

# Homework 11: Image Classification

Dear 109ers,

This homework assignment requires very little coding but it will take up to 3x3 hours to train convolutional neural networks (CNN) on your laptop. We shift the due date back by three days to give you a bit more time. You might let your laptop do the hard work while you sleep.

We assume you have completed this week's lab assignment and have a working code that allows you to classify arbitrary set of images by training CNNs.

From bCourses please download and unzip these files, which will require 2x2.8 GB of disk space:

```
241795449 Nov 5 12:53:02 2025 images_for_homework_bikes_and_scooters_8_0.zip
218199307 Nov 5 12:54:47 2025 images_for_homework_bikes_and_scooters_8_1.zip
226916067 Nov 5 12:56:31 2025 images_for_homework_bikes_and_scooters_8_2.zip
233071021 Nov 5 12:58:20 2025 images_for_homework_bikes_and_scooters_8_3.zip
119862601 Nov 5 12:27:07 2025 images_for_homework_buildings_8_4.zip
115042029 Nov 5 12:28:38 2025 images_for_homework_buildings_8_5.zip
109132335 Nov 5 12:30:06 2025 images_for_homework_buildings_8_6.zip
140166644 Nov 5 12:46:55 2025 images_for_homework_cars_8_0.zip
138271039 Nov 5 12:48:18 2025 images_for_homework_cars_8_1.zip
134039977 Nov 5 12:49:41 2025 images_for_homework_cars_8_2.zip
139030129 Nov 5 12:51:10 2025 images_for_homework_cars_8_3.zip
87193429 Nov 5 12:36:04 2025 images_for_homework_doors_8_7.zip
91827745 Nov 6 10:42:54 2025 images_for_homework_doors_8_8.zip
92671145 Nov 5 12:39:21 2025 images_for_homework_doors_8_9.zip
189395912 Nov 5 12:20:17 2025 images_for_homework_plants_8_4.zip
195202545 Nov 5 12:22:06 2025 images_for_homework_plants_8_5.zip
192165496 Nov 5 12:23:44 2025 images_for_homework_plants_8_6.zip
99044475 Nov 5 12:41:14 2025 images_for_homework_windows_8_7.zip
91011663 Nov 6 10:41:23 2025 images_for_homework_windows_8_8.zip
101060249 Nov 5 12:44:03 2025 images_for_homework_windows_8_9.zip
```

(1) Depending on the last digit of your student ID, you uploaded images for two categories:

ID = 0, 1, 2, or 3	Cars vs bicycles/scooters
ID = 4, 5, or 6	Plants vs buildings
ID = 7, 8, or 9	Doors vs windows

These zip files above were generated by selecting images at random from your submissions. If your images were submitted on time and their format were found to be appropriate, you will likely recognize some of your own. They were made a bit smaller and converted to the .jpg format with quality factor 6 to save you all some disk space.

From the two .zip files to which you likely contributed to (if you last ID were **6**, this would be images\_for\_homework\_plants\_8\_6.zip and images\_for\_homework\_buildings\_8\_6.zip), select **2x(200+50) images** of your choice and copy them into usual training/testing folder structure:

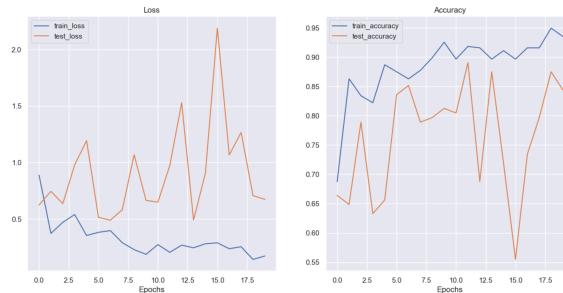
```
$dd/training_set/plants    <<< 200
$dd/training_set/buildings <<< 200
$dd/test_set/plants        <<< 50
$dd/test_set/buildings    <<< 50
```

```

5% | Epoch: 1 | train_loss: 0.7898 | train_acc: 0.7893 | test_loss: 1.7538 | test_acc: 0.3986 | 1/28 [02:16<43:11, 136.41s/it]
18% | Epoch: 2 | train_loss: 0.6324 | train_acc: 0.7885 | test_loss: 0.8036 | test_acc: 0.4844 | 2/28 [04:32<48:47, 136.88s/it]
Epoch: 3 | train_loss: 0.3475 | train_acc: 0.8558 | test_loss: 0.8557 | test_acc: 0.5469 | 3/28 [06:45<50:15, 135.82s/it]
20% | Epoch: 4 | train_loss: 0.2995 | train_acc: 0.8774 | test_loss: 0.2912 | test_acc: 0.8359 | 4/28 [08:59<50:51, 134.47s/it]
25% | Epoch: 5 | train_loss: 0.2191 | train_acc: 0.8894 | test_loss: 0.2090 | test_acc: 0.9069 | 5/28 [11:13<53:33, 134.13s/it]
30% | Epoch: 6 | train_loss: 0.1438 | train_acc: 0.8534 | test_loss: 0.4052 | test_acc: 0.8847 | 6/28 [13:26<51:16, 134.89s/it]
35% | Epoch: 7 | train_loss: 0.5179 | train_acc: 0.8078 | test_loss: 0.3212 | test_acc: 0.8986 | 7/28 [15:14<50:18, 134.16s/it]
40% | Epoch: 8 | train_loss: 0.4823 | train_acc: 0.8774 | test_loss: 0.3239 | test_acc: 0.8281 | 8/28 [17:52<50:56, 134.75s/it]
45% | Epoch: 9 | train_loss: 0.4439 | train_acc: 0.8786 | test_loss: 0.1414 | test_acc: 0.9531 | 9/28 [19:13<24:47, 135.22s/it]
50% | Epoch: 10 | train_loss: 0.3882 | train_acc: 0.8552 | test_loss: 0.3282 | test_acc: 0.8750 | 10/28 [21:28<52:32, 135.25s/it]
55% | Epoch: 11 | train_loss: 0.3474 | train_acc: 0.8822 | test_loss: 0.1144 | test_acc: 0.9862 | 11/28 [24:43<50:15, 135.89s/it]
60% | Epoch: 12 | train_loss: 0.3354 | train_acc: 0.8901 | test_loss: 0.2227 | test_acc: 0.9453 | 12/28 [26:58<50:18, 135.11s/it]
65% | Epoch: 13 | train_loss: 0.3736 | train_acc: 0.9088 | test_loss: 0.1741 | test_acc: 0.9531 | 13/28 [29:15<51:48, 135.48s/it]
70% | Epoch: 14 | train_loss: 0.2831 | train_acc: 0.8822 | test_loss: 0.2359 | test_acc: 0.8986 | 14/28 [31:47<52:45, 137.62s/it]
75% | Epoch: 15 | train_loss: 0.2127 | train_acc: 0.8956 | test_loss: 0.1008 | test_acc: 0.9069 | 15/28 [33:59<51:34, 138.84s/it]
80% | Epoch: 16 | train_loss: 0.3428 | train_acc: 0.9838 | test_loss: 0.1587 | test_acc: 0.8984 | 16/28 [36:33<50:22, 140.55s/it]
85% | Epoch: 17 | train_loss: 0.3973 | train_acc: 0.8918 | test_loss: 0.2652 | test_acc: 0.9375 | 17/28 [38:47<50:18, 141.29s/it]
90% | Epoch: 18 | train_loss: 0.2141 | train_acc: 0.9807 | test_loss: 0.2389 | test_acc: 0.8828 | 18/28 [41:49<50:43, 141.78s/it]
95% | Epoch: 19 | train_loss: 0.2193 | train_acc: 0.9159 | test_loss: 0.3854 | test_acc: 0.9297 | 19/28 [43:23<50:19, 139.27s/it]
100% | Epoch: 20 | train_loss: 0.2241 | train_acc: 0.9231 | test_loss: 0.3811 | test_acc: 0.9862 | 20/28 [45:37<50:18, 136.85s/it]
Total training time: 2737.075 seconds

```

Then start the training of your CNN with 20 epochs, which may take between 40 minutes and 3 hours (see above). **Only if one epoch takes longer than 9 minutes on your laptop, you are permitted to reduce the number of images from the standard number of 2x(200+50) so that the whole training process does not take you much longer than 3 hours.** However, if you reduce the image number even more, points will be taken off most likely. Once the training is complete, generate the usual loss/accuracy graph:



You will be asked to submit your three notebooks for this assignment but we are also asking you to submit a separate *results.docx* file to make the grading process more efficient. **Please insert loss/accuracy graph into your results file and state which two image categories it corresponds to.**

(2) Now that you have trained your CNN, we want to subject it to a real test with many more images than you have used in your test set before. To make this process efficient we want you to first **write a function** that uses the existing CNN to classify all images in a given folder and then returns two numbers that specify how many images were predicted to be in each class. Your function may look like this:

```

def ClassifyImagesInFolder(folder):
    # 1. Get all image paths (* means "any combination")
    image_path_list = glob.glob(folder+"/*")
    print(image_path_list)

    number_of_images_per_class = np.zeros((2)) # number of images identified in each class

    i = 0
    for custom_image_path in image_path_list:
        i += 1
        name = os.path.basename(custom_image_path) # name of image
        print(i, number_of_images_per_class, name)

    ...

    number_of_images_per_class[custom_image_pred_label.cpu()] += 1

    ...

    return number_of_images_per_class

```

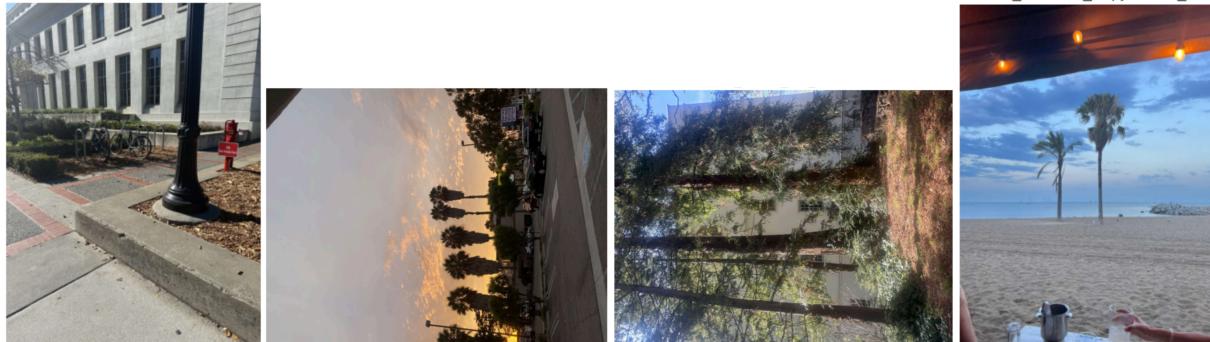
(3) For the current image classes like planets/buildings, we want to you to copy all files from the 2x2 other folders to two new folders. If your ID was 6 this would be images in folders `images_for_homework_plants_8_7`, `images_for_homework_buildings_8_8`, `images_for_homework_buildings_8_7` and `images_for_homework_plants_8_8`.

Then classify all these images by calling the function above. In your *results* files, please report how many images in each class were identified correct and what their fractions were. So you should have two outputs for this format, one for each class:

[367. 28.]  
[0.92911392 0.07088608]

(4) Modify the function above so that it automatically displays three images that were classified incorrectly. **Insert these three images into your *results* file.** Say if you see any reason or not why they were misclassified. Please also upload your final notebook.

(5) If you find any images that in your opinion should not have been placed inside any of these folders, **please report their names (one per line) in a separate text file named `problematic_images.txt` that we ask you to submit with your homework** so that we can remove them next year. Here are some problematic examples:



(6) Repeat every steps 1-5 for the two remaining pairs of images categories. If you ID was 6, this would planets vs buildings and cars vs bikes/scooters. Insert your findings into your *results* file and upload your two final notebooks as well.

(7) Add a concluding paragraph to your *results* file where you compare the results of step 3 for all three pairs of image categories. Say which classification problem was the easiest and which one was the hardest.

(If you know of a way to modify the existing Python code to make the training more robust, please contact the instructor.)