Exercises 13: Image analysis with deep learning

Introduction to data science and Python programming for medical research (2025)

For this exercise you will need to install the pytorch package.
You should install the CPU version.
Note that the package is no longer distributed with conda.
So to install it you will have to use pip.
`https://pytorch.org/`
Download the example code and the data from the course website.

1 UNet theory

The UNet is a common deep learning architecture for image analysis and image segmentation.

It is based on convolutional layers with successive downsampling in the encoding part of the network and upsampling in the decoding part. Between these two parts the layers also have so called skip connections. The network is named from the U-like shape it has when drawn schematically.

Look at the implementation in the class `UNet` in `unet.py` and in Figure 1 and be prepared to answer the following questions as well as you can:

1. How are the skip connections implemented in the Python code?
2. How many convolutions are performed in the first layer? Each convolution will result in an image, how many images are output from this layer? How many do the second layer expect as input?
3. The kernel size is set to 3 for all convolutional layers. How many parameters do you think the first layer, `self.conv1_down` has?
4. What is the purpose of the pooling operations after the convolutions in the first part of the network?

2 UNet results comparison

Run the `exercises-13-unet-segmentation-comparisons.py` script.

It will load two pre-trained UNet models, calculate evaluation scores and metrics for both as well as plot three resulting examples for both. The plots will be saved as two files in the same directory as the script.

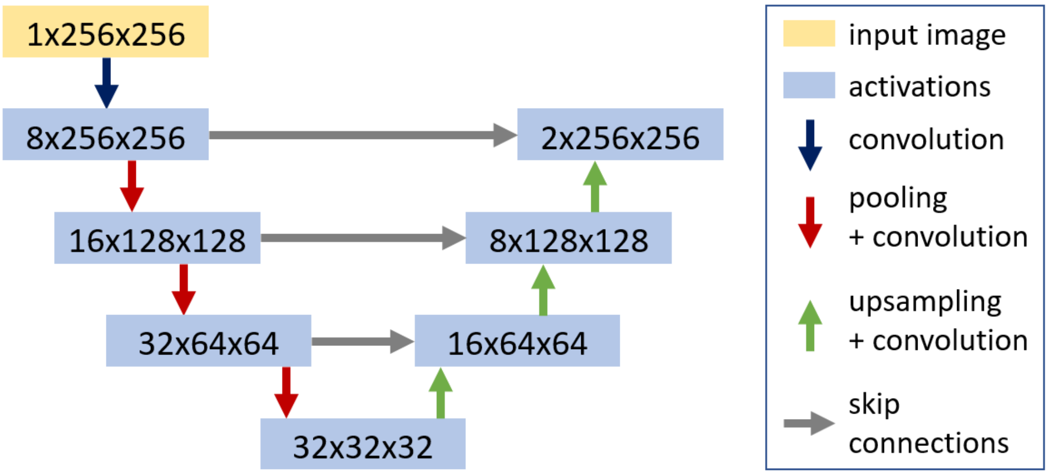

Figure 1: The UNet architecture used in this example. Original source: https://github.com/MachineLearningBiomedicalApplications

Note: If you are using a notebook you can try to change the `show_plots` variable at line 52 to `True` to have the plots show in your environment.

Inspect the output and the two plots and be prepared to answer the following questions as well as you can:

1. Based on the scores reported, which of the two models seem to perform the best?
2. What is the difference between the scores reported in the print output and at the bottom of the two plots? (How can we have *e.g.* a precision score for one *sample*/image mask?)
3. The precision, recall and F1 scores tells us about the *global* discrepancy between the ground truth segmentation and the prediction output by the UNet model. What do the Hasudorff distance tell us?

3 Links

Original source for the example, UNet code and data:
Machine Learning Biomedical Applications, Kings College, London
https://github.com/MachineLearningBiomedicalApplications
https://github.com/MachineLearningBiomedicalApplications/notebooks/tree/main/Chapter%2011%20-%20Convolutional%20networks
https://www.youtube.com/@emmarobinson8262/playlists
https://www.youtube.com/watch?v=n9VHABwn9Y0&list=PLc95HdN2_wt-D04bu96Md0TeUkIJn0jV0