

CMP 717 – Practical 3: Convolutional Neural Networks

Due date: Tuesday, 15-12-2020, 11:59 PM.

Overview

The goal of this assignment is to get you familiarize with convolutional neural networks (CNNs). In this assignment, you will use CNNs to colorize a grayscale image..

Problem

In this assignment, we will train a convolutional neural network to solve two classic image processing tasks: image colourization and super-resolution. First, we will focus on image colourization. That is, given a greyscale image, we wish to predict the colour at each pixel. Image colourization is a difficult problem for many reasons, one of which being that it is ill-posed: for a single greyscale image, there can be multiple, equally valid colourings.

Setting Up

We recommend that you use Colab (<https://colab.research.google.com/>) for the assignment, as all the assignment notebooks have been tested on Colab. Otherwise, if you are working on your own environment, you will need to install Python 2, PyTorch (<https://pytorch.org>), iPython Notebooks, SciPy, NumPy and scikit-learn. Check out the websites of the course and relevant packages for more details.

From the assignment zip file, you will find two python notebook files: `colour_regression.ipynb`, `colourization.ipynb`. To setup the Colab environment, you will need to upload the two notebook files using the upload tab at <https://colab.research.google.com/>.

Dataset

We will use the CIFAR-10 data set, which consists of images of size 32×32 pixels. For most of the questions we will use a subset of the dataset. The data loading script is included with the handout, and should download automatically the first time it is loaded. If you have trouble downloading the file, you can also do so manually from:

<http://www.cs.toronto.edu/~kriz/cifar-10-python.tar.gz>.

To make the problem easier, we will only use the “Horse” category from this data set. Now let’s learn to colour some horses!

Problem 1: Colourization as Regression

There are many ways to frame the problem of image colourization as a machine learning problem. One naïve approach is to frame it as a regression problem, where we build a model to predict the RGB intensities at each pixel given the greyscale input. In this case, the outputs are continuous, and so squared error can be used to train the model.

In this section, you will get familiar with training neural networks using cloud GPUs. Open the notebook `colour_regression.ipynb` on Colab and answer the following questions.

1. Describe the model `RegressionCNN`. How many convolution layers does it have? What are the filter sizes and number of filters at each layer? Construct a table or draw a diagram.

2. Run all the notebook cells in `colour_regression.ipynb` on Colab (No coding involved). You will train a CNN, and generate some images showing validation outputs. How many epochs are we training the CNN model in the given setting?
3. Re-train a couple of new models using a different number of training epochs. You may train each new models in a new code cell by copying and modifying the code from the last notebook cell. Comment on how the results (output images, training loss) change as we increase or decrease the number of epochs.
4. A colour space¹ is a choice of mapping of colours into three-dimensional coordinates. Some colours could be close together in one colour space, but further apart in others. The RGB colour space is probably the most familiar to you, the model used in `colour_regression.py` computes squared error in RGB colour space. But, most state of the art colourization models do not use RGB colour space. How could using the RGB colour space be problematic? Your answer should relate how human perception of color is different than the squared distance. You may use the Wikipedia article on color space to help you answer the question.
5. How does framing colourization as a classification problem alleviate the above problem?

Problem 2: Colourization as Classification

We will select a subset of 24 colours and frame colourization as a pixel-wise classification problem, where we label each pixel with one of 24 colours. The 24 colours are selected using k-means clustering² over colours, and selecting cluster centers. This was already done for you, and cluster centers will be downloaded automatically by the notebook. For simplicity, we still measure distance in RGB space. This is not ideal but reduces the software dependencies for this assignment.

Open the notebook `colourization.ipynb` on Colab and answer the following questions.

1. Complete the model `CNN` in `colourization.ipynb`. This model should have the same layers and convolutional filters as the `RegressionCNN`, with the exception of the output layer. Continue to use PyTorch layers like `nn.ReLU`, `nn.BatchNorm2d` and `nn.MaxPool2d`, however we will not use `nn.Conv2d`. We will use our own convolution layer `MyConv2d` included in the file to better understand its internals.
2. Run `main training loop of CNN` in `colourization.ipynb` on Colab. This will train a CNN for a few epochs using the cross-entropy objective. It will generate some images showing the trained result at the end. How do the results compare to the previous regression model?

Problem 3: Skip Connections

A skip connection in a neural network is a connection which skips one or more layer and connects to a later layer. We will introduce skip connections.

1. Add a skip connection from the first layer to the last, second layer to the second last, etc. That is, the final convolution should have both the output of the previous layer and the initial greyscale input as input. This type of skip-connection is introduced by Ronneberger et al. (MICCAI 2015), and is called a “UNet”. Following the CNN class that you have completed, complete the `__init__` and `forward` methods of the `UNet` class.
Hint: You will need to use the function `torch.cat`.

¹https://en.wikipedia.org/wiki/Color_space

²https://en.wikipedia.org/wiki/K-means_clustering

2. Train the "UNet" model for the same amount of epochs as the previous CNN and plot the training curve using a batch size of 100. How does the result compare to the previous model? Did skip connections improve the validation loss and accuracy? Did the skip connections improve the output qualitatively? How? Give at least two reasons why skip connections might improve the performance of our CNN models.
3. Re-train a few more "UNet" models using different mini batch sizes with a fixed number of epochs. Describe the effect of batch sizes on the training/validation loss, and the final image output.

Problem 4: Super-Resolution

Many classic image processing problems are to transform the input images into an output image via a transformation pipeline, e.g. colourization, denoising, and super-resolution. These image processing tasks share many similarities, where the inputs are lower quality images and the outputs are the restored high-quality images. Instead of hand-design the transformations, one approach is to learn the transformation pipeline from a training dataset using supervised learning. Previously, you have trained conv nets for colourization. In this question, you will use the same conv net models to solve super-resolution tasks. In the super-resolution task, we aim to recover a high-resolution image from a low-resolution input.

1. Take a look at the data `process` function process. What is the resolution difference between the downsized input image and output image?
2. Bilinear interpolation³ is one of the basic but widely used resampling techniques in image processing. Run `super-resolution` with both CNN and UNet. Are there any difference in the model outputs? Also, comment on how the neural network results (images from the third row) differ from the bilinear interpolation results (images from the fourth row). Give at least two reasons why conv nets are better than bilinear interpolation.

Problem 5: Visualizing Intermediate Activations

We will visualize the intermediate activations for several inputs. Run the visualization block in the `colourization.ipynb` that has already been written for you. For each model, a list of images will be generated and be stored in `cmp717/pa3/outputs/model_name/act0/` folder in the Colab environment. You will need to use the left side panel (the "Table of contents" panel) to find these images under the Files tab.

1. Visualize the activations of the CNN for a few test examples. How are the activation in the first few layers different from the later layers? You do not need to attach the output images to your writeup, only descriptions of what you see.
2. Visualize the activations of the colourization UNet for a few test examples. How do the activations differ from the CNN activations?
3. Visualize the activations of the super-resolution UNet for a few test examples. Describe how the activations differ from the colourization models?

Problem 6: Conceptual Problems

1. We also did not tune any hyperparameters for this assignment other than the number of epochs and batch size. What are some hyperparameters that could be tuned? List five.

³https://en.wikipedia.org/wiki/Bilinear_interpolation

2. In the `RegressionCNN` model, `nn.MaxPool2d` layers are applied after `nn.ReLU` activations, comment on how the output of CNN changes if we switch the order of the max-pooling and ReLU?
3. In `colourization.ipynb`, we have trained a few different image processing convolutional neural networks on input and output image size of 32×32 . In the test time, the desired output size is often different than the one used in training. Describe how we can modify the trained models in this assignment to colourize test images that are larger than 32×32 .

Grading

The assignment will be graded out of 4: 0 (*no submission*), 1 (*an attempt at a solution*), 2 (*a partially correct solution*), 3 (*a mostly correct solution*), 4 (*a correct solution*), 5 (*a particularly creative or insightful solution*).

What to Hand In

You are required to submit all your code along with a report in HTML format. The codes you will submit should be well commented. Your report should be self-contained and should contain a brief overview of the problem, the details of your implemented solution and your comments about the experimental analysis on the aforementioned issues. Finally, prepare a ZIP file named `name-surname(s)-pa3.zip` containing

- `README.txt` - text file containing anything about the assignment that you want to tell but is not appropriate for the writeup webpage
- `code/` - directory containing all your code for this assignment
- `html/` - directory containing all your html report for this assignment, including images

The ZIP file will be submitted via email to `erkut@cs.hacettepe.edu.tr`.

Late policy

You may use up to five *extension* days (in total) over the course of the semester for the three programming assignments. Any additional unapproved late submission will be weighted by 0.5.

Academic Integrity

All work on assignments must be done individually unless stated otherwise. You are encouraged to discuss with your other classmates about the given assignments, but these discussions should be carried out in an abstract way. That is, discussions related to a particular solution to a specific problem (either in actual code or in the pseudocode) will not be tolerated. In short, turning in someone else's work, in whole or in part, as your own will be considered as a violation of academic integrity. Please note that the former condition also holds for the material found on the web as everything on the web has been written by someone else.