

Building medical image classifiers with very limited data using segmentation networks[☆]

Ken C. L. Wong, Tanveer Syeda-Mahmood, Mehdi Moradi*

IBM Research – Almaden Research Center, San Jose, CA, USA

Abstract

Deep learning has shown promising results in medical image analysis, however, the lack of very large annotated datasets confines its full potential. Although transfer learning with ImageNet pre-trained classification models can alleviate the problem, constrained image sizes and model complexities can lead to unnecessary increase in computational cost and decrease in performance. As many common morphological features are usually shared by different classification tasks of an organ, it is greatly beneficial if we can extract such features to improve classification with limited samples. Therefore, inspired by the idea of curriculum learning, we propose a strategy for building medical image classifiers using features from segmentation networks. By using a segmentation network pre-trained on similar data as the classification task, the machine can first learn the simpler shape and structural concepts before tackling the actual classification problem which usually involves more complicated concepts. Using our proposed framework on a 3D three-class brain tumor type classification problem, we achieved 82% accuracy on 191 testing samples with 91 training samples. When applying to a 2D nine-class cardiac semantic level classification problem, we achieved 86% accuracy on 263 testing samples with 108 training samples. Comparisons with ImageNet pre-trained classifiers and classifiers trained from scratch are presented.

Keywords: image segmentation, image classification, deep learning, transfer learning, convolutional neural network, fully convolutional network

1. Introduction

With the availability of big data and GPU computing, deep learning has become a powerful technique which raises the benchmarks of classification and detection challenges in computer vision (Schmidhuber, 2015). Using huge databases such as ImageNet which has more than a million annotated images (Deng et al., 2009), deep convolutional neural networks (CNNs) such as AlexNet (Krizhevsky et al., 2012), VGGNet (Simonyan and Zisserman, 2014), and GoogLeNet (Szegedy et al., 2015) were proposed with impressive performance

in visual pattern recognition. In medical imaging, however, such large datasets are usually unavailable especially in 3D analysis. Given the large variety of possible clinical conditions in an imaging modality, it is very challenging to build a sufficiently large dataset with samples of desired abnormalities, and the datasets are usually highly unbalanced. If such datasets are directly used to train classifiers, low classification performance can be expected. Furthermore, different from the annotation tasks of ImageNet, annotations of medical images require radiologists or trained experts to ensure the clinical relevance. Therefore, medical image annotations can be time consuming and expensive. As a consequence, it is highly beneficial if we can reduce the number of annotated samples without decreasing the performance of the classification task.

[☆]This paper was accepted by Medical Image Analysis. The publication is available via <https://doi.org/10.1016/j.media.2018.07.010>.

*Corresponding author.

Email addresses: clwong@us.ibm.com (Ken C. L. Wong), stf@us.ibm.com (Tanveer Syeda-Mahmood), mmoradi@us.ibm.com (Mehdi Moradi)

1.1. Related work

To address the above difficulties, transfer learning using ImageNet pre-trained CNNs has been a common approach in medical image analysis (Litjens et al., 2017). In van Ginneken et al. (2015), the 4096 off-the-shelf features from the first fully-connected layer of a pre-trained OverFeat model were extracted for pulmonary nodule detection from computed tomography (CT) images through a linear support vector machine (SVM). In Moradi et al. (2016), hand-crafted features such as the histogram of oriented gradients were combined with features from selected convolutional and fully-connected layers of a pre-trained CNN-M model for cardiac semantic level classification from computed tomography angiography (CTA) images using a SVM. In Tajbakhsh et al. (2016), four different medical imaging applications on three imaging modalities were used to study the differences between CNNs trained from scratch and CNNs fine-tuned from pre-trained models. Using the AlexNet architecture, this study showed that the fine-tuned CNNs performed better as the number of training samples reduced. In Shin et al. (2016), by studying different CNN architectures such as AlexNet, GoogLeNet, and VGGNet on thoraco-abdominal lymph node detection and interstitial lung disease classification, it was shown that models trained from scratch or fine-tuned from ImageNet pre-trained CNNs consistently outperform applications using off-the-shelf CNN features.

Recently, algorithms for building medical image applications on limited training data without ImageNet pre-trained CNNs have been studied. In Roy et al. (2017), auxiliary labels generated by an existing automated software tool were used to pre-train a segmentation network. The pre-trained network was fine tuned by error corrective boosting which selectively focuses on classes with erroneous segmentations. In Uzunova et al. (2017), by learning a statistical appearance model from few sample images, a large number of synthetic image pairs were generated with the associated ground-truth deformation. These synthetic data were used to fine-tune a pre-trained FlowNet for dense image registration, and it was shown that data-driven, model-based augmentation approach outperforms generic but highly unspecific methods. In Madani et al. (2018b), a semi-supervised learning architecture based on the generative adversarial nets was proposed for cardiac abnormality classification in chest X-rays. This architecture allows both labeled

and unlabeled data to contribute to the model thus the required number of labeled data can be largely reduced.

There are several limitations of using ImageNet pre-trained CNNs on medical image analysis. First of all, if the features of the fully-connected layers are used, the input images need to be resized to match the training images (e.g. $224 \times 224 \times 3$ for VGGNet). This leads to unnecessary increase in computation for smaller images, or reduced details for larger images. Secondly, the size of the pre-trained model may be unnecessarily large for medical image applications. Using VGGNet as an example, its architecture was proposed to classify 1000 classes of non-medical images. Such a large number of classes is uncommon in medical image analysis and thus such a large model may be unnecessary. Thirdly, to the best of our knowledge, there are no publicly available models equivalent to ImageNet pre-trained CNNs for 3D image analysis. Datasets with millions of annotated 3D medical images are publicly unavailable, and the computational cost of training or using an equivalent model in 3D can be computationally difficult.

1.2. Our framework

To overcome the above limitations, here we propose to use a pre-trained segmentation network as the feature source to reduce the number of samples required to train a classification network. By using a segmentation network pre-trained on data similar to those of the classification task, we can achieve high classification performance with very limited training data. The power of this framework can be partly explained by curriculum learning (Bengio et al., 2009). Similar to the learning of humans and animals, machine learning can benefit from presenting concepts which are meaningfully organized from simple to complex. In our framework, instead of a classification task which involves complex and abstract concepts such as disease categories, we first train the machine to perform a segmentation task which involves simpler concepts such as shapes and structures. This is similar to radiologists who need to learn anatomical and functional concepts before being able to perform diagnosis. In fact, when pixel-level semantic labels are provided by radiologists or automated software tools for training, the segmentation network learns the anatomical knowledge from experts.

There are several reasons to use a segmentation network as the feature source. First of all, differ-

ent 2D and 3D segmentation networks are available (Ronneberger et al., 2015; Milletari et al., 2016; Mehta and Sivaswamy, 2017; Dou et al., 2017), and all of them show impressive capabilities of segmenting complicated biological or medical structures with relatively few training samples and straightforward training strategy. In Zhou et al. (2016), it is shown that classification networks trained on image-level labels can provide useful localization information. For segmentation networks, there are tens of thousands of pixel-level labels providing the shape, structural, and semantic information per image. Therefore, the amount of useful features in a segmentation network can be large even with relatively few training samples.

Secondly, compared with generative models such as the generative adversarial nets (Goodfellow et al., 2014; Mirza and Osindero, 2014; Pathak et al., 2016; Madani et al., 2018a,b), the trainings of segmentation networks are more straightforward and require less data. Furthermore, using features from pre-trained segmentation networks is as simple as using those from pre-trained classification networks. Although training segmentation networks requires pixel-level labels, as our concentration is not on segmentation accuracy, semantic labels generated by existing software tools or even non-semantic labels generated by techniques such as intensity thresholding can provide useful features depending on the complexities of the classification problems.

Using our framework on a three-class brain tumor classification problem of 3D magnetic resonance (MR) images, we achieved 82% accuracy on 191 testing samples with 91 training samples. When applying to a nine-class semantic level classification problem of 2D cardiac CTA images, we achieved 86% accuracy on 263 testing samples with 108 training samples.

The preliminary results of this work were reported in Wong et al. (2017). Further improvements and experiments are reported in this paper:

- The segmentation network in Mehta and Sivaswamy (2017) is utilized which provides more features than the segmentation network in Ronneberger et al. (2015) used in Wong et al. (2017). A new classification framework is proposed to utilize the features with reduced model size and complexity. Image augmentations are also introduced.
- Experiments were performed on 3D brain MR

images and 2D cardiac CTA images, not only on disease detection but also on cardiac semantic level classification. This shows that our framework can be applied to different image dimensions, modalities, and classification tasks. The classification problems and reported data in Wong et al. (2017) were entirely different.

- For the cardiac CTA images, instead of training the segmentation network using 2D slices at a particular viewpoint (e.g. the four chamber view), the network is trained on all axial slices of the 3D images. This provides more useful features for complicated tasks such as the cardiac semantic level classification.
- Apart from a classification network trained from scratch, we also compare our framework to another network which uses features from ImageNet pre-trained VGGNet. This compares our framework with one of the commonly used transfer learning techniques.

In this paper, the network architectures of the proposed and tested frameworks and the corresponding training strategy are presented in Section 2. Section 3 presents the experiments on brain tumor classification and Section 4 presents the experiments on cardiac semantic level classification. Section 5 discusses the findings from the experiments and Section 6 provides the conclusion.

2. Methodology

Our classification framework includes a pre-trained segmentation network. The architecture of the segmentation network and the architectures of the proposed and tested classification frameworks are presented.

2.1. Architecture of segmentation network

The segmentation network in Mehta and Sivaswamy (2017) (M-Net) is modified to serve as a feature source for image classification (Fig. 1). While the network was originally proposed for segmenting 3D images by stacking 2D segmentations, the modified architecture can be used for both 2D and 3D segmentation. The M-Net is modified from the U-Net (Ronneberger et al., 2015), which is a fully convolutional network that works with few training samples. Apart from the encoding and decoding paths of the U-Net, the M-Net also has two

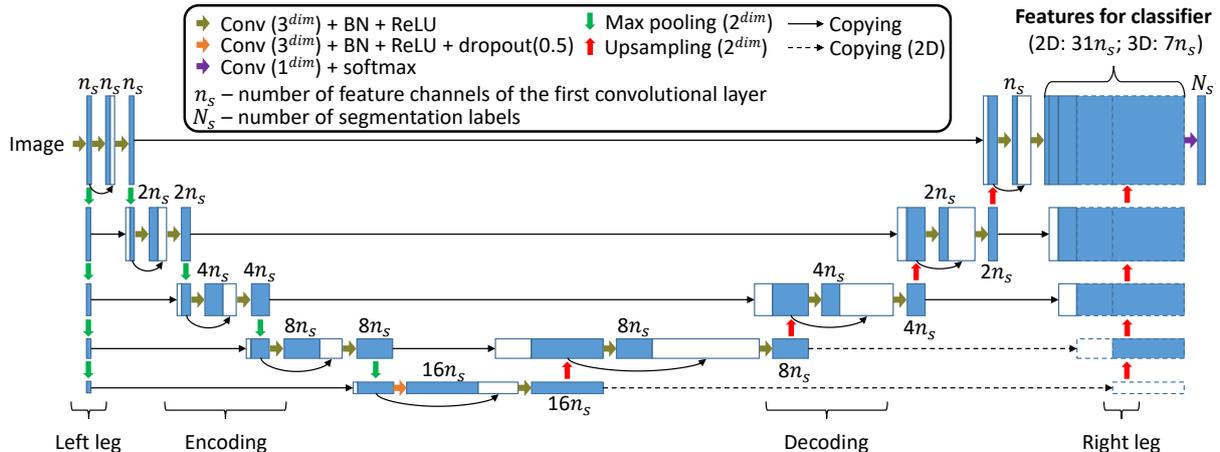


Figure 1: Segmentation network architecture. Blue blocks represent operation outputs and white blocks represent copied data. $dim \in \{2, 3\}$ is the input image dimension. At the right leg, dashed arrows and blocks are not used in 3D image analysis for computational feasibility. For the 2D cardiac CTA images, $n_s = 16$ and the number of features for the classifiers was $31n_s = 496$. For the 3D brain MR features, $n_s = 12$ and the number of features for the classifiers was $7n_s = 84$.

side paths (left and right legs) which provide functionality of deep supervision (Lee et al., 2015).

For our segmentation network, the M-Net is modified so that the number of feature channels of each convolutional layer systematically evolves with max pooling and upsampling. The number of channels is doubled after each pooling and halved after each upsampling. The encoding path consists of repetitions of two cascaded 3^{dim} convolutional layers and a 2^{dim} max pooling layer, with $dim \in \{2, 3\}$ the input image dimension. A dropout layer with 0.5 probability is introduced after the last pooling layer to reduce overfitting. The decoding path consists of repetitions of a 2^{dim} upsampling layer, a concatenation with the corresponding feature channels from the encoding path, and two cascaded 3^{dim} convolutional layers. Padding is used for each convolution to ensure segmentation image of the same size. Each 3^{dim} convolution associates with batch normalization (BN) (Ioffe and Szegedy, 2015) and rectified linear units (ReLU) (Nair and Hinton, 2010), and skip connections are introduced between cascaded convolutional layers for better features learning (Srivastava et al., 2015). Similar to Ronneberger et al. (2015), we have four max pooling and four upsampling layers. For the two side paths, the left leg operates on the feature channels of the first convolutional layer with four 2^{dim} max pooling layers, and the outputs are input to the corresponding encoding layers. The right leg consists of upsampling and concatenation that combine the

outputs of the decoding layers. The feature channels of the last 3^{dim} convolution and of the right leg are concatenated and input to the final 1^{dim} convolutional layer with the softmax function to provide the segmentation probabilities. For computational feasibility, only parts of the outputs from the decoding path are used for the right leg for 3D segmentation.

This segmentation network can be pre-trained to provide useful features for classification. Compared with the last concatenation layer in the modified U-Net in Wong et al. (2017), the right leg of this modified M-Net can provide more features (e.g. $3n_s$ in U-Net vs. $31n_s$ in M-Net for 2D images, Fig. 1). Therefore, this modified M-Net is used as the feature source in our classification framework.

2.2. Architectures of classification networks

2.2.1. Proposed architecture with a pre-trained segmentation network

The proposed classification network with a pre-trained segmentation network is shown Fig. 2(a). Instead of combining with an existing complex architecture such as VGGNet (Simonyan and Zisserman, 2014), we found that only a few convolutional layers with max pooling and fully-connected (FC) layers can provide high classification performance. Our classification network takes the features from a pre-trained segmentation network as inputs. The features are first downsampled by 2^{dim} max pooling, followed by r repetitions of a conv-pooling pair

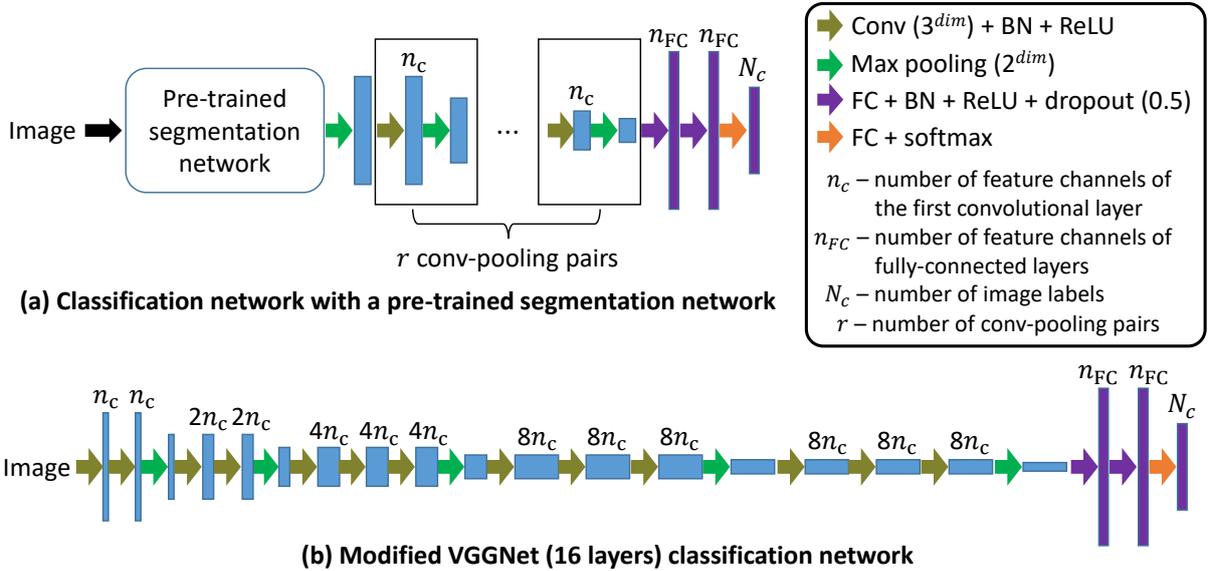


Figure 2: Classification network architectures. $dim \in \{2, 3\}$ is the input image dimension. (a) The proposed classification network architecture with a pre-trained segmentation network. For the 2D cardiac CTA images, $n_c = 16$, $n_{FC} = 100$, and $r = 3$. For the 3D brain MR images, $n_c = 32$, $n_{FC} = 200$, and $r = 4$. (b) A modified VGGNet architecture, which reproduces VGGNet when $n_c = 64$, $n_{FC} = 4096$, and with batch normalization removed. For the 2D cardiac CTA images, $n_c = 16$ and $n_{FC} = 100$. For the 3D brain MR images, $n_c = 16$ and $n_{FC} = 200$.

(a 3^{dim} convolutional layer followed by a 2^{dim} max pooling layer). Similar to the segmentation network, each 3^{dim} convolution associates with BN and ReLU. The output from the final max pooling layer is flattened and passed to three cascaded FC layers. The first two FC layers associate with BN, ReLU, and dropout with 0.5 probability. The last FC layer with the softmax function provides the classification probabilities.

2.2.2. Architecture of modified VGGNet

We compare the performance of our proposed classification network to a modified VGGNet (16 layers) architecture. We found in the early experiments that for tasks with relatively few classes compared to the 1000 classes in ImageNet, fewer feature channels can provide better performance with limited training data. Therefore, we modify the VGGNet architecture as shown in Fig. 2(b), which is mostly the same as the original VGGNet architecture but using BN with the convolutional and FC layers. We also generalize the evolution of the numbers of channels of the convolutional layers with respect to the max pooling layers.

For 2D image classification, we also compare our framework to a classification network with features

from an ImageNet pre-trained VGGNet model. To utilize the features, the FC layers of the pre-trained VGGNet are replaced by those of the modified VGGNet architecture with $n_{FC} = 100$. Only the weights of the FC layers are modified during training. This allows us to compare our framework with a commonly used transfer learning strategy.

2.3. Training strategy

Similar training strategies are used for both segmentation and classification networks. For the segmentation network, weighted cross entropy is used to compute the loss function as:

$$L = - \sum_{\mathbf{x}} w_{l(\mathbf{x})} \ln(p_{l(\mathbf{x})}) \quad (1)$$

where \mathbf{x} is the pixel position and $l(\mathbf{x})$ is the corresponding ground-truth label. $p_{l(\mathbf{x})}$ is the softmax probability of the channel corresponds to $l(\mathbf{x})$. $w_{l(\mathbf{x})}$ is the label weight for reducing the influences of more frequently seen labels. For the classification networks, the summation in (1) is removed.

The label weights can be computed from the training data as:

$$w_l = \left(\frac{\sum_k f_k}{f_l} \right)^{0.5} \quad (2)$$

with f_k the number of pixels with label k in segmentation, or the number of images with label k in classification. The exponent of 0.5 is used to avoid over penalty as the differences among label occurrences can be large. The stochastic optimization algorithm Adam is used for fast convergence (Kingma and Ba, 2014). Glorot uniform initialization is used for both the convolutional and FC layers (Glorot and Bengio, 2010).

With limited training data, image augmentation is important for learning invariant features and reducing overfitting. As computing realistic deformation of anatomical structures can be difficult and time consuming, while unrealistic deformation may introduce undesirable knowledge to the networks, we limit the image augmentation to rigid transformations, including rotation, shifting, zooming, and flipping (for symmetric structures such as the brain). In each epoch, each image has a probability (0.6 in our experiments) to be transformed by random transformations within given ranges, thus the number of augmented images is proportional to the number of epochs. Same setting is applied to all tested frameworks.

For the classification network with the ImageNet pre-trained VGGNet model, additional image pre-processing is required to comply with the input formats. Firstly, the images need to be resized to $224 \times 224 \times 3$ to match the pre-training image size. Secondly, in medical images, the image intensity may not be within 0 and 255 as the ImageNet data do. In such situations, the intensity is normalized and shifted by -128 to resemble the pre-training condition. Contrast limited adaptive histogram equalization (CLAHE) is used for the normalization which also improves image contrast (Zuiderveld, 1994). In fact, this normalization step may be unnecessary with the trainable FC layers, but we do this to ensure the network provides its best performance in the experiments.

2.4. Implementation

The Python deep learning library Keras was used with the TensorFlow backend (Chollet et al., 2017). The DIGITS DevBox of NVIDIA was used for the experiments, which has four TITAN X GPUs with 12 GB of memory per GPU. Only one GPU was used for each experiment.

2.5. Tested frameworks for image classification

With the black-box nature of CNNs, it is difficult to perform fundamental or theoretical analysis. Al-

though fundamental analysis such as that in Bengio et al. (2009) exists, it is very challenging to develop a similar analysis for our framework. Therefore, we demonstrate the advantages of our proposed framework through comparisons with other frameworks on different problems, organs, and image modalities. The tested frameworks include:

- **MANUAL**: the classification network in Fig. 2(a) (Section 2.2.1) with the segmentation network pre-trained on samples with manual segmentation by radiologists.
- **THRESHOLD**: the classification network in Fig. 2(a) (Section 2.2.1) with the segmentation network pre-trained on samples with segmentation obtained by intensity thresholding, i.e. k -means clustering.
- **VGG**: the classification network in Fig. 2(b) with the convolutional features from the ImageNet pre-trained VGGNet model (Section 2.2.2). As the 3D counterpart of the ImageNet pre-trained VGGNet model does not exist, this framework was only tested on 2D image classification.
- **SCRATCH**: the classification network in Fig. 2(b) (Section 2.2.2) trained from scratch without any pre-trained features.

The MANUAL framework is the optimal configuration of our proposed framework with pre-trained features from semantic segmentation. As manual segmentation can be expensive and time consuming, we use the THRESHOLD framework to study the use of non-semantic labels from image thresholding. The VGG framework represents a commonly used transfer learning technique (Yosinski et al., 2014; Shin et al., 2016). The SCRATCH framework acts as a baseline for comparison. These frameworks were tested on tumor classification and cardiac semantic level classification.

3. Tumor classification on brain MR images

Our goal is to distinguish among no tumor, low-grade gliomas (LGG), and glioblastomas. Gliomas are brain tumors which can be classified into Grade I (least malignant) through Grade IV (most malignant) using the World Health Organization classification system (Louis et al., 2007). LGG consist of Grade I and II gliomas while glioblastomas are

Table 1: Semantic brain segmentation. Semantic labels and their corresponding Dice coefficients between prediction and ground truth on the testing data. CVL represents cerebellar vermal lobules. The average Dice coefficient was 69%.

Cerebral grey	3rd ventricle	4th ventricle	Brainstem	CVL I-V
83%	59%	75%	80%	73%
CVL VI-VII	CVL VIII-X	Accumbens	Amygdala	Caudate
61%	71%	45%	63%	63%
Cerebellar grey	Cerebellar white	Cerebral white	Hippocampus	Inf. lateral vent.
82%	77%	86%	58%	46%
Lateral ventricle	Pallidum	Putamen	Thalamus	
84%	65%	71%	76%	

Grade IV (van den Bent et al., 2011; Stark et al., 2012). In T1-weighted MR images, LGG are often homogeneous with low image intensity, while glioblastomas are more heterogeneous and contrast enhanced (Fig. 3 (b) and (c)).

3.1. Segmentation network pre-training

3.1.1. Data and training

A dataset of 43 3D images from patients without brain tumors was neuroanatomically labeled to provide the training and testing data for the segmentation network. The images were produced by the T1-weighted, magnetization-prepared rapid gradient-echo (MP-RAGE) pulse sequence which provides good tissue contrast for anatomical segmentation. They were manually segmented by highly trained experts with the results reviewed by a consulting neuroanatomist. The 19 anatomical labels used are shown in Table 1.

From the dataset, 21 volumes were used for training and 22 were used for testing. All volumes were resized to $128 \times 128 \times 128$ pixels. Different from the Hounsfield unit in CT, the intensity values of MR images do not directly characterize material properties, and the variations can be large with different settings. Therefore, we normalized the image intensity between 0 and 255 using CLAHE for consistency. As mentioned in Section 2.3, image augmentation was used to learn invariant features and reduce overfitting. The segmentation network was trained with the batch size of 1, learning rate of 10^{-3} , and 100 epochs. We used $n_s = 12$ (Fig. 1) as this produced a relatively small model that can fit in a GPU with 12 GB of memory during training. $N_s = 20$ with the background included. These hyperparameters were chosen as they provided the best performance among the tested parameters.

3.1.2. Semantic segmentation results

As the segmentation performance is not the main concern of this paper, only the results of the trained segmentation network that was used by the classification network are presented. The Dice coefficient between the prediction and the ground truth of each semantic label was computed for each volume, and the average values are reported in Table 1.

Some structures with Dice coefficients $\leq 63\%$, such as the nucleus accumbens, amygdala, and inferior lateral ventricle, were small structures which were more prone to segmentation error. Structures which were relatively large or with simpler patterns, such as the cerebral grey and white matters, brainstem, and lateral ventricle, were well segmented with Dice coefficients $\geq 80\%$. The average Dice coefficient was 69%.

Fig. 3 shows examples of segmentation on unseen images. For the segmentation on the testing data (Fig. 3(a)), the results were generally consistent with the ground truths but with some missing details. On the other hand, as our goal is to apply the features from the trained segmentation network for tumor classification, we also show some segmentation results on images with LGG and glioblastoma in Fig. 3(b) and (c). As the semantic labels did not account for tumors, they were not segmented. Nevertheless, the segmented structures around the tumors were visually distinguishable from their non-tumorous counterparts, and the overall segmentations were less symmetric. Therefore, useful information to distinguish between non-tumorous and tumorous cases was embedded in the feature channels. Note that we have $7n_s = 84$ features for the classification (Fig. 1).

3.1.3. Non-semantic segmentation results

As mentioned in Section 2.5, we also tested a framework with a segmentation network trained

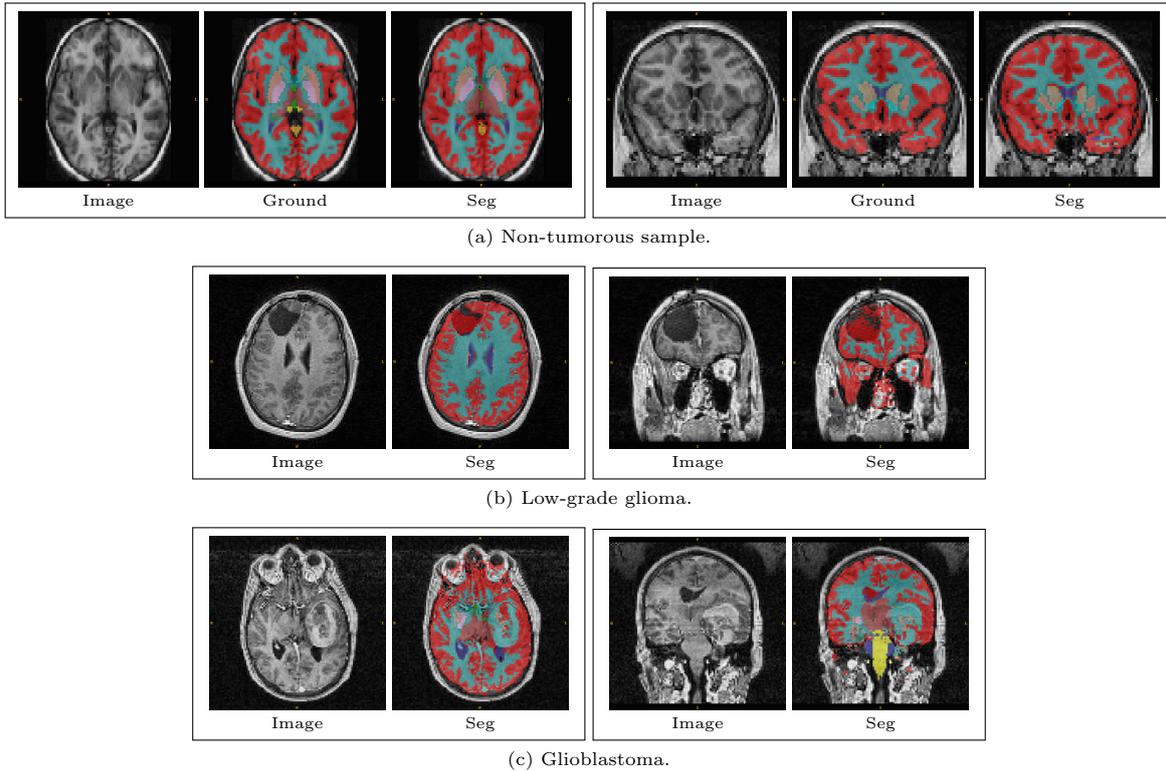


Figure 3: Semantic brain segmentation. (a) An unseen non-tumorous sample from the segmentation testing data in axial and coronal views. (b) and (c) Samples with low-grade glioma and glioblastoma (no ground truth).

from non-semantic labels (THRESHOLD). The same images and training strategy in Section 3.1.1 were used to train this segmentation network, with the five ground-truth labels provided by the k -means method. The number of classes was chosen so that segmented key structures such as grey and white matters were visually distinguishable. The segmentation results on unseen tumorous images are shown in Fig. 4. Although the segmentations were noisy, similar to the semantic segmentation, the segmented structures around the tumors were partially distinguishable.

3.2. Tumor classification

3.2.1. Data and training

The tumorous images were obtained from The Cancer Genome Atlas (TCGA) from patients with LGG and glioblastomas (TCGA, 2017). As the MR images for training the segmentation network were produced by the MP-RAGE pulse sequence, images from the TCGA dataset produced by the MP-RAGE or spoiled gradient recalled (SPGR) pulse sequences which produce similar tissue con-

trast were used. This provided 155 and 125 3D images for LGG and glioblastomas, respectively. For both LGG and glioblastomas, we used 50 samples for training and the rest for testing. This resulted in 50 and 105 images for training and testing for LGG, and 50 and 75 for training and testing for glioblastomas.

The testing data of the segmentation network provided the non-tumorous samples. By splitting the dataset in half, 11 images were produced for each training and testing.

Similar to the segmentation network, all images were resized to $128 \times 128 \times 128$ pixels, and the image intensity was normalized between 0 and 255 using CLAHE. Image augmentation was used to learn invariant features and reduce overfitting. The classification networks were trained with the batch size of 4, learning rate of 10^{-4} , and 50 epochs. For the MANUAL and THRESHOLD frameworks (Fig. 2(a)), $n_c = 32$, $n_{FC} = 200$, and $r = 4$. For the SCRATCH framework, $n_c = 16$ and $n_{FC} = 200$. $N_c = 3$ for the three classes of no tumor, LGG, and glioblastomas. These hyperparameters were chosen

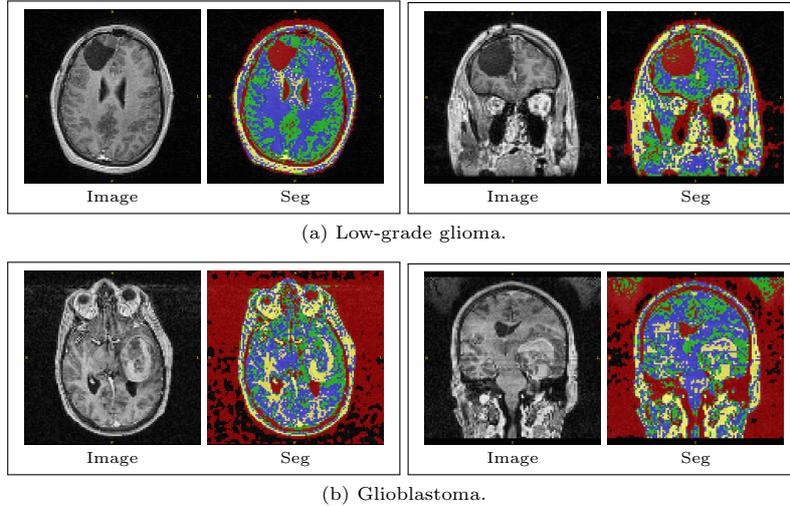


Figure 4: Non-semantic brain segmentation with low-grade glioma and glioblastoma in axial and coronal views (no ground truth).

so that all frameworks converged and achieved the best performances when all training data were used. To study the performance with limited samples, the tumorous training samples of each of the LGG and glioblastomas were gradually increased from five to 50 with strides of five (i.e. 10 to 100 with strides of 10 for all tumorous samples). All non-tumorous training samples (11 3D images) were used given the small number. Five repetitions were performed on randomized splits of training and testing data. Identical data and setting were applied to all frameworks.

3.2.2. Classification results

Fig. 5 shows the results of the tested frameworks. Their performances were assessed by the classification accuracy, Cohen’s kappa, and F1 score. The Cohen’s kappa ($\in[-1, 1]$) measures the inter-rater agreement between the ground-truth and predicted classes. The F1 score ($\in[0, 1]$) is the harmonic average of the recall and precision. For each framework, the accuracy and Cohen’s kappa were computed for the overall performance, while the F1 scores were computed for each class.

The accuracies and Cohen’s kappas show that the frameworks with a pre-trained segmentation network (MANUAL and THRESHOLD) outperformed the SCRATCH framework when the number of tumorous samples was small. All frameworks started to have similar performances at 70 tumorous samples. The standard deviations of the

SCRATCH framework were larger than those of the other frameworks in general. The MANUAL framework performed slightly better than the THRESHOLD framework, with the average absolute differences in accuracy and Cohen’s kappa as 3% and 5%, respectively. On the contrary, the average absolute differences between the MANUAL and the SCRATCH frameworks in accuracy and Cohen’s kappa were 10% and 12%, respectively. The MANUAL framework converged at 80 tumorous samples with the accuracy and Cohen’s kappa as 82% and 67%, respectively.

The F1 scores of each class provide further insights. The differences were larger in the no tumor and LGG classes, for which the MANUAL and THRESHOLD frameworks performed better than the SCRATCH framework in general. For the no tumor class, although the number of normal training samples did not change, the F1 scores improved with the number of tumorous samples, though the fluctuations and standard deviations were relatively large. The MANUAL and THRESHOLD frameworks performed well with LGG even with small numbers of training samples, for example, the MANUAL framework had F1 score of 80% with 40 tumorous samples. In contrast, the SCRATCH framework had relatively low scores before 70 tumorous samples with much larger fluctuations and standard deviations. For glioblastomas, the performances of all frameworks were similar.

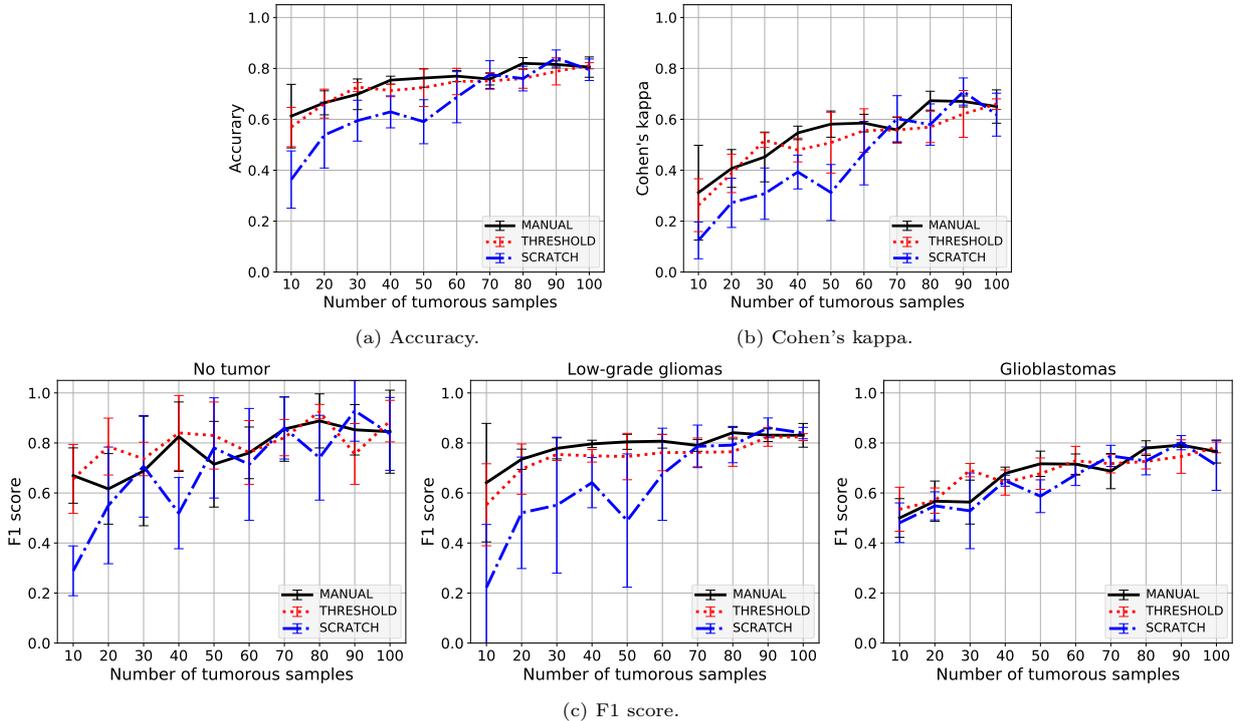


Figure 5: Brain tumor classification. Results of the tested frameworks (Section 2.5). Each data point corresponds to five repeated experiments, with the lines representing the average values and the error bars representing the standard deviations. The x -axis represents the number of tumorous training samples (five to 50 for each class), and the number of non-tumorous samples was fixed at 11. The accuracies and Cohen's kappas were computed from all testing samples, while the F1 scores were computed for each class.

4. Semantic axial level classification on cardiac CTA images

Apart from disease classification, we also tested our frameworks on semantic axial level classification on cardiac CTA images. With the advancements of medical imaging, 3D cardiac CTA images with hundreds of slices are common. Although these images provide valuable clinical information, they also pose challenges on computational time and accuracy for computer-aided diagnosis, especially on 3D analysis. As a result, it is beneficial to divide the entire volume into subregions. Following the work in Moradi et al. (2016), we tested the frameworks on classifying 2D cardiac CTA slices into nine semantic axial levels (Fig. 6). This is a challenging task as seven of the levels are over the relatively narrow span of the heart.

4.1. Segmentation network pre-training

4.1.1. Data and training

A dataset of 48 3D cardiac CTA images was manually labeled by a radiologist to provide the train-

ing and testing data for the segmentation network. There were 16 anatomical labels in each volume (Table 2).

The volumes were split in half for training and testing, resulting in 3910 and 5215 2D slices for training and testing, respectively. All slices were resized to 128×128 pixels, and we directly used the image intensity in Hounsfield unit. The segmentation network was trained with a batch size of 5, learning rate of 10^{-3} and 20 epochs. Image augmentation was used. We used $n_s = 16$ and $N_s = 17$ with the background included (Fig. 1). These hyperparameters were chosen as they provided the best performance among the tested parameters.

4.1.2. Semantic segmentation results

As segmentation performance is not the main concern of this paper, only the results of the trained segmentation network that was used by the classification network are presented. Table 2 shows the Dice coefficients between the prediction and ground truth on the 5215 testing slices. As not all seman-

Table 2: Semantic heart segmentation. Semantic labels and their corresponding Dice coefficients between prediction and ground truth on the testing data. PA stands for pulmonary artery. The average Dice coefficient was 78%.

Sternum	Ascending aorta	Descending aorta	Right PA
69%	78%	81%	81%
Vertebrae	Right atrium	Left atrium	Right ventricle
74%	85%	88%	81%
Left ventricle	Myocardium	Left PA	PA trunk
80%	87%	81%	82%
Aortic root	Aortic arch	Superior vena cava	Inferior vena cava
62%	76%	79%	62%

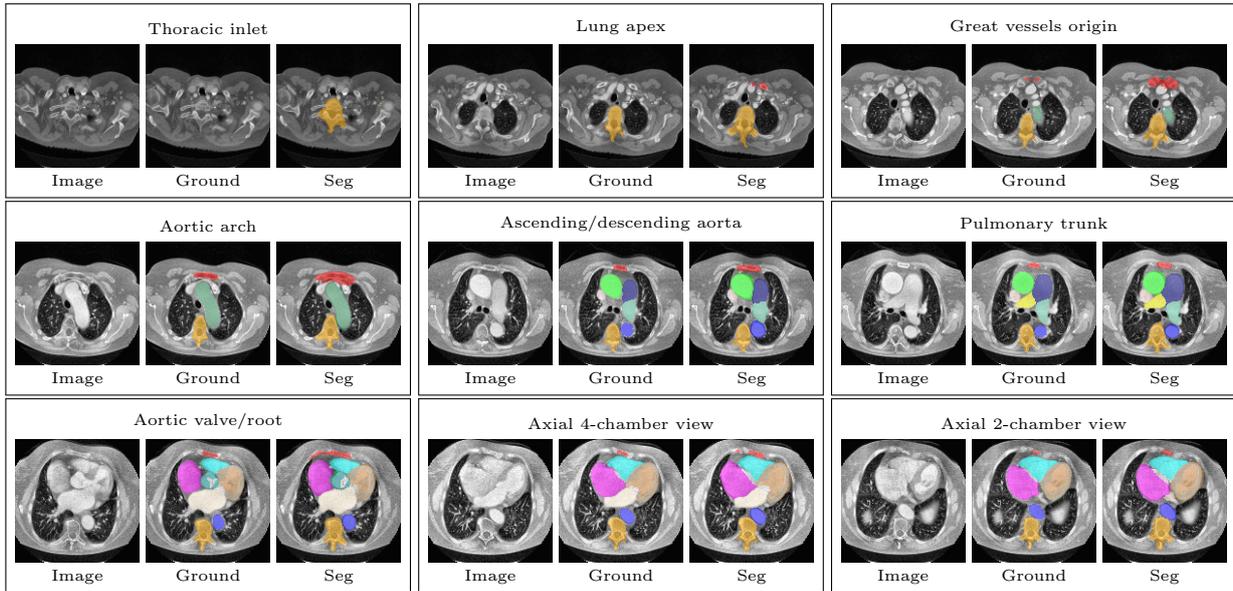


Figure 6: Semantic cardiac CTA image segmentation on unseen samples with the semantic axial level shown on top of each sub-figure. The segmentations were performed on image intensity in Hounsfield unit, while the images shown in this figure were contrast-enhanced by CLAHE for visualization.

tic labels were available on every slice, we used the pixel labels of all slices together when computing the Dice coefficients to avoid ambiguity.

There were 13 labels with Dice coefficients $\geq 74\%$ and the rests were $\geq 62\%$. The average Dice coefficient was 78%. The atria and the myocardium had the best Dice coefficients, followed by the pulmonary arteries, ventricles, and aorta. The sternum, aortic root, and inferior vena cava were challenging for the segmentation network. As the network could provide segmentations similar to those of the ground truth given the complicated structures across different levels of the cardiac CTA images, it contained useful features for the semantic axial level classification.

Fig. 6 shows the segmentation results on unseen samples. Although the segmentation network

could not provide segmentations as detailed as the ground truth, the segmented structures were consistent with those of the ground truth with good overlaps. For the thoracic inlet, the segmentation network correctly segmented the vertebrae even it did not exist in the ground truth. These show that the network learned the structural features.

4.1.3. Non-semantic segmentation results

A non-semantic segmentation network was also trained for the THRESHOLD framework (Section 2.5). The same images and training strategy in Section 4.1.1 were used to train this segmentation network, with the ground-truth labels provided by the k -means method with ten classes. The number of classes was chosen so that some segmented structures such as the blood pools and myocardium were

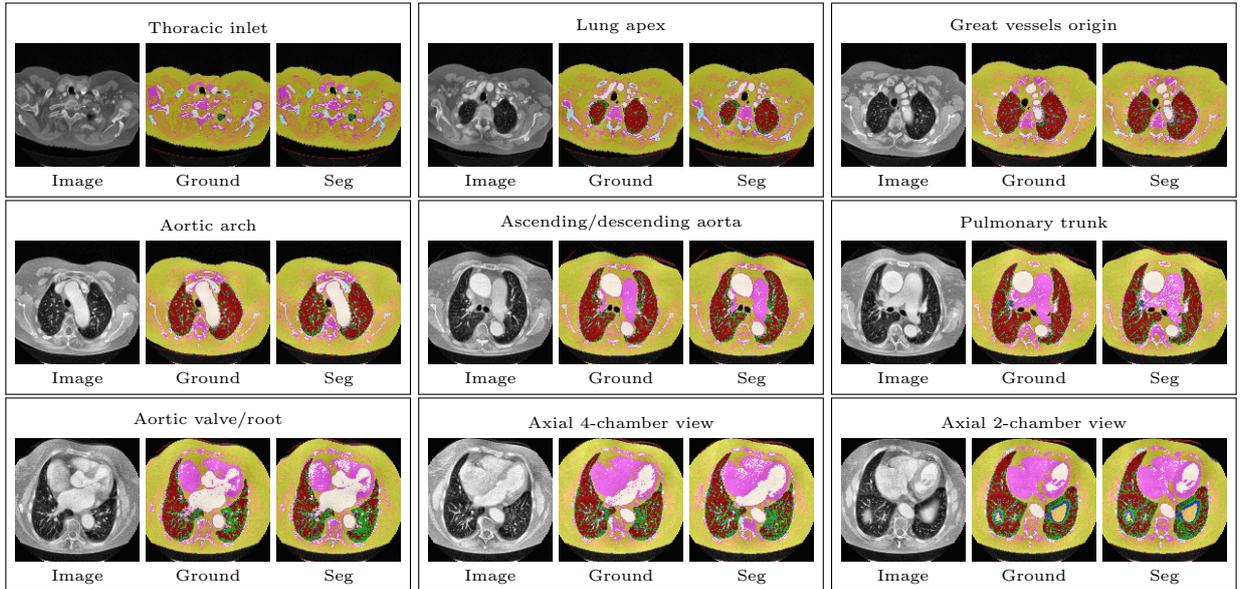


Figure 7: Non-semantic cardiac CTA image segmentation on unseen samples with the semantic axial level shown on top of each sub-figure. The segmentations were performed on image intensity in Hounsfield unit, while the images shown in this figure were contrast-enhanced by CLAHE for visualization.

visually distinguishable. Nevertheless, this was not an easy task as different structures share similar intensity values. The segmentation results on unseen samples are shown in Fig. 7. Only a few structures such as the blood pools were visible and most classes were semantically irrelevant.

4.2. Cardiac semantic axial level classification

4.2.1. Data and training

There were 425 2D slices annotated by a radiologist into nine classes of semantic axial levels for training and testing. Using 18 training samples for each class, we had 162 and 263 images for training and testing, respectively.

Similar to the segmentation network, all images were resized to 128×128 pixels, and the image intensity in Hounsfield unit was used except for the VGG framework (Section 2.3). Image augmentation was used to learn invariant features and reduce overfitting. The classification networks were trained with the batch size of 64, learning rate of 5×10^{-4} , and 300 epochs. We used $r = 3$ for the MANUAL and THRESHOLD frameworks (Fig. 2(a)), and used $n_c = 16$ and $n_{FC} = 100$ for all frameworks (Fig. 2). $N_c = 9$ for the nine classes. These hyperparameters were chosen as they provided the best performance among the tested parameters. To study the performance with limited

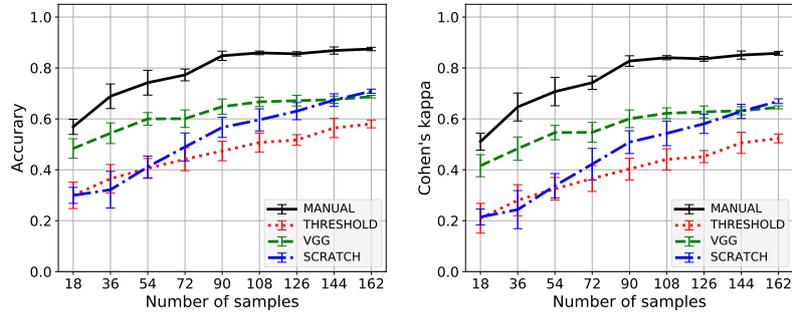
numbers of samples, we gradually increased the training samples of each class from two to 18 with strides of two (i.e. from 18 to 162 with strides of 18 for all classes). Five repetitions were performed on randomized splits of training and testing data. Identical data and setting were applied to all frameworks.

4.2.2. Classification results

Fig. 8 shows the results of the tested frameworks. The performances were assessed by the classification accuracy, Cohen’s kappa, and F1 score. For each framework, the accuracies and Cohen’s kappas were computed for the overall performance, while the F1 scores were computed for each class.

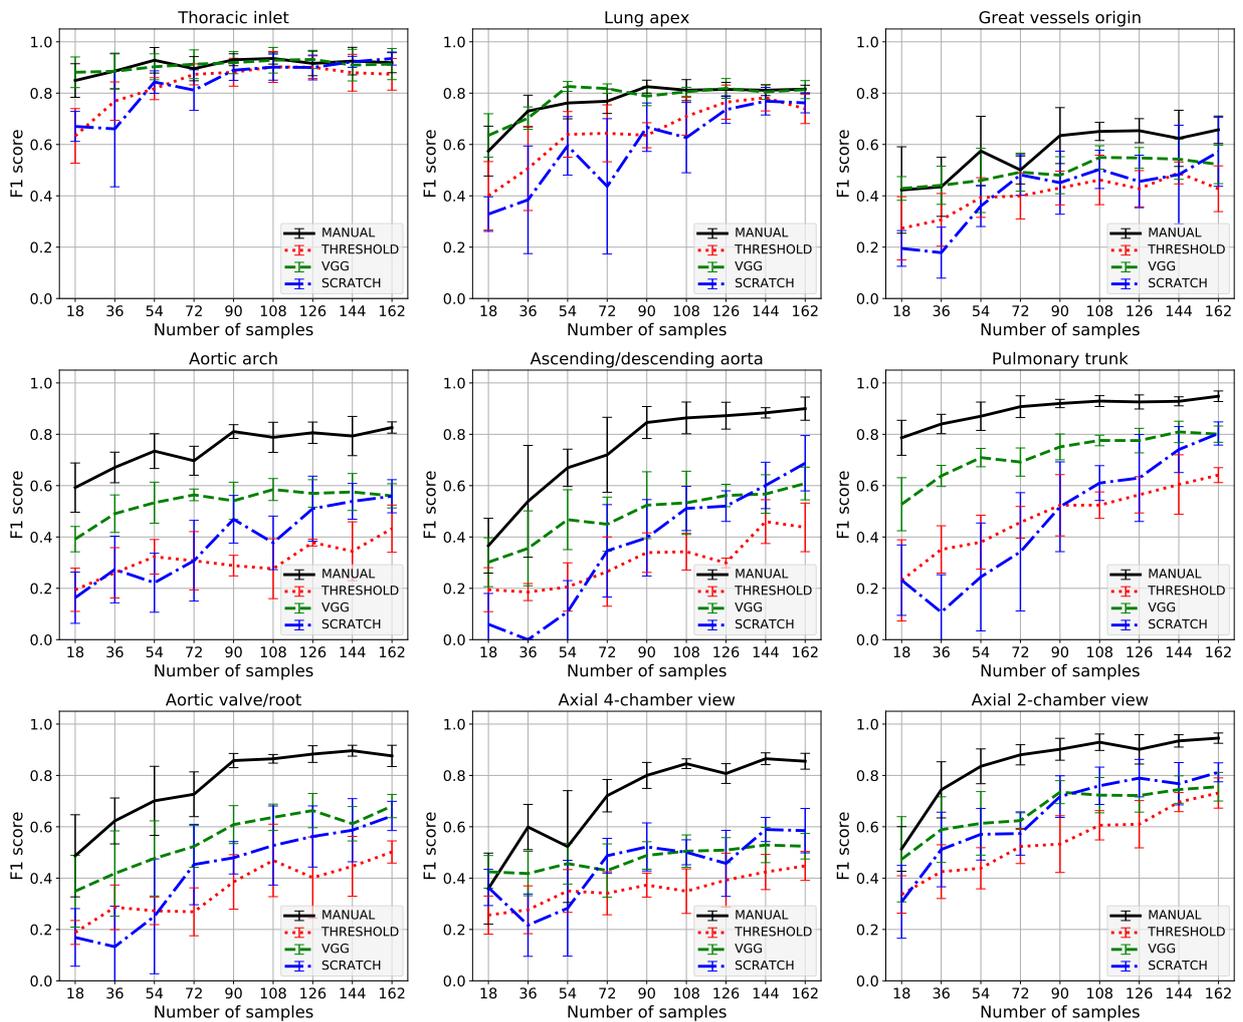
Both accuracies and Cohen’s kappas show that the MANUAL framework had the best performance. The VGG framework outperformed the THRESHOLD and SCRATCH frameworks with small training samples, but was caught up by the SCRATCH framework at 144 samples. Both MANUAL and VGG frameworks converged at about 108 samples, though the MANUAL framework had better performance (86% accuracy, 84% Cohen’s kappa) than the VGG framework (67% accuracy, 62% Cohen’s kappa).

The F1 scores at different semantic axial levels provide further comparisons. The MANUAL



(a) Accuracy.

(b) Cohen's kappa.



(c) F1 score.

Figure 8: Cardiac semantic axial level classification. Results of the tested frameworks (Section 2.5). Each data point corresponds to five repeated experiments, with the lines representing the average values and the error bars representing the standard deviations. The x -axis represents the number of training samples (two to 18 for each class). The accuracies and Cohen's kappas were computed from all testing samples, while the F1 scores were computed for each class.

framework outperformed the other frameworks in general, but the degrees were different among classes. For example, all frameworks performed best at the thoracic inlet level and achieved almost 90% of F1 scores. A similar trend can be observed at the lung apex level. On the other hand, the differences between the MANUAL and the other frameworks were large at the ascending/decending aorta level and the 4-chamber view level. At most levels, the SCRATCH framework started with relatively low F1 scores, but gradually outperformed the THRESHOLD framework and caught up with the VGG framework with increased training samples. At 90 samples, the MANUAL framework had high average F1 scores between 80% and 93% except at the great vessels origin level (63%). The MANUAL framework was more consistent than the other frameworks given its small standard deviations at all levels after convergence at 90 samples.

5. Discussion

The experimental results in Section 3 and 4 provide several interesting observations for discussion.

As segmentation accuracy is not our main concern, we simply modified an existing network architecture (M-Net) as a feature source with respect to the computational feasibility. Although we examined different hyperparameters and selected those provided the best results, our segmentation accuracy could not be as good as those tailored for specific problems (Mehta and Sivaswamy, 2017; Roy et al., 2017). Nevertheless, given the complicated anatomical structures with more than 16 labels, the average Dice coefficients of 69% on 3D brain MR images and 78% on 2D cardiac CTA images show that the segmentation networks learned useful features that improved classification performance with limited data.

From our experience on the presented segmentation and classification problems on medical images, we may not need large networks and big data to provide useful clinical applications. Using the 2D cardiac semantic axial level classification as an example, compared with the ~ 138 million parameters of VGGNet, there were only ~ 3 million, ~ 17 million, and ~ 1 million parameters for the MANUAL, VGG, and SCRATCH frameworks, respectively. This is mainly due to the fact that we usually do not have problems with 1000 classes in medical image analysis. Although we may not require a lot of training data, a reasonably large amount of

testing data should be used to ensure the practical applicability of the trained models.

The MANUAL and THRESHOLD frameworks can be very useful and important especially for 3D image classification, as 3D models pre-trained from hundreds of thousands of 3D images are not publicly available to the best of our knowledge. Furthermore, notice that for the brain segmentation networks, although only trained on 21 volumes of normal cases, the learned features improved the performance on tumor type classification. This shows that instead of using hundreds of thousands of image-based annotations to provide pre-trained feature sources, we can use fewer images by using more complex pixel-based annotations. This also shows that segmentation networks trained on normal images can improve disease classification.

The results of the brain and heart classification in Fig. 5 and 8 show some fluctuations of the curves when the numbers of samples were small, especially for the SCRATCH framework. One of the reasons was overfitting as there were not enough data for the networks to learn probably. As the SCRATCH framework does not have pre-trained features, it had larger fluctuations and standard deviations in general. On the other hand, frameworks with pre-trained features had relatively smooth curves. This shows that using pre-trained features can reduce overfitting.

The experiments on the THRESHOLD framework were used to investigate alternatives to the time consuming and expensive manual segmentation. This framework had different performances on different datasets, and this is probably related to the natures of the problems. For the brain tumor classification, detailed brain structures may be unnecessary as long as the tumor can be distinguished, and Fig. 4 shows that this purpose was achieved by image thresholding. On the other hand, the semantic level classification required more structural details, and image thresholding could not provide them. As potential future work, instead of using image thresholding, similar to Roy et al. (2017), we can use labels generated by an existing automated software to provide good enough training data to avoid manual segmentation.

The VGG framework tested in this paper may not be optimal for the presented classification problem. In Yosinski et al. (2014), it was shown that features of deeper layers may be too specific to be used by another problem. As it is nontrivial to know which layers should be used and the main goal of this pa-

per is not to explore the best way of using pre-trained features from ImageNet, we used the setting that is commonly used with limited data.

Given the large number of segmentation and classification deep learning architectures available, we could only choose the well-known architectures such as the VGGNet in this study. As different hyperparameters may lead to different results, for each framework on each dataset, we performed experiments to search for the appropriate parameters. Using the 2D cardiac semantic axial level classification as an example, the original VGGNet architecture with $n_c = 64$ and $n_{FC} = 4096$ (Fig. 2(b)) was trained from scratch, and it gave very low accuracy as expected as there were not enough data to train the large number of parameters. After trying different parameters, we found that $n_c = 16$ and $n_{FC} = 100$ gave the best results and thus we report this setting in the paper. A similar procedure was performed on the 3D brain tumor type classification to get $n_c = 16$ and $n_{FC} = 200$. As there are still no computationally feasible algorithms for hyperparameter optimization to the best of our knowledge, there is no guarantee that the hyperparameters are optimal, and this is a common problem for deep learning related studies. On the other hand, we tried our best to ensure meaningful comparisons. For example, the MANUAL and THRESHOLD frameworks were identical in architecture, hyperparameters, and training, with the only difference as the ground-truth annotations for training the segmentation networks.

Although the experimental results were promising, the proposed end-to-end framework may have reduced performance when applying to image classification tasks with subtle morphological differences. In such situations, algorithms that use localized CNN features such as that in Li et al. (2017) can be good alternatives.

6. Conclusion

In this paper, we propose a strategy for building medical image classifiers from pre-trained segmentation networks. Using a segmentation network pre-trained on data similar to those of the classification task, we can build classifiers that achieve high performance with very limited data. The proposed framework improved accuracy on 3D brain tumor classification from MR images, and outperformed the framework with ImageNet pre-trained VGGNet on 2D cardiac semantic level classification

from CTA images. As this framework is applicable to 3D medical image analysis in which models equivalent to ImageNet pre-trained CNNs are unavailable, this framework provides a useful and practical alternative.

References

- Bengio, Y., Louradour, J., Collobert, R., Weston, J., 2009. Curriculum learning, in: International Conference on Machine Learning, pp. 41–48.
- van den Bent, M.J., Wefel, J.S., Schiff, D., Taphoorn, M.J., Jaeckle, K., Junck, L., Armstrong, T., Choucair, A., Waldman, A.D., Gorlia, T., Chamberlain, M., Baumert, B., Vogelbaum, M., Macdonald, D., Reardon, D., Wen, P., Chang, S., Jacobs, A., 2011. Response assessment in neuro-oncology (a report of the RANO group): assessment of outcome in trials of diffuse low-grade gliomas. *The Lancet Oncology* 12, 583–593.
- Chollet, F., et al., 2017. Keras: The Python deep learning library. <https://keras.io/>. Accessed on October 2, 2017.
- Deng, J., Dong, W., Socher, R., Li, L.J., Li, K., Fei-Fei, L., 2009. ImageNet: A large-scale hierarchical image database, in: IEEE Conference on Computer Vision and Pattern Recognition, pp. 248–255.
- Dou, Q., Yu, L., Chen, H., Jin, Y., Yang, X., Qin, J., Heng, P.A., 2017. 3D deeply supervised network for automated segmentation of volumetric medical images. *Medical Image Analysis* 41, 40–54.
- van Ginneken, B., Setio, A.A., Jacobs, C., Ciompi, F., 2015. Off-the-shelf convolutional neural network features for pulmonary nodule detection in computed tomography scans, in: IEEE International Symposium on Biomedical Imaging, pp. 286–289.
- Glorot, X., Bengio, Y., 2010. Understanding the difficulty of training deep feedforward neural networks, in: International Conference on Artificial Intelligence and Statistics, pp. 249–256.
- Goodfellow, I., Pouget-Abadie, J., Mirza, M., Xu, B., Warde-Farley, D., Ozair, S., Courville, A., Bengio, Y., 2014. Generative adversarial nets, in: Advances in Neural Information Processing Systems, pp. 2672–2680.
- Ioffe, S., Szegedy, C., 2015. Batch normalization: Accelerating deep network training by reducing internal covariate shift, in: International Conference on Machine Learning, pp. 448–456.
- Kingma, D.P., Ba, J.L., 2014. Adam: A method for stochastic optimization. arXiv:1412.6980 [cs.LG].
- Krizhevsky, A., Sutskever, I., Hinton, G.E., 2012. Imagenet classification with deep convolutional neural networks, in: Advances in Neural Information Processing Systems, pp. 1097–1105.
- Lee, C.Y., Xie, S., Gallagher, P.W., Zhang, Z., Tu, Z., 2015. Deeply-supervised nets, in: International Conference on Artificial Intelligence and Statistics, pp. 562–570.
- Li, Z., Wang, Y., Yu, J., Guo, Y., Cao, W., 2017. Deep Learning based Radiomics (DLR) and its usage in non-invasive IDH1 prediction for low grade glioma. *Scientific Reports* 7, 1–11.
- Litjens, G., Kooi, T., Bejnordi, B.E., Setio, A.A.A., Ciompi, F., Ghafoorian, M., van der Laak, J.A., van Ginneken, B., Sánchez, C.I., 2017. A survey on deep learning in medical image analysis. *Medical Image Analysis* 42, 60–88.

- Louis, D.N., Ohgaki, H., Wiestler, O.D., Cavenee, W.K., Burger, P.C., Jouvet, A., Scheithauer, B.W., Kleihues, P., 2007. The 2007 WHO classification of tumours of the central nervous system. *Acta Neuropathologica* 114, 97–109.
- Madani, A., Moradi, M., Karargyris, A., Syeda-Mahmood, T., 2018a. Chest X-ray generation and data augmentation for cardiovascular abnormality classification, in: *Medical Imaging 2018: Image Processing*, p. 105741M.
- Madani, A., Moradi, M., Karargyris, A., Syeda-Mahmood, T., 2018b. Semi-supervised learning with generative adversarial networks for chest X-ray classification with ability of data domain adaptation, in: *IEEE International Symposium on Biomedical Imaging*, pp. 1038–1042.
- Mehta, R., Sivaswamy, J., 2017. M-net: A convolutional neural network for deep brain structure segmentation, in: *IEEE International Symposium on Biomedical Imaging*, pp. 437–440.
- Milletari, F., Navab, N., Ahmadi, S.A., 2016. V-Net: Fully convolutional neural networks for volumetric medical image segmentation, in: *IEEE International Conference on 3D Vision*, pp. 565–571.
- Mirza, M., Osindero, S., 2014. Conditional generative adversarial nets. *arXiv:1411.1784 [cs.LG]*.
- Moradi, M., Gur, Y., Wang, H., Prasanna, P., Syeda-Mahmood, T., 2016. A hybrid learning approach for semantic labeling of cardiac CT slices and recognition of body position, in: *IEEE International Symposium on Biomedical Imaging*, pp. 1418–1421.
- Nair, V., Hinton, G.E., 2010. Rectified linear units improve restricted Boltzmann machines, in: *International Conference on Machine Learning*, pp. 807–814.
- Pathak, D., Krahenbuhl, P., Donahue, J., Darrell, T., Efros, A.A., 2016. Context encoders: Feature learning by inpainting, in: *IEEE Conference on Computer Vision and Pattern Recognition*, pp. 2536–2544.
- Ronneberger, O., Fischer, P., Brox, T., 2015. U-Net: Convolutional networks for biomedical image segmentation, in: *International Conference on Medical Image Computing and Computer-Assisted Intervention*, pp. 234–241.
- Roy, A.G., Conjeti, S., Sheet, D., Katouzian, A., Navab, N., Wachinger, C., 2017. Error corrective boosting for learning fully convolutional networks with limited data, in: *International Conference on Medical Image Computing and Computer-Assisted Intervention*, pp. 231–239.
- Schmidhuber, J., 2015. Deep learning in neural networks: An overview. *Neural Networks* 61, 85–117.
- Shin, H.C., Roth, H.R., Gao, M., Lu, L., Xu, Z., Nogues, I., Yao, J., Mollura, D., Summers, R.M., 2016. Deep convolutional neural networks for computer-aided detection: CNN architectures, dataset characteristics and transfer learning. *IEEE Transactions on Medical Imaging* 35, 1285–1298.
- Simonyan, K., Zisserman, A., 2014. Very deep convolutional networks for large-scale image recognition. *arXiv:1409.1556 [cs.CV]*.
- Srivastava, R.K., Greff, K., Schmidhuber, J., 2015. Training very deep networks, in: *Advances in Neural Information Processing Systems*, pp. 2377–2385.
- Stark, A.M., van de Bergh, J., Hedderich, J., Mehdorn, H.M., Nabavi, A., 2012. Glioblastoma: Clinical characteristics, prognostic factors and survival in 492 patients. *Clinical Neurology and Neurosurgery* 114, 840–845.
- Szegedy, C., Liu, W., Jia, Y., Sermanet, P., Reed, S., Anguelov, D., Erhan, D., Vanhoucke, V., Rabinovich, A., 2015. Going deeper with convolutions, in: *IEEE Conference on Computer Vision and Pattern Recognition*, pp. 1–9.
- Tajbakhsh, N., Shin, J.Y., Gurudu, S.R., Hurst, R.T., Kendall, C.B., Gotway, M.B., Liang, J., 2016. Convolutional neural networks for medical image analysis: Full training or fine tuning? *IEEE Transactions on Medical Imaging* 35, 1299–1312.
- TCGA, 2017. TCGA Research Network. <https://cancergenome.nih.gov/>. Accessed on October 2, 2017.
- Uzunova, H., Wilms, M., Handels, H., Ehrhardt, J., 2017. Training CNNs for image registration from few samples with model-based data augmentation, in: *International Conference on Medical Image Computing and Computer-Assisted Intervention*, pp. 223–231.
- Wong, K.C.L., Karargyris, A., Syeda-Mahmood, T., Moradi, M., 2017. Building disease detection algorithms with very small numbers of positive samples, in: *International Conference on Medical Image Computing and Computer-Assisted Intervention*, pp. 471–479.
- Yosinski, J., Clune, J., Bengio, Y., Lipson, H., 2014. How transferable are features in deep neural networks?, in: *Advances in Neural Information Processing Systems*, pp. 3320–3328.
- Zhou, B., Khosla, A., Lapedriza, A., Oliva, A., Torralba, A., 2016. Learning deep features for discriminative localization, in: *IEEE Conference on Computer Vision and Pattern Recognition*, pp. 2921–2929.
- Zuiderveld, K., 1994. Contrast limited adaptive histogram equalization, in: *Graphics Gems IV*, Academic Press, Inc.. pp. 474–485.