The CIOU loss function performs decently better than the MSE loss which makes sense since the MSE loss isn't specific to this task and thus

has a very very small loss. There is certainly lots of room for improvement. Given more time, hyper parameter tuning could be performed in order to find the optimal learning rate, momentum, batch size, etc. I did a little bit of this, but did not have the time or resources to get the most out of it. Furthermore, changing the dataset to include images with only a single annotation regardless of size would likely help the bounding box success. Finally, the network layers could certainly be designed better should I have more time and knowledge. Someone with more experience could better determine which types of hidden layers to use and what the size of each should be.

## 5 Source Code

The source code was broken across two different files: `custom_dataset.py` and `hw5_network.py`.

### 5.1 `custom_dataset.py`

```
1  '''
2  The code in this file is used to generate the custom COCO dataset
3  '''
4  from pycocotools.coco import COCO
5  import numpy as np
6  import matplotlib.pyplot as plt
7  from PIL import Image
8  import cv2
9  import json
10 import skimage
11 import json
12 from skimage import data, io, filters
13
14 def get_images_with_dominant_obj(coco, catIds):
15     """Find all the images that have a dominant object in them
16     """
17     return_images = {}
18
19     # get the unique images for the categories
20     uniqueImgIds = set()
21     for catId in catIds:
22         imgIds = coco.getImgIds(catIds=catId)
23         uniqueImgIds |= set(imgIds)
```

```

24
25 # loop through the categories
26 for i, img in enumerate(coco.loadImgs(list(uniqueImgIds))):
27     annIds = coco.getAnnIds(imgIds=img['id'], catIds=catIds, iscrowd=False)
28     anns = coco.loadAnns(annIds)
29
30     curr_ann = []
31     # find the images with one dominant object
32     for ann in anns:
33         if ann['category_id'] in catIds and ann['area'] > 40000:
34             curr_ann.append(ann)
35     # make sure there is only one dominant object
36     if len(curr_ann) == 1:
37         assert img['id'] not in return_images, img['id']
38         return_images[img['id']] = {'coco_url': img['coco_url'], 'annotations': curr_ann[0]}
39
40     return return_images
41
42 def save_images(images, folder_name, coco_labels_inverse):
43     """Save images into custom dataset folder
44     """
45     labels = {}
46     # loop through all the images
47     for i, (k, img) in enumerate(images.items()):
48         if i % 100 == 0:
49             print(i)
50
51     # read image from coco API and convert
52     I = io.imread(img['coco_url'])
53     if len(I.shape) == 2:
54         I = skimage.color.gray2rgb(I)
55     image = np.uint8(I)
56     h, w, c = image.shape
57
58     # resize image and bounding boxes
59     xFactor = 256 / w
60     yFactor = 256 / h
61     image = cv2.resize(image, (256, 256))
62     curr_bbox = img['annotations']['bbox']
63     new_bbox = [curr_bbox[0]*xFactor, curr_bbox[1]*yFactor,
64                 curr_bbox[2]*xFactor, curr_bbox[3]*yFactor]

```



```

65     label = coco_labels_inverse[img['annotations']['category_id']]
66
67     # save image and label for image
68     image_name = f'{i}_{label}.jpg'
69     assert image_name not in labels
70     labels[image_name] = {'label': label, 'bbox': new_bbox}
71     image = cv2.cvtColor(image, cv2.COLOR_BGR2RGB)
72     cv2.imwrite(f'custom_dataset/{folder_name}/{i}_{label}.jpg', image)
73
74     return labels
75
76 def display_class_images(imageBoxes, class_num, class_list):
77     """Display class_num number of images for a given class
78     """
79     imgs = []
80     for k, v in imageBoxes.items():
81         if len(imgs) == 4:
82             break
83         if v['label'] == class_num:
84             image = np.uint8(Image.open(f'custom_dataset/train/{k}'))
85             [x, y, w, h] = v['bbox']
86             # draw bounding boxes
87             image = cv2.rectangle(image, (int(x), int(y)), (int(x + w), int(y + h)), (36,255,12), 2)
88             image = cv2.putText(image, class_list[v['label']], (int(x), int(y - 10)),
89                                 cv2.FONT_HERSHEY_SIMPLEX, 0.8, (36,255,12), 2)
90             imgs.append(image)
91
92     fig, ax = plt.subplots(2,2)
93     ax[0,0].imshow(imgs[0])
94     ax[0,1].imshow(imgs[1])
95     ax[1,0].imshow(imgs[2])
96     ax[1,1].imshow(imgs[3])
97     fig.suptitle(f"4 images for {class_list[class_num]} class", fontsize=16)
98
99 if __name__ == '__main__':
100     class_list = ['pizza', 'bus', 'cat']
101
102     cocoTrain = COCO('annotations/instances_train2014.json')
103     cocoVal = COCO('annotations/instances_val2014.json')
104
105     catIdsTrain = cocoTrain.getCatIds(catNms=class_list)

```

```

106 catIdsVal = cocoVal.getCatIds(catNms=class_list)
107
108 # create train and validation sets
109 trainImages = get_images_with_dominant_obj(cocoTrain, catIdsTrain)
110 valImages = get_images_with_dominant_obj(cocoVal, catIdsVal)
111
112 # pulled from the homework assignment example code
113 categories = cocoTrain.loadCats(catIdsTrain)
114 coco_labels_inverse_train = {}
115 for idx, in_class in enumerate(class_list):
116     for c in categories:
117         if c['name'] == in_class:
118             coco_labels_inverse_train[c['id']] = id
119
120 # pulled from the homework assignment example code
121 categories = cocoVal.loadCats(catIdsVal)
122 coco_labels_inverse_val = {}
123 for idx, in_class in enumerate(class_list):
124     for c in categories:
125         if c['name'] == in_class:
126             coco_labels_inverse_val[c['id']] = id
127
128 # save images and get annotations
129 trainImageBoxes = save_images(trainImages, 'train', coco_labels_inverse_train)
130 valImageBoxes = save_images(valImages, 'val', coco_labels_inverse_val)
131
132 # save annotations
133 with open("custom_dataset/train_labels.json", "w") as f:
134     json.dump(trainImageBoxes, f)
135 with open("custom_dataset/val_labels.json", "w") as f:
136     json.dump(valImageBoxes, f)
137
138 print(len(trainImages))
139 print(len(valImages))
140
141 display_class_images(trainImageBoxes, 0, class_list)
142 display_class_images(trainImageBoxes, 1, class_list)
143 display_class_images(trainImageBoxes, 2, class_list)

```

## 5.2 hw5\_network.py

```
1  '''
2  The code in this file lays out the model architecture and the training routine
3  '''
4  import numpy as np
5  import matplotlib.pyplot as plt
6  from PIL import Image
7  import os
8  import glob
9  from sklearn.metrics import confusion_matrix, accuracy_score
10 import seaborn as sns
11 import torchvision
12 import torchvision.transforms as tvf
13 import torch
14 import torch.nn as nn
15 import torch.nn.functional as F
16 import cv2
17 import json
18 from tqdm.notebook import trange, tqdm
19
20 class MyDataset(torch.utils.data.Dataset):
21     """Class to load in custom COCO dataset
22     """
23     def __init__(self, root, label_file):
24         super().__init__()
25         self.root = root
26         self.label_dict = {0:'pizza', 1:'bus', 2: 'cat'}
27
28         with open(label_file, 'r') as f:
29             self.labels = json.load(f)
30
31         # referenced https://stackoverflow.com/questions/26392336/
32         # importing-images-from-a-directory-python-to-list-or-dictionary
33         # to determine how to find all image file names
34         self.image_files = glob.glob(os.path.join(self.root, '*.jpg'))
35
36     def __len__(self):
37         return len(self.image_files)
38
39     def __getitem__(self, index):
```

```

40     # open image and get label
41     pil_img = Image.open(self.image_files[index])
42     curr_label_info = self.labels[os.path.basename(self.image_files[index])]
43     label = torch.tensor(curr_label_info['label'])
44
45     # get bbox info
46     bbox = curr_label_info['bbox']
47     [x, y, w, h] = bbox
48     bbox = torch.tensor([x,y,x+w,y+h])
49
50     # perform RGB and tensor transforms
51     if pil_img.mode != "RGB":
52         pil_img = pil_img.convert(mode="RGB")
53     transforms = tvn.Compose([
54         tvn.ToTensor()
55     ])
56     transformed_img = transforms(pil_img)
57     assert transformed_img.shape == torch.Size([3,256,256])
58
59     # normalize bbox to be in range [0,1]
60     bbox = torch.div(bbox, transformed_img.shape[1])
61     assert torch.max(bbox) <= 1
62
63     return transformed_img, label.squeeze(), bbox.squeeze()
64
65 class SkipBlock(nn.Module):
66     """Skip Connection Layer with option to downsample tensor
67     """
68     def __init__(self, in_ch, out_ch, downsample=False):
69         super(SkipBlock, self).__init__()
70         self.downsample = downsample
71         if self.downsample:
72             self.downsampler = nn.Conv2d(in_ch, out_ch, (3,3), stride=2)
73
74         self.in_ch = in_ch
75         self.out_ch = out_ch
76         self.conv1 = nn.Conv2d(in_ch, out_ch, (3,3), stride=1, padding=1)
77         self.bn1 = nn.BatchNorm2d(out_ch)
78
79         # different logic for when in_ch == out_ch
80         if self.in_ch == self.out_ch:

```

```

81     self.conv2 = nn.Conv2d(in_ch, out_ch, (3,3), stride=1, padding=1)
82     self.bn2 = nn.BatchNorm2d(out_ch)
83     # downsample layer sizes are different too
84     if self.downsample:
85         self.downsampler_out = nn.Conv2d(in_ch, out_ch, (3,3), stride=2)
86         self.downsampler_identity = nn.Conv2d(in_ch, out_ch, (3,3), stride=2)
87     else:
88         self.conv2 = nn.Conv2d(out_ch, out_ch, (3,3), stride=1, padding=1)
89         self.bn2 = nn.BatchNorm2d(out_ch)
90         # downsample layer sizes are different too
91         if self.downsample:
92             self.downsampler_out = nn.Conv2d(out_ch, out_ch, (3,3), stride=2)
93             self.downsampler_identity = nn.Conv2d(in_ch, out_ch, (3,3), stride=2)
94
95     self.relu = nn.ReLU(True)
96     return
97
98 def forward(self, x):
99     # store input for skip connection
100    identity = x
101
102    # run input through two conv and bn layers
103    out = self.conv1(x)
104    out = self.bn1(out)
105    out = self.relu(out)
106    out = self.conv2(out)
107    out = self.bn2(out)
108
109    # downsample output and identity if necessary
110    if self.downsample:
111        out = self.downsampler_out(out)
112        identity = self.downsampler_identity(identity)
113
114    # combine input with output
115    out += identity
116    out = self.relu(out)
117
118    return out
119
120 class FlattenForLinear(nn.Module):
121     """Custom module to flatten tensor for input into fc layers

```

```

122     """
123     def __init__(self):
124         super(FlattenForLinear, self).__init__()
125     def forward(self, x):
126         return x.view(x.shape[0], -1)
127
128     class HW5Net(nn.Module):
129         """Resnet-based encoder that consists of a few downsampling + several Resnet
130         blocks as the backbone and two prediction heads.
131         NOTE: THIS FUNCTION IS DERIVED FROM THE ONE PROVIDED IN THE HOMEWORK.
132         MODIFICATIONS WERE MADE IN THE NECESSARY PARTS
133         """
134         def __init__(self, input_nc, output_nc, ngf=8, n_blocks=4):
135             assert (n_blocks >= 0)
136             super(HW5Net, self).__init__()
137             # The first conv layer
138             model = [nn.ReflectionPad2d(3),
139                     nn.Conv2d(input_nc, ngf, kernel_size=7, padding=0),
140                     nn.BatchNorm2d(ngf),
141                     nn.ReLU(True)]
142
143             # Add downsampling layers
144             n_downsampling = 4
145             for i in range(n_downsampling):
146                 mult = 2 ** i
147                 model += [nn.Conv2d(ngf * mult, ngf * mult * 2
148                                     , kernel_size=3, stride=2, padding=1),
149                           nn.BatchNorm2d(ngf * mult * 2),
150                           nn.ReLU(True)]
151
152             # Add your own ResNet blocks
153             ## most of code below is all my own (code above is from assignment) ##
154             mult = 2 ** n_downsampling
155             for i in range(n_blocks):
156                 model += [SkipBlock(ngf * mult, ngf * mult, downsample=False)]
157             self.model = nn.Sequential(*model)
158
159             ##### The classification head #####
160             n_downsampling = 3
161             skip_out = ngf * mult
162             class_head = []

```

```

163
164     # use skip blocks to downsample tensor before linear layers
165     for i in range(n_downsampling):
166         mult = 2 ** i
167         class_head += [SkipBlock(skip_out * mult, skip_out * mult * 2,
168                                 downsample=True)]
169     # linear layers to get output classes
170     class_head += [FlattenForLinear(),
171                   nn.Linear(skip_out * mult * 2, int(skip_out * mult / 16)),
172                   nn.ReLU(True),
173                   nn.Linear(int(skip_out * mult / 16), output_nc)]
174
175     self.class_head = nn.Sequential(*class_head)
176
177     #### The bounding box regression head ####
178     bbox_head = []
179     n_downsampling = 3
180
181     # use skip blocks to downsample tensor before linear layers
182     for i in range(n_downsampling):
183         mult = 2 ** i
184         bbox_head += [SkipBlock(skip_out * mult, skip_out * mult * 2,
185                                 downsample=True)]
186     # linear layers to get output classes
187     bbox_head += [FlattenForLinear(),
188                   nn.Linear(skip_out * mult * 2, int(skip_out * mult / 16)),
189                   nn.ReLU(True),
190                   nn.Linear(int(skip_out * mult / 16), 4)]
191     self.bbox_head = nn.Sequential(*bbox_head)
192
193     def forward(self, input):
194         ft = self.model(input)
195         cls = self.class_head(ft.clone())
196         bbox = self.bbox_head(ft.clone())
197         return cls, bbox
198
199     def train(net, data_loader, criterion_bbox_type, device, num_epochs=20):
200         all_epoch_loss_cls = []
201         all_epoch_loss_bbox = []
202
203         net = net.to(device)

```

```

204 net.train()
205 criterion_cls = torch.nn.CrossEntropyLoss()
206 optimizer = torch.optim.Adam(
207     net.parameters(), lr=1e-4, betas=(0.9, 0.99))
208
209 # run training for all epochs (tqdm displays progress bar)
210 for epoch in tqdm(range(num_epochs), desc=" epochs", position=0):
211     # keep track of different running losses
212     running_loss_cls = 0.0
213     running_loss_bbox = 0.0
214     epoch_loss_cls = 0.0
215     epoch_loss_bbox = 0.0
216
217     # loop through data in train dataset
218     pbar = tqdm(data_loader, desc=" data loader", position=1)
219     for i, data in enumerate(pbar):
220         inputs, labels, bboxes = data
221         inputs = inputs.to(device)
222         labels = labels.to(device)
223         bboxes = bboxes.to(device)
224
225         optimizer.zero_grad()
226         outputs_cls, outputs_bbox = net(inputs)
227
228         # compute loss for classification
229         loss_cls = criterion_cls(outputs_cls, labels)
230         loss_cls.backward(retain_graph=True)
231
232         # choose correct loss function for regression
233         if criterion_bbox_type == 'MSE':
234             criterion_bbox = torch.nn.MSELoss()
235             loss_bbox = criterion_bbox(outputs_bbox, bboxes)
236         elif criterion_bbox_type == 'CIOU':
237             loss_bbox = torchvision.ops.complete_box_iou_loss(outputs_bbox, bboxes,
238                                                                 reduction='mean')
239         else:
240             assert False, 'Invalid Criterion Type for bbox'
241
242         loss_bbox.backward()
243         optimizer.step()
244

```



```

245     # update running losses
246     running_loss_cls += loss_cls.item()
247     running_loss_bbox += loss_bbox.item()
248     epoch_loss_cls += loss_cls.item()
249     epoch_loss_bbox += loss_bbox.item()
250
251     # display progress as it trains
252     pbar.set_description(f'loss_cls: {epoch_loss_cls/(i+1):.3f} \
253                        loss_bbox: {epoch_loss_bbox/(i+1):.3f}')
254     if (i+1) % 100 == 0:
255         print("[epoch: %d, batch: %5d] loss_cls: %.3f loss_bbox: %.3f" \
256              % (epoch + 1, i + 1, running_loss_cls / 100, running_loss_bbox / 100))
257         running_loss_cls = 0.0
258         running_loss_bbox = 0.0
259     all_epoch_loss_cls.append(epoch_loss_cls / (i+1))
260     all_epoch_loss_bbox.append(epoch_loss_bbox / (i+1))
261
262     return all_epoch_loss_cls, all_epoch_loss_bbox
263
264 def graph_loss(epoch_loss_cls, epoch_loss_bbox):
265     """This function graphs the loss for the classification and regression
266     """
267     num_epochs = len(epoch_loss_cls)
268     fig, ax = plt.subplots(1,2)
269     fig.set_size_inches(9.5, 5.5)
270
271     # plot classification loss
272     ax[0].plot(list(range(1,num_epochs+1)), epoch_loss_cls)
273     ax[0].set_xlabel('Num Epochs')
274     ax[0].set_ylabel('Loss')
275     ax[0].set_title('Classification Loss')
276
277     # plot regression loss
278     ax[1].plot(list(range(1,num_epochs+1)), epoch_loss_bbox)
279     ax[1].set_xlabel('Num Epochs')
280     ax[1].set_ylabel('Loss')
281     ax[1].set_title('Bounding Box Loss')
282     plt.show()
283
284 class EvaluateModel():
285     """Class to perform various evaluations on the validation set

```

```

286 """
287 def __init__(self, net, name, data_loader, epoch_loss_cls, epoch_loss_bbox, num_comp=4):
288     self.net = net
289     self.name = name
290     self.data_loader = data_loader
291     self.labels = ['pizza', 'bus', 'cat']
292     self.epoch_loss_cls = epoch_loss_cls
293     self.epoch_loss_bbox = epoch_loss_bbox
294     self.num_epochs = len(self.epoch_loss_cls)
295     self.num_comp = num_comp
296
297 def perform_inference(self):
298     self.net.eval()
299
300     self.y_true = []
301     self.y_pred = []
302     self.bbox_true = []
303     self.bbox_pred = []
304
305     self.compare_dict = {'pizza': [], 'bus': [], 'cat': []}
306
307     for i, data in enumerate(self.data_loader):
308         inputs, labels, bboxes = data
309         inputs = inputs.to(device)
310
311         # put data through model
312         outputs_cls, outputs_bbox = self.net(inputs)
313         outputs_cls = torch.argmax(outputs_cls, dim=1)
314
315         # move outputs to numpy on cpu
316         labels = labels.numpy()
317         bboxes = bboxes.numpy()
318         outputs_cls = outputs_cls.detach().cpu().numpy()
319         outputs_bbox = outputs_bbox.detach().cpu().numpy()
320         inputs = inputs.detach().cpu().numpy()
321
322         # track image, label, bbox groupings for later display
323         for label, image, true_bbox, pred_bbox, pred_label in \
324             zip(labels, inputs, bboxes, outputs_bbox, outputs_cls):
325             curr_label = self.labels[label]
326             # only save the number desired by the user

```

```

327         if len(self.compare_dict[curr_label]) < self.num_comp:
328             self.compare_dict[curr_label].append({'img': image.transpose(1,2,0).copy(),
329                                                     'pred_label': self.labels[pred_label],
330                                                     'true_bbox': true_bbox*image.shape[1],
331                                                     'pred_bbox': pred_bbox*image.shape[1]})
332
333         # track data for confusion matrix
334         self.y_true.extend(labels)
335         self.y_pred.extend(outputs_cls)
336
337         # track data for IOU metrics
338         self.bbox_true.extend(bboxes)
339         self.bbox_pred.extend(outputs_bbox)
340
341     def confusion_matrix(self):
342         # plot confusion matrix and accuracy score
343         conf_mat = confusion_matrix(self.y_true, self.y_pred)
344         acc_score = accuracy_score(self.y_true, self.y_pred)
345
346         sns.heatmap(conf_mat, cmap='Blues', annot=True, fmt='g',
347                    xticklabels=self.labels, yticklabels=self.labels)
348         plt.xlabel('Predicted Label')
349         plt.ylabel('True Label')
350         plt.title(f'Confusion Matrix for {self.name}\nAccuracy: {acc_score*100:.2f}%')
351         plt.show()
352
353     def IOU_metrics(self):
354         # calculate mean iou
355         running_iou = 0
356         for true, pred in zip(self.bbox_true, self.bbox_pred):
357             true = torch.tensor(true)
358             pred = torch.tensor(pred)
359             running_iou += torchvision.ops.box_iou(true.unsqueeze(0), pred.unsqueeze(0))
360
361         mean_iou = running_iou / len(self.bbox_true)
362         print(f'Mean IOU: {float(mean_iou.squeeze())}')
363
364     def show_bbox_inference(self):
365         # loop through saved images from validation loop
366         for label, images in self.compare_dict.items():
367             # fig, ax = plt.subplots(int(self.num_comp / 2), int(self.num_comp / 2))

```

```

368     fig, ax = plt.subplots(1, self.num_comp)
369     fig.set_size_inches(9.5, 5.5)
370     ax = ax.flatten()
371
372     # loop through image for a given class
373     for i, img_dict in enumerate(images):
374         img = img_dict['img']
375
376         # draw ground truth bbox
377         [x1_gt, y1_gt, x2_gt, y2_gt] = img_dict['true_bbox']
378         img = cv2.rectangle(img, (int(x1_gt), int(y1_gt)),
379                             (int(x2_gt), int(y2_gt)), (0,1,0), 2)
380         img = cv2.putText(img, f'truth', (int(x1_gt), int(y1_gt - 10)),
381                             cv2.FONT_HERSHEY_SIMPLEX, 0.8, (0,1,0), 2)
382
383         # draw prediction bbox
384         [x1_pred, y1_pred, x2_pred, y2_pred] = img_dict['pred_bbox']
385         img = cv2.rectangle(img, (int(x1_pred), int(y1_pred)),
386                             (int(x2_pred), int(y2_pred)), (1,0,0), 2)
387         img = cv2.putText(img, f'pred',
388                             (int(x2_pred-40), int(y1_pred - 10)),
389                             cv2.FONT_HERSHEY_SIMPLEX, 0.8, (1,0,0), 2)
390         img[img > 1] = 1
391         img[img < 0] = 1
392         ax[i].imshow(img)
393         ax[i].set_title(f'Truth: {label}\nPred: {img_dict["pred_label"]}')
394
395     plt.tight_layout()
396     plt.show()
397
398     if __name__ == '__main__':
399         device = torch.device("cuda" if torch.cuda.is_available() else "cpu")
400         print(device)
401
402         ##### MSE NETWORK #####
403         train_dataset = MyDataset('custom_dataset/train', 'custom_dataset/train_labels.json')
404         val_dataset = MyDataset('custom_dataset/val', 'custom_dataset/val_labels.json')
405
406         train_batch_size = 16
407         train_data_loader = torch.utils.data.DataLoader(train_dataset, batch_size=train_batch_size, shuffle=True)
408

```

```

409 net_mse = HW5Net(3, 3)
410 num_layers = len(list(net_mse.parameters()))
411 print(f'Num layers: {num_layers}')
412
413 # train network
414 num_epochs = 20
415 all_epoch_loss_cls, all_epoch_loss_bbox = train(net_mse, train_data_loader,
416                                                'MSE', device, num_epochs=num_epochs)
417
418 # perform validation
419 val_batch_size = 16
420 val_data_loader = torch.utils.data.DataLoader(val_dataset,
421                                                batch_size=val_batch_size,
422                                                shuffle=True)
423
424 eval = EvaluateModel(net_mse, 'MSE', val_data_loader,
425                       all_epoch_loss_cls, all_epoch_loss_bbox)
426 graph_loss(all_epoch_loss_cls, all_epoch_loss_bbox)
427 eval.perform_inference()
428 eval.confusion_matrix()
429 eval.show_bbox_inference()
430 eval.IOU_metrics()
431
432 ##### CIOU NETWORK #####
433 net_ciou = HW5Net(3, 3)
434 num_layers = len(list(net_ciou.parameters()))
435 print(f'Num layers: {num_layers}')
436
437 # train network
438 num_epochs = 20
439 all_epoch_loss_cls, all_epoch_loss_bbox = train(net_ciou, train_data_loader,
440                                                'CIOU', device, num_epochs=num_epochs)
441
442 # perform validation
443 eval = EvaluateModel(net_ciou, 'CIOU', val_data_loader,
444                       all_epoch_loss_cls, all_epoch_loss_bbox)
445 graph_loss(all_epoch_loss_cls, all_epoch_loss_bbox)
446 eval.perform_inference()
447 eval.confusion_matrix()
448 eval.show_bbox_inference()
449 eval.IOU_metrics()

```