

Workshop 28/04/2024

▼ Object Detection with a custom dataset in Pytorch

This Google Colab notebook demonstrates the process of training an object detection network using a manually annotated dataset created using [CVAT](#).

Object detection is a computer vision task that involves identifying and localizing objects within an image or video. Transfer learning, on the other hand, is a technique that allows us to reuse pre-trained models trained on large-scale datasets to solve new tasks with limited labeled data.

In this notebook, we will be using a popular deep learning framework, PyTorch, to implement transfer learning for object detection. We will start with a pre-trained model, typically trained on a large dataset such as COCO (Common Objects in Context), and fine-tune it on our specific dataset, which has been manually annotated using CVAT.

CVAT is a powerful annotation tool that enables us to label objects in images or videos with bounding boxes, polygons, or keypoints. By manually annotating our dataset, we can provide ground truth labels that are crucial for training an accurate object detection model.

The steps involved in this notebook include:

1. Preparing the dataset: Organizing the annotated dataset in a suitable format for training the model.
2. Setting up the deep learning framework and installing any required dependencies.
3. Loading the pre-trained model: Importing a pre-trained model, such as Faster R-CNN or YOLO, that has been pre-trained on a large-scale dataset.
4. Fine-tuning the model: Training the model on our annotated dataset, allowing it to learn to detect objects specific to our task.
5. Evaluating the model: Assessing the performance of the trained model on a validation set or testing dataset.
6. Inference: Applying the trained model to new images or videos to detect objects in real-world scenarios.

By following along with this notebook, you will gain a practical understanding of how to use transfer learning for object detection and how to leverage the power of manually annotated datasets created using CVAT.

▼ 0. Set up environment

This code block imports the necessary Python modules and mounts your Google Drive in Colab to access your data files.

When running this code block, you will be prompted to authorize access to your Google Drive. Follow the instructions provided and enter the authorization code to mount your drive.

The mounted Google Drive allows you to easily load and save files during your Colab session, providing convenient access to your data.

```

1 import os
2 import torch
3 import torch.utils.data
4 import torchvision
5 from PIL import Image
6 from pycocotools.coco import COCO
7 from torchvision.models.detection.faster_rcnn import FastRCNNPredictor
8 from torch.utils.data import random_split
9 from torch.utils.data import random_split
10 from tqdm import tqdm
11 from torch.utils.tensorboard import SummaryWriter
12 from torchvision.models.detection.rpn import AnchorGenerator
13 from torchvision.models.detection import FasterRCNN
14 from torch.utils.tensorboard import SummaryWriter
15
16 dataset_dir = '/content/Dataset'
17 image_data_dir = os.path.join(dataset_dir, 'images')
18 annotations_dir = os.path.join(dataset_dir, 'annotations')
19 annotations_json = os.path.join(annotations_dir, 'instances_default.json')
20
21 os.makedirs(dataset_dir, exist_ok=True)
22 os.makedirs(annotations_dir, exist_ok=True)
23 os.makedirs(image_data_dir, exist_ok=True)

```

24
25

▼ Creating Dataset and DataLoader

In this code block, we create our own custom dataset and set up a data loader for training our object detection network.



- 1. Creating the Dataset:** We instantiate `CVATDataset`, our custom dataset class, by passing the `root` directory of the dataset, the annotation file, and the desired `transforms`. This dataset class is responsible for loading images and their corresponding annotations.
- 2. Removing Images Without Annotations:** To ensure that our training dataset contains only images with relevant annotations, we call `coco_remove_images_without_annotations` function. It takes the `my_dataset` as input and removes any images that do not have valid annotations. This helps us maintain a high-quality dataset for training.
- 3. Setting the Batch Size:** We define the `train_batch_size` variable to determine the number of samples per batch during training. Adjust this value based on your available resources and the model's memory requirements.
- 4. Creating the DataLoader:** Finally, we create the `data_loader` using `torch.utils.data.DataLoader`. It takes the `dataset` object, the `batch_size`, `shuffle=True` for randomizing the order of samples, `num_workers` for parallel data loading, and `collate_fn` for defining how to collate the batch.

By creating a custom dataset and setting up the data loader, we ensure that our data is efficiently loaded and processed in batches during training, facilitating the training of our object detection network.

```

1 class CVATDataset(torch.utils.data.Dataset):
2     def __init__(self, root, annotation, transforms=None):
3         self.root = root
4         self.transforms = transforms
5         self.coco = COCO(annotation)
6         self.ids = list(sorted(self.coco.imgs.keys()))
7
8     def __getitem__(self, index):
9         # Own coco file
10        coco = self.coco
11        # Image ID
12        img_id = self.ids[index]
13        # List: get annotation id from coco
14        ann_ids = coco.getAnnIds(imgIds=img_id)
15        # Dictionary: target coco_annotation file for an image
16        coco_annotation = coco.loadAnns(ann_ids)
17        # path for input image
18        path = coco.loadImgs(img_id)[0]['file_name']
19        # open the input image
20        img = Image.open(os.path.join(self.root, path))
21
22        # number of objects in the image
23        num_objs = len(coco_annotation)
24
25        # Bounding boxes for objects
26        # In coco format, bbox = [xmin, ymin, width, height]
27        # In pytorch, the input should be [xmin, ymin, xmax, ymax]
28        boxes = []
29        for i in range(num_objs):
30            xmin = coco_annotation[i]['bbox'][0]
31            ymin = coco_annotation[i]['bbox'][1]
32            xmax = xmin + coco_annotation[i]['bbox'][2]
33            ymax = ymin + coco_annotation[i]['bbox'][3]
34            boxes.append([xmin, ymin, xmax, ymax])
35        boxes = torch.as_tensor(boxes, dtype=torch.float32)
36        # Labels (In my case, I only one class: target class or background)
37        labels = torch.ones((num_objs,), dtype=torch.int64)
38        # Tensorise img_id
39        img_id = torch.tensor([img_id])
40        # Size of bbox (Rectangular)
41        areas = []
42        for i in range(num_objs):
43            areas.append(coco_annotation[i]['area'])
44        areas = torch.as_tensor(areas, dtype=torch.float32)
45        # Iscrowd
46        iscrowd = torch.zeros((num_objs,), dtype=torch.int64)
47
48        # Annotation is in dictionary format
49        my_annotation = {}
50        my_annotation["boxes"] = boxes
51        my_annotation["labels"] = labels
52        my_annotation["image_id"] = img_id
53        my_annotation["area"] = areas
54        my_annotation["iscrowd"] = iscrowd

```

train_ratio: 0.8 
test_ratio: 0.1 
validation_ratio: 0.1 

```

55
56     if self.transforms is not None:
57         img = self.transforms(img)
58
59     return img, my_annotation
60
61     def __len__(self):
62         return len(self.ids)
63
64     def collate_fn(batch):
65         return tuple(zip(*batch))
66
67     def coco_remove_images_without_annotations(dataset, cat_list=None):
68         def _has_only_empty_bbox(anno):
69             return all(any(o <= 1 for o in obj["bbox"][2:]) for obj in anno)
70
71         def _count_visible_keypoints(anno):
72             return sum(sum(1 for v in ann["keypoints"][2::3] if v > 0) for ann in anno)
73
74         min_keypoints_per_image = 10
75
76         def _has_valid_annotation(anno):
77             # if it's empty, there is no annotation
78             if len(anno) == 0:
79                 return False
80             # if all boxes have close to zero area, there is no annotation
81             if _has_only_empty_bbox(anno):
82                 return False
83             # keypoints task have a slight different criteria for considering
84             # if an annotation is valid
85             if "keypoints" not in anno[0]:
86                 return True
87             # for keypoint detection tasks, only consider valid images those
88             # containing at least min_keypoints_per_image
89             if _count_visible_keypoints(anno) >= min_keypoints_per_image:
90                 return True
91             return False
92
93         # assert isinstance(dataset, torchvision.datasets.CocoDetection)
94         ids = []
95         for ds_idx, img_id in enumerate(dataset.ids):
96             ann_ids = dataset.coco.getAnnIds(imgIds=img_id, iscrowd=None)
97             anno = dataset.coco.loadAnns(ann_ids)
98             if cat_list:
99                 anno = [obj for obj in anno if obj["category_id"] in cat_list]
100             if _has_valid_annotation(anno):
101                 ids.append(ds_idx)
102
103         dataset = torch.utils.data.Subset(dataset, ids)
104         return dataset
105
106     def get_transform():
107         custom_transforms = []
108         custom_transforms.append(torchvision.transforms.ToTensor())
109         return torchvision.transforms.Compose(custom_transforms)
110
111     # create own Dataset
112     my_dataset = CVATDataset(root=image_data_dir,
113                             annotation=annotations_json,
114                             transforms=get_transform()
115                             )
116
117     # collate_fn needs for batch
118     dataset = coco_remove_images_without_annotations(my_dataset)
119
120     # Define the ratios for train, test, and validation sets
121     train_ratio = 0.8 # @param {type:"number"}
122     test_ratio = 0.1 # @param {type:"number"}
123     validation_ratio = 0.1 # @param {type:"number"}
124
125     # Extra check
126     total_ratio = train_ratio + test_ratio + validation_ratio
127
128     if abs(total_ratio - 1) > 0.0001:
129         print("Warning: Ratios do not add up to 1. Using default values.")
130         train_ratio = 0.8
131         test_ratio = 0.1
132         validation_ratio = 0.1
133
134     # Split the dataset into train, test, and validation sets
135     train_size = int(train_ratio * len(dataset))
136     test_size = int(test_ratio * len(dataset))

```

```

137 validation_size = len(dataset) - train_size - test_size
138
139 train_dataset, test_dataset, validation_dataset = random_split(dataset, [train_size, test_size, validation_size])
140
141 # Batch size
142 train_batch_size = 2
143
144 # Create data loaders for train, test, and validation sets
145 train_loader = torch.utils.data.DataLoader(train_dataset,
146                                             batch_size=train_batch_size,
147                                             shuffle=True,
148                                             num_workers=2,
149                                             collate_fn=collate_fn)
150
151 test_loader = torch.utils.data.DataLoader(test_dataset,
152                                           batch_size=train_batch_size,
153                                           shuffle=False,
154                                           num_workers=2,
155                                           collate_fn=collate_fn)
156
157 validation_loader = torch.utils.data.DataLoader(validation_dataset,
158                                                  batch_size=train_batch_size,
159                                                  shuffle=False,
160                                                  num_workers=2,
161                                                  collate_fn=collate_fn)
162
163 data_loaders = {
164     "Train": train_loader,
165     "Test": test_loader,
166     "Validation": validation_loader
167 }
168
169 print("Data Loader Information:")
170 print("-----")
171
172 for split, loader in data_loaders.items():
173     dataset_size = len(loader.dataset)
174     batch_size = loader.batch_size
175     num_batches = len(loader)
176
177     print(f"{split} Dataset Size: {dataset_size}")
178     print(f"{split} Batch Size: {batch_size}")
179     print(f"Number of {split} Batches: {num_batches}")
180     print("-----")
181
182

```

loading annotations into memory...

Done (t=0.00s)

creating index...

index created!

Data Loader Information:

Train Dataset Size: 25

Train Batch Size: 2

Number of Train Batches: 13

Test Dataset Size: 3

Test Batch Size: 2

Number of Test Batches: 2

Validation Dataset Size: 4

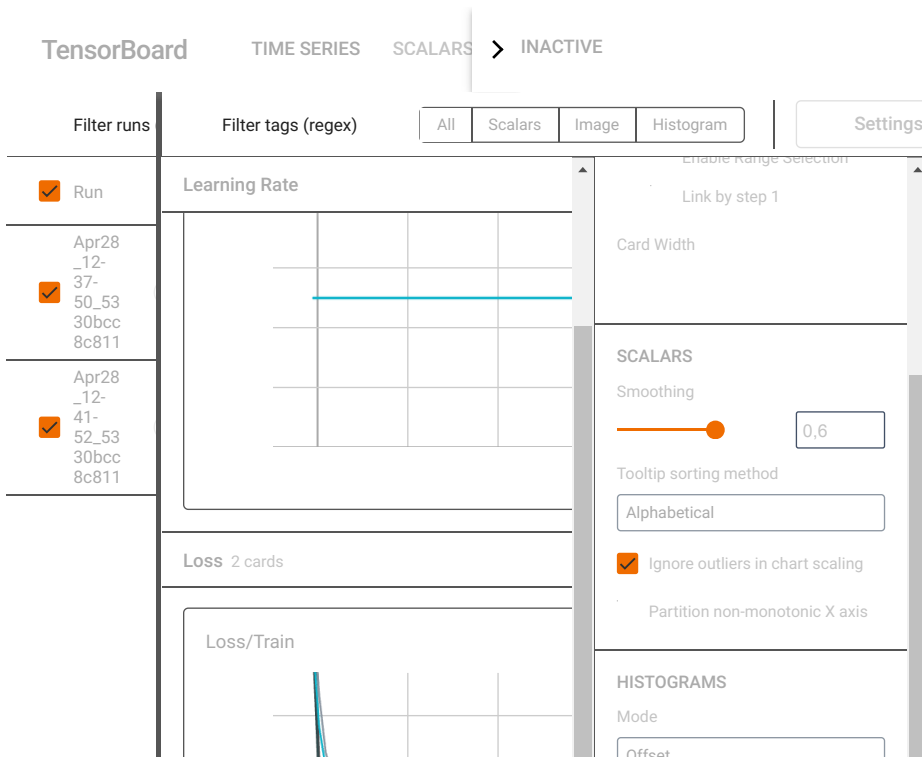
Validation Batch Size: 2

Number of Validation Batches: 2

```

1 %load_ext tensorboard
2 %tensorboard --logdir=runs
3

```

▼ Training Configuration

In this code cell, we configure the training parameters for our model.

- `num_epochs` is set to 20, which represents the number of times the entire training dataset will be passed through the model during training.
- `num_classes` is set to 2, indicating the number of classes in our classification problem. This typically includes the target class and the background class.

These parameters will be used in subsequent code cells to train a model and evaluate its performance.

```
1 num_epochs = 40 #@param {type:"number"}
2 num_classes = 2 #@param {type:"number"}
```

`num_epochs:`

`num_classes:`

The code in this cell demonstrates the training process for an instance segmentation model using the Faster R-CNN architecture.

1. **Defining the Model:** The `get_model_instance_segmentation` function is defined to create an instance segmentation model. It loads a pre-trained Faster R-CNN model and replaces the head with a new one to match the desired number of classes.
2. **Setting up the Model:** We specify the number of classes (`num_classes`) as 2 (target class or background) and the number of training epochs (`num_epochs`) as 10. The `get_model_instance_segmentation` function is called to initialize the model.
3. **Moving the Model to the Device:** The model is moved to the appropriate device (e.g., GPU) for training using the `to(device)` method.
4. **Defining Optimizer and Parameters:** The optimizer is defined as SGD (Stochastic Gradient Descent) with specific learning rate, momentum, and weight decay parameters. The trainable parameters of the model are selected using `model.parameters()` and stored in `params`.
5. **Training Loop:** The training loop begins with an outer loop over the specified number of epochs. Within each epoch, the model is set to training mode (`model.train()`). Then, an inner loop iterates over the data loader, retrieving batches of images and annotations.
6. **Moving Data to the Device:** The images and annotations are moved to the same device as the model (`imgs to device`). The annotations are converted to a list of dictionaries, ensuring all tensors are on the device.
7. **Forward Pass and Loss Computation:** The forward pass of the model is performed using `model(imgs, annotations)`, and the resulting losses are computed. The losses are accumulated and summed across all loss components.
8. **Backpropagation and Optimization:** The optimizer's gradients are reset (`optimizer.zero_grad()`), and the losses are backpropagated through the model (`losses.backward()`). Finally, the optimizer updates the model's parameters based on the computed gradients (`optimizer.step()`).
9. **Logging Progress:** Within the inner loop, the iteration number and loss value are printed to monitor the training progress.

This code cell demonstrates the training process, iterating over the data loader for a specified number of epochs, computing the loss, and updating the model's parameters.

```

1 def get_model_instance_segmentation(num_classes):
2     # load an instance segmentation model pre-trained on COCO
3     model = torchvision.models.detection.fasterrcnn_resnet50_fpn(pretrained=True)
4     # get the number of input features for the classifier
5     in_features = model.roi_heads.box_predictor.cls_score.in_features
6     # replace the pre-trained head with a new one
7     model.roi_heads.box_predictor = FastRCNNPredictor(in_features, num_classes)
8
9     return model
10
11 model = get_model_instance_segmentation(num_classes)
12
13 device = torch.device('cuda') if torch.cuda.is_available() else torch.device('cpu')
14
15 # move the model to the right device
16 model.to(device)
17
18 # parameters
19 params = [p for p in model.parameters() if p.requires_grad]
20 optimizer = torch.optim.SGD(params, lr=0.005, momentum=0.9, weight_decay=0.0005)
21 lr_scheduler = torch.optim.lr_scheduler.StepLR(optimizer, step_size=3, gamma=0.1)
22
23 # Create a SummaryWriter for TensorBoard logging
24 writer = SummaryWriter()
25
26 best_loss = float('inf')
27
28 # Define the total number of iterations
29 total_iterations = num_epochs * len(train_loader)
30
31 #####
32 # Starting the actual training
33 #####
34
35 for epoch in range(num_epochs):
36     model.train()
37     epoch_loss = 0.0 # Accumulate the loss within the epoch
38     loop = tqdm(train_loader) # Create a progress bar for training iterations
39
40     for imgs, annotations in loop:
41         imgs = list(img.to(device) for img in imgs) # Move images to the device
42         annotations = [{k: v.to(device) for k, v in t.items()} for t in annotations] # Move annotations to the device
43         loss_dict = model(imgs, annotations) # Forward pass: compute the loss
44         losses = sum(loss for loss in loss_dict.values()) # Calculate the total loss
45         epoch_loss += losses.item() # Accumulate the loss within the epoch
46
47         optimizer.zero_grad() # Clear gradients
48         losses.backward() # Backward pass: compute gradients
49         optimizer.step() # Update model parameters using gradients
50
51         loop.set_description(f"Epoch [{epoch}/{num_epochs}]") # Update progress bar description
52         loop.set_postfix(loss=losses.item()) # Update progress bar with current loss
53
54     # Calculate average loss for the epoch
55     average_loss = epoch_loss / len(train_loader)
56
57     # # Check if the average loss is decreasing and save the model
58     # if average_loss < best_loss:
59     #     best_loss = average_loss
60     #     torch.save(model.state_dict(), 'best_model.pth')
61
62     print(f'\n Epoch: {epoch + 1}/{num_epochs}, Average Loss: {average_loss}')
63
64     # Write average loss to TensorBoard
65     writer.add_scalar('Loss/Train', average_loss, epoch + 1)
66
67     # Log learning rate to TensorBoard
68     current_lr = optimizer.param_groups[0]['lr']
69     writer.add_scalar('Learning Rate', current_lr, epoch + 1)
70
71     # Validation
72     val_loss = 0.0
73     with torch.no_grad():
74         for val_imgs, val_annotations in validation_loader:
75             val_imgs = list(img.to(device) for img in val_imgs) # Move validation images to the device
76             val_annotations = [{k: v.to(device) for k, v in t.items()} for t in val_annotations] # Move validation annotations to th
77             val_loss_dict = model(val_imgs, val_annotations) # Compute validation loss
78             val_losses = sum(val_loss for val_loss in val_loss_dict.values()) # Calculate total validation loss
79             val_loss += val_losses.item() # Accumulate validation loss

```

```

70         val_loss += val_losses.item() # Accumulate validation loss
71
72     average_val_loss = val_loss / len(validation_loader)
73
74     # # Check if the average validation loss is decreasing and save the model
75     if average_val_loss < best_loss:
76         best_loss = average_val_loss
77         torch.save(model.state_dict(), 'best_model.pth')
78         print('Validation loss has improved, saving model')
79
80     print(f'Epoch: {epoch + 1}/{num_epochs}, Validation Loss: {average_val_loss} \n')
81
82     # Write validation loss to TensorBoard
83     writer.add_scalar('Loss/Validation', average_val_loss, epoch + 1)
84
85     print('\n Training completed.')
86
87 # Close the SummaryWriter
88 writer.close()
89

```

```

Epoch [0/40]: 100%|██████████| 13/13 [00:08<00:00, 1.55it/s, loss=0.219]

Epoch: 1/40, Average Loss: 0.3978698941377493
Validation loss has improved, saving model
Epoch: 1/40, Validation Loss: 0.2571355700492859

Epoch [1/40]: 100%|██████████| 13/13 [00:07<00:00, 1.66it/s, loss=0.148]
Epoch: 2/40, Average Loss: 0.22072126085941607

Validation loss has improved, saving model
Epoch: 2/40, Validation Loss: 0.23643583059310913

Epoch [2/40]: 100%|██████████| 13/13 [00:07<00:00, 1.65it/s, loss=0.15]

Epoch: 3/40, Average Loss: 0.19092857665740526
Validation loss has improved, saving model
Epoch: 3/40, Validation Loss: 0.22996443510055542

Epoch [3/40]: 100%|██████████| 13/13 [00:08<00:00, 1.46it/s, loss=0.18]
Epoch: 4/40, Average Loss: 0.18917847482057717

Validation loss has improved, saving model
Epoch: 4/40, Validation Loss: 0.17629437893629074

Epoch [4/40]: 100%|██████████| 13/13 [00:08<00:00, 1.52it/s, loss=0.0861]

Epoch: 5/40, Average Loss: 0.17212307567779833
Epoch: 5/40, Validation Loss: 0.1853400617837906

Epoch [5/40]: 100%|██████████| 13/13 [00:08<00:00, 1.62it/s, loss=0.118]
Epoch: 6/40, Average Loss: 0.14468934386968613

Validation loss has improved, saving model
Epoch: 6/40, Validation Loss: 0.1453757882118225

Epoch [6/40]: 100%|██████████| 13/13 [00:07<00:00, 1.65it/s, loss=0.0477]
Epoch: 7/40, Average Loss: 0.12882111680049163

Epoch: 7/40, Validation Loss: 0.1643618792295456

Epoch [7/40]: 100%|██████████| 13/13 [00:08<00:00, 1.57it/s, loss=0.0631]

Epoch: 8/40, Average Loss: 0.1067372104869439
Epoch: 8/40, Validation Loss: 0.19414902478456497

Epoch [8/40]: 100%|██████████| 13/13 [00:08<00:00, 1.51it/s, loss=0.0808]

Epoch: 9/40, Average Loss: 0.10013907746626781
Validation loss has improved, saving model
Epoch: 9/40, Validation Loss: 0.11546885222196579

Epoch [9/40]: 100%|██████████| 13/13 [00:08<00:00, 1.58it/s, loss=0.0569]

Epoch: 10/40, Average Loss: 0.07861400309663552
Epoch: 10/40, Validation Loss: 0.1648208238184452

Epoch [10/40]: 100%|██████████| 13/13 [00:07<00:00, 1.70it/s, loss=0.0288]

```

```

1 # Load the best model for inference
2 inference_model = get_model_instance_segmentation(num_classes)
3 inference_model.load_state_dict(torch.load('best_model.pth'))
4 inference_model.to(device)
5
6 # Prepare the new image for inference
7 image_name = '0003_hypercube.jpg' #@param {type:"string"}
8 pred_image_path = os.path.join(image_data_dir, image_name)

```

image_name: "0003_hypercube.jpg"

threshold: 0.8




threshold: 0.8






```




9 threshold = 0.8 #@param {type:"number"}
10 new_image = Image.open(pred_image_path)
11 new_image_tensor = get_transform()(new_image).unsqueeze(0).to(device)
12
13 # Set the model to evaluation mode
14 inference_model.eval()
15
16 # Perform inference on the new image
17 with torch.no_grad():
18     predictions = inference_model(new_image_tensor)
19
20
21 import matplotlib.pyplot as plt
22
23 # Function to plot bounding boxes on the image
24 def plot_bounding_boxes(image, predictions):
25     # Get the image dimensions
26     image_width, image_height = image.size
27
28     # Create a figure and axes
29     fig, ax = plt.subplots(1)
30
31     # Plot the image
32     ax.imshow(image)
33
34     # Delete axis
35     ax.axis('off')
36
37     # Plot the bounding boxes
38     for i, pred in enumerate(predictions['boxes']):
39         if predictions['scores'][i].tolist() > threshold:
40             # Get the coordinates of the bounding box
41             xmin, ymin, xmax, ymax = pred.tolist()
42
43             # Calculate the width and height of the bounding box
44             width = xmax - xmin
45             height = ymax - ymin
46
47             # Create a rectangle patch for the bounding box
48             rect = plt.Rectangle((xmin, ymin), width, height, fill=False, edgecolor='red')
49
50             # Add the rectangle patch to the axes
51             ax.add_patch(rect)
52
53     # Show the plot
54     plt.show()
55
56 # Call the function to plot bounding boxes on the new image
57 # plot_bounding_boxes(new_image, predictions[0])
58
59 listdir = os.listdir(image_data_dir)
60
61 for im in listdir:
62     threshold = 0.8 #@param {type:"number"}
63     pred_image_path = os.path.join(image_data_dir, im)
64     new_image = Image.open(pred_image_path)
65     new_image_tensor = get_transform()(new_image).unsqueeze(0).to(device)
66     with torch.no_grad():
67         predictions = inference_model(new_image_tensor)
68     plot_bounding_boxes(new_image, predictions[0])
69




```





```
/usr/local/lib/python3.10/dist-packages/torchvision/models/_utils.py:208: UserWarning: The warnings.warn(  
/usr/local/lib/python3.10/dist-packages/torchvision/models/_utils.py:223: UserWarning: Arg warnings.warn(msg)
```











































































```
1 from google.colab import drive
2 drive.mount('/content/drive')
```



```
1 os.listdir(image_data_dir)
```

```
'0002_hypercube.jpg',
'0016_hypercube.jpg',
'0057_hypercube.jpg',
'0028_hypercube.jpg',
'0055_hypercube.jpg',
'0051_hypercube.jpg',
'0015_hypercube.jpg',
'0046_hypercube.jpg',
'0070_hypercube.jpg',
'0008_hypercube.jpg',
'0041_hypercube.jpg',
'0006_hypercube.jpg',
'0039_hypercube.jpg',
'0036_hypercube.jpg',
'0029_hypercube.jpg',
'0048_hypercube.jpg',
'0001_hypercube.jpg',
'0052_hypercube.jpg',
'0021_hypercube.jpg',
'0011_hypercube.jpg',
'0013_hypercube.jpg',
'0066_hypercube.jpg',
'0064_hypercube.jpg',
'0045_hypercube.jpg',
'0058_hypercube.jpg',
'0023_hypercube.jpg',
'0014_hypercube.jpg',
'0056_hypercube.jpg',
'0009_hypercube.jpg',
'0003_hypercube.jpg',
'0060_hypercube.jpg',
'0020_hypercube.jpg',
'0047_hypercube.jpg',
'0035_hypercube.jpg',
'0032_hypercube.jpg',
'0068_hypercube.jpg',
'0065_hypercube.jpg',
'0043_hypercube.jpg',
'0024_hypercube.jpg',
'0034_hypercube.jpg',
'0026_hypercube.jpg',
'0022_hypercube.jpg',
'0018_hypercube.jpg',
'0031_hypercube.jpg',
'0030_hypercube.jpg',
'0038_hypercube.jpg',
'0033_hypercube.jpg',
'0050_hypercube.jpg',
'0040_hypercube.jpg',
'0044_hypercube.jpg',
'0017_hypercube.jpg',
'0007_hypercube.jpg',
'0067_hypercube.jpg',
'0010_hypercube.jpg',
'0069_hypercube.jpg',
'0063_hypercube.jpg',
'0005_hypercube.jpg',
'0042_hypercube.jpg']
```

✓ 1 m 15 s voltooid om 15:03



Er kan geen verbinding worden gemaakt met de reCAPTCHA-service. Controleer je internetverbinding en laad de pagina opnieuw om een reCAPTCHA-uitdaging te ontvangen.