Bone Segmentation in 3D Knee MRI Images Using U-Net

Tomer Y Alon, Juan Shan, Maolin Zhang, Joe Delvecchio, Ming Zhang

Abstract— **Many diagnostic applications critically depend on the successful localization of bone structure. In this work, we trained an end-to-end U-Net convolutional neural network to perform bone segmentation on 3D knee magnetic resonance images. Without any preprocessing or post-processing of the images, the trained model generated very promising segmentation results with Dice coefficient 96.5% on the testing dataset.**

I. Data & Method

Our dataset contains 99 cases of 3D magnetic resonance images (MRI) of the knee. The images were downloaded from the public OAI database [1]. Each case included 160 2D slices with size 384x384 pixels. We selected all the slices that contained bone by marking the starting slice where bone began to show and the ending slice where bone disappeared, and used all the images between them. In total, there were 11,701 DICOM images selected from the 99 cases. Figure 1(a) provides an example of DICOM images. For each image, the bone area was manually delineated, as shown in Figure 1(b). Using a MATLAB script, we produced binary mask image for every labeled DICOM image, as shown in Figure 1(c). Every raw DICOM image was paired with its corresponding mask image.

The deep learning model employed in this work is U-Net, which is a special U-shaped convolutional neural network designed for biomedical images [2]. The advantage of U-Net is that it can generate high segmentation accuracy with limited training data, which is suitable for many medical image related tasks.

To prepare training and testing data for the U-Net model, the 99 cases were randomly divided into three distinct sets: 70% for training, 15% for validation, and 15% for testing. The testing set was held back until the end of the study. Be noted that the data separation was done at case level, i.e., all images from the same knee MRI case belong to the same set.

The code was run on the PyCharm integrated development environment (IDE) using TensorFlow, while utilizing the power of the graphics processing unit (GPU) NVIDIA Titan X to speed up the process. The training process took about 3 hours and 45 minutes with 25 epochs. With the Keras Application Programming Interface (API), we were able to implement a callback function in the code, called EarlyStopping, which stops training the model if a monitored quantity – such as the Dice coefficient – no longer improves after a number of epochs. This saved us a significant amount of time, and avoided overfitting at the same time.

II. Results

The peak performance of the validation set in Dice coefficient was 0.971– recorded at the 20^{th} epoch. The metric did not improve thereafter, and thus the training stopped at 25^{th} epoch. Upon finishing the training, the model was evaluated on the testing set, generating average Dice as 0.965. Figure 1(d) shows an example of the model's output after 25 epochs.

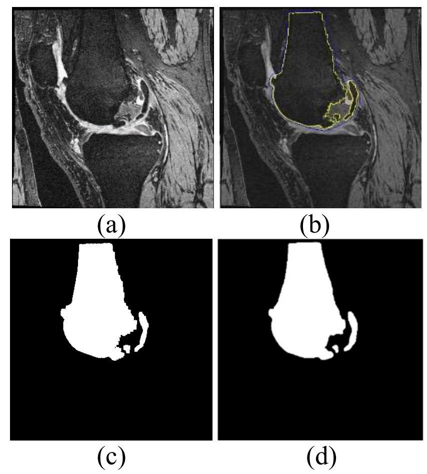

Figure 1. (a) Raw image. (b) Manual segmentation. (c) Mask image generated from manual segmentation. (d) Segmentation result from U-net.

As the results of the testing dataset show, the model's prediction was accurate for most cases. We encountered some predictions that included parts of the original image that were not of the actual bone. Such noise regions only appeared in a few of the predictions and can be handled by post-processing.

III. Conclusions

Our U-Net model for bone segmentation achieved very remarkable results. This preliminary study can serve as the critical step for segmenting other knee structures, such as cartilage, bone marrow lesion, effusion, and meniscus. Without bone identification, direct segmentation of these complicated structures from MRI is a more challenging task.

References

[1] https://oai.nih.gov

[2] O. Ronneberger, P. Fischer, and T. Brox, "U-net: Convolutional networksfor biomedical image segmentation," *International Conference on Medical Image Computing and Computer-assisted Intervention*, pp. 234–241, Springer, 2015.

T. Alon and J. Shan are with the Department of Computer Science, Seidenberg School of CSIS, Pace University, New York, NY 10038 USA (email: ta28815n@pace.edu, jshan@pace.edu).

Maolin Zhang is with Hopewell Valley Central High School, Pennington, NJ 08534 USA (email: maolinzhang@hvrsd.org).

J. Delvecchio and Ming Zhang are with Wentworth Institute of Technology, Boston, MA 02115. Ming Zhang is also with Division of Rheumatology, Tufts Medical Center, Boston, MA 02111 USA (email: delvecchioj@wit.edu, zhangm1@wit.edu)