Segmenting liver scans using Deep Learning

UPPSALA UNIVERSITET

Project in Computational Science

Fall Term 2019

Project members:

Dag Lindgren

Lowe Lundin

Andreas Wallin

Supervisor:

Camilla Englund

Carl Sjöberg

Aim

The aim of our project is to build a deep learning pipeline and to train models that segments livers from MRI scans. These can then be used to calculate liver volume and fat fraction values. The end-goal-product is one that can be fed with scans previously unseen from different cameras and segment it to a degree of satisfaction high enough for shipping after minor or no interference from a human operator.

Background

When trialing new a medicine, there are multiple ways of evaluating the effects. Seeing how it effects the inner organs is one very important aspect, which is often done using MRT- or CT-scans. This is one the services Antaros Medical provides and our project aims to enhance their ability to do so.

For our project the two interesting measures were the fat fraction of the liver as well as the volume of the liver. It is considered unhealthy having a fat-fraction over 5.5%, which is a condition called "fatty liver disease". This will cause an enlarged liver, which is also unhealthy. A medicine trying to alleviate/cure this condition will want to figure out how the liver changes before and after taking the medicine.

When calculating the fat fraction it is very important not to include anything outside of the liver. Therefore two models were trained, one which tries to segment the volume correctly and one which is much more conservative.

The Data

Two datasets of MRI images captured with different types of cameras (see the difference in quality between figure 4 and 6) was used to train the models. They consisted of 76 and 67 unique patients respectively. 11% of the data was set aside for testing and was only used during final run when all hyper-parameters had been tuned. All patients have four types of scans and two segmentations. Three types were used to train the fat fraction model and one type was used to train the volume model.

The Model

The model we used was based on the U-Net architecture, which is a Fully Convolutional network containing long skip connections. Local information such as texture can be passed through the long skip connections, and high-level features such as shapes can be extracted deeper down in the U-Net. The model was fed 1 or 3 types of scans depending on if it was a liver volume or a fat fraction model.

To kickstart the training we initialized the model with weights from a model that was trained on the huge ImageNet dataset.

The U-Net has a major limitation in that it uses 2D information and hence misses out on potential patterns in the third dimension. To not waste this information, we fed our networks with neighboring slices of the slice that is currently being segmented.

We also used elastic deformation, which in contrast to many other augmentation techniques can create novel, believable shapes of the organs. We also added short skip connections to allow the network to bypass convolutional blocks as it pleases.

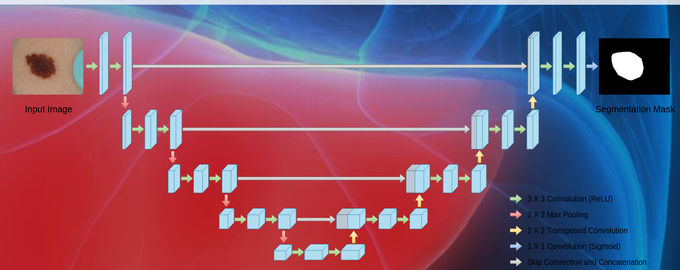

Figure 1, the computational structure of the original U-net

Results

The fat fraction model managed to predict all but two of the 16 patients in the test data to the correct fat fraction value. The ones where it differed, had the smallest possible difference of one percentage point. While the volume model predicted All patient's liver volume within or right at the 5% between the predicted volume and the ground truth value difference that is acceptable. The results of the volume model is shown in figure 2-7.

Figure 2, The volume model evaluated on the test data while training.

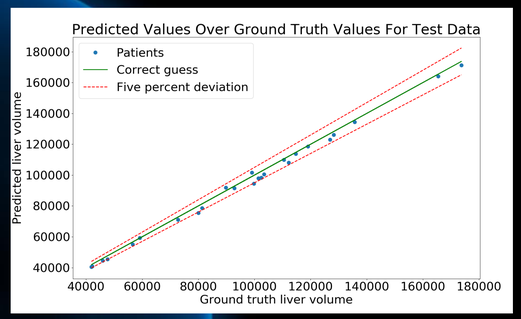

Figure 3, the predicted and ground truth liver volume of the patients in the test dataset. A difference of 5% between the predicted volume and the ground truth value is considered acceptable.

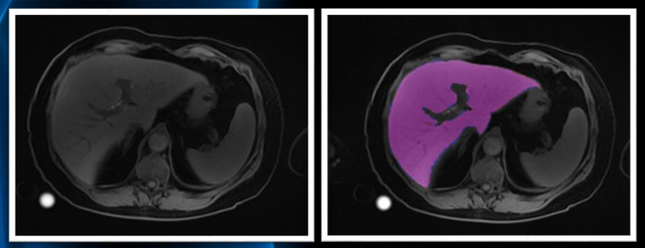

Figure 4 and 5, the two figures above show the unsegmented MRI scan on the left and the proposed liver volume segmentation from our model. Where purple is correct, red is under-segmentation and blue is over-segmentation. Here the proposed segmentation overlaps almost perfectly.

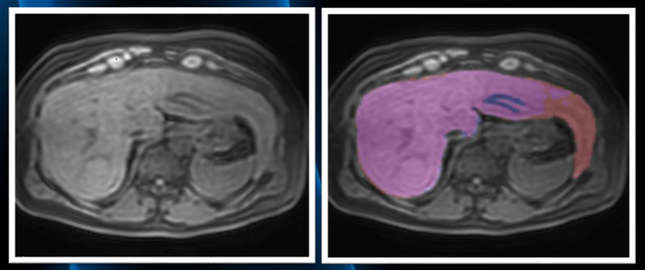

Figure 6 and 7, the two figures above show the unsegmented MRI scan on the left and the proposed liver volume segmentation from our model. Where purple is correct, red is under-segmentation and blue is over-segmentation. Notice how the proposed segmentation is missing most of the right part of the liver because of its unusual shape.

Conclusions

The models that we developed where successful in estimating the fat fraction and volumes to a degree where little or no correction is needed before shipping.

Antaros Medical