Unil-CHUV Team, Lausanne, Switzerland, Task 1

Pre-processing:

1. If the subjects are still zipped, the first function that is called is *unzip_subjects* which loops over all the subjects in the input path, unzips them and deletes the zip files. If instead the subjects are already unzipped, the function does nothing.
2. Brain Extraction: the fsl Brain Extraction Tool is performed on all subjects in the input path via the function *skull_strip_subs*. The skull-stripped TOF volume is saved in the "pre" folder.
3. Registration: a statistical vessel atlas (Mouches et al.) is registered from MNI to TOF space with ANTs. Since a direct registration would have been too difficult, we decomposed the MNI → TOF registration into MNI → STRUCT (either T1 or T2), followed by STRUCT → TOF. The registered statistical atlas was used to only extract training patches in areas of the brain where there are vessels (see more below). The whole registration operation is carried out with the function *mni_2_tof_registration.*

Analysis:

We addressed the problem of aneurysm detection first with a patch-wise and then with a patient-wise approach.

- Patch-wise analysis: we trained a 3D-Unet to segment 32x32x32-wide patches with and without aneurysms. The sampling of negative patches was not random within the patients' brains. On the contrary, we extracted from all subjects only negative patches either in correspondence of precise points in the brain vasculature recurrent for aneurysms (24 **landmark points**) or just in correspondence of vessels. The **landmark points**' locations were selected by one radiologist of our institute in MNI space and they were registered to each subject using the registration parameters of the above-mentioned registration (see Preprocessing 3.). Positive patches were instead extracted around the provided masks (ground truth) in a non-centered fashion, but always ensuring that the whole mask was included in the patch. Extensive data augmentation was performed on positive patches in order to balance the two classes, namely rotations, horizontal and vertical flipping and contrast adjustement. In addition, some in-house (from our hospital) patients were also added.
- Patient-wise analysis: we create a U-Net identical to the one from the patch-wise analysis and we load its trained parameters. Then, we scan each new subject in a "sliding window" approach: in other words, we loop over all the skull-stripped patient's brain with non-overlapping 32x32x32 patches. Since most of the patches do not include vessels (and therefore aneurysms), we only retain patches that:
    - Have reasonable min and mean distances from at least one of the **landmark points** (the distance ranges were computed empirically)
    - Have reasonable mean intensity (i.e. are not too dark). Again, the intensity values were computed empirically by looking at the mean intensity of positive patches.

  In this way, all the patches that are retained at the end of the sliding window approach are reasonable candidates for containing aneurysm (provided that the registration was successful). The retained patches are then fed to the trained U-Net which outputs probabilistic volumes out of them. These probabilistic volumes are thresholded (again, the threshold was chosen empirically) and all the patches that are non-empty after the binarization are considered positive patches (i.e. with aneurysm). The center of the lesion is extracted by finding the center of mass of the largest connected component of the binarized volume. To reduce the number of false positives (FP), we set a max number of allowed FP per patient. When more than max aneurysms are predicted, we only keep the most probable by looking at the average brightness of the connected components. Also, if two (or more) adjacent patches predict aneurysms, we average the centers because it sometime happened that an aneurysm was split among neighboring patches.