

CAMELYON 2017: EXPANDING CONVOLUTIONAL FILTERS FOR ROBUST METASTASIS DETECTION IN LYMPH NODES

Arjun Moorthy, Arun Moorthy, Sujay Nair, Suraj Nair

ABSTRACT

This project is part of a submission in the 2017 CAMELYON pathology competition organized by the Diagnostic Image Analysis Group (DIAG) and Department of Pathology of the Radboud University Medical Center (Radboudumc) in Nijmegen, The Netherlands. The competition explores new methods for breast cancer metastasis detection using image analysis of lymph node sections. This entry in the competition presents an algorithm for such detection, utilizing advanced preprocessing, a convolutional neural network with expanding filters, and a rule based diagnosis system that makes use of the network output.

Index Terms— Metastasis Detection, Bio-Imaging, Pathology, Image Analysis

1. INTRODUCTION

This paper will begin by explaining the problem statement, and the details of the CAMELYON 2017 challenge. Then, it will present the algorithm, including all steps of preprocessing, modeling, and output generation. Lastly, it will examine the results of the algorithm.

2. BACKGROUND

This project is part of a submission in the 2017 CAMELYON competition. The purpose of the challenge was to "evaluate new and existing algorithms for automated detection and classification of breast cancer metastases in whole-slide images of histological lymph node sections" [1]. Specifically, the data was provided in the form of annotated whole slide images.

2.1. Whole Slide Images

The whole slide images of lymph node sections were provided for training and testing. The training data contained expert labeled annotations, exposing the lymph node, and in some cases metastasis sections. 5 different slides were provided for each of 100 anonymized patients. The slides varied in size and shape, and were available with multiple levels of down sampling.

Thanks to DIAG and Radboudumc for hosting the competition and cve-dia for providing data retrieval/preprocessing tools.

The whole slide images contained annotations which distinguished certain sections of the image. Specifically, each training image contained annotations which consisted of an expert drawn polygon which highlighted the areas of the image with tissue, lymph nodes, and metastases for the examples which contained metastases. The annotations could be used to isolate certain parts of the image for training. However, these annotations did not exist for the testing data, and separate preprocessing had to be applied.

2.2. Labels

2.2.1. Node Labels

Each node in the training data consisted of a label from an expert. Specifically, the images are classified as *negative*, *macro*, *micro*, or *itc*. An image is labeled *macro* if it contains macro-metastases, or metastases greater than 2.0 mm. Similarly, an image is labeled *micro* if it contains micro-metastases, or metastases between .2 and 2 mm, or having at least 200 cells, and contains no macro-metastases.

An image is labeled *itc* if it contains isolated tumour cells (ITC). They are not exactly metastases, and instead are "single tumour cells or a cluster of tumour cells smaller than 0.2 mm or less than 200 cells. Lymph nodes containing only ITC are therefore not counted as positive lymph nodes" [1]. If no metastases are found, pathologists are required to report ITC.

Lastly, an image is labeled *negative* if no metastases or ITC are found.

2.2.2. Patient Labels

Given the labels for each node, the following set of rules is used to determine the final label for the patient. Specifically, a simplified version of the pN-staging system is used [1]:

pN0: No micro-metastases or macro-metastases or ITCs found.

pN0(i+): Only ITCs found.

pN1mi: Micro-metastases found, but no macro-metastases found.

pN1: Metastases found in 1-3 lymph nodes, of which at least one is a macro-metastasis.

pN2: Metastases found in 4-9 lymph nodes, of which at least one is a macro-metastasis.

The final submission consists of a patient label for all 100 test patients and 5 node labels for each patient.

3. DATA ACQUISITION

Both the training and testing data were acquired using the CVEDIA tool. This tool allows for image processing, filtering, and augmentation prior to download.

3.1. Training Data

For the training data, each node image was selected and segmented depending on the label.

Specifically, if the image was labeled as negative, then the entire lymph node was segmented out. The reason for this is that is because the entire lymph node contains all the "healthy sections" which we want to focus on in labeling a lymph node as healthy. This segmentation was done using the expert annotations that were provided with the training data. Then, for each image, 50 256 x 256 pixel crops were generated, each with the "negative" label. By doing this, we create a large dataset with examples of healthy lymph nodes and the "negative" label.

If the image was labeled as "macro", "micro", or "itc", we select only the metastases section of the node image. Again, this is done using the expert annotations that were provided with the training data. Then, for each image, 50 256 x 256 pixel crops were generated, each with the appropriate "macro", "micro", or "itc" label. By doing this, we create a large dataset with examples of metastases or ITC.

Note, that due to the random cropping, there may be some random crops which contain some healthy tissue, but are part of a larger image which have metastases and have the corresponding label. Thus, these samples will be misleading, but these examples should be rare in the larger dataset.

3.2. Testing Data

The same CVEDIA tool was used for the processing of the testing data. Unlike the training data, the testing data has no annotations. Thus, for each node image, 50 256 x 256 random crops are generated. Of course, this leads to many crops containing nothing but background. These crops are filtered out using the mean intensity of the image because background will be exclusively black or white.

Thus, for each testing patient and node combination we have several random crops which can be used for prediction.

4. ALGORITHM

The overall algorithm consists of multiple stages, including data acquisition, preprocessing, modeling, and labeling.

As can be seen from figure 1, CVEDIA is used to acquire basic training data. Then Tensorflow, specifically Tflern, is

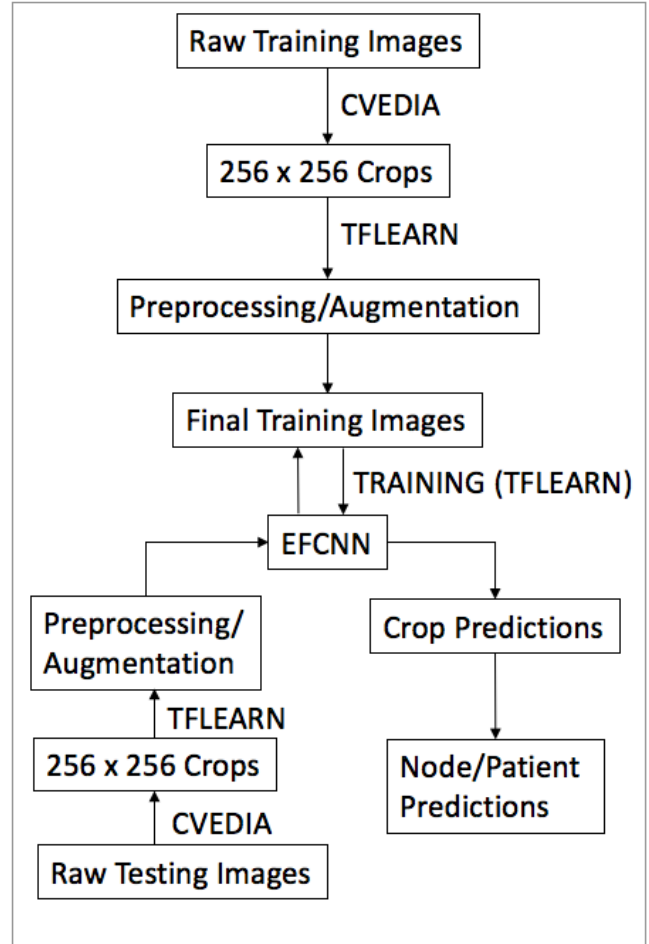


Fig. 1. Outline of full algorithm

used for further image augmentation. Then, these final training crops are used to train our EFCNN (Expanding Filter Convolutional Neural Network) Model

Similar crops are then regenerated for the testing data using CVEDIA. Tflern is used to apply final augmentations to these crops, which are fed into the EFCNN model. Then, based on the set of predictions for each node, we assign a node label. These node labels are then used to generate the final patient labels.

5. PREPROCESSING AND AUGMENTATION

Before random crops were used for training or testing, some basic preprocessing and augmentation was applied.

5.1. Preprocessing

The first step in image preprocessing was applying feature wise zero center. Specifically, this zero centers every sample with the mean over the full dataset. Next, the standard devia-

tion is evaluated over the full dataset and this is used to scale each sample.

The goal of this is make it easier for the model to differentiate lighter parts of the image from darker parts, and is a standard technique in image processing.

5.2. Augmentation

The goal of image augmentation was to add noise to the images to help generalize the model. Specifically, three different kinds of noise were added to the image.

5.2.1. Random Flip

Every image was flipped from left to right with some probability. This will help the model learn to recognize metastases regardless of angle.

5.2.2. Random Rotation

Every image is rotated by an angle uniformly sampled from -20 degrees to 20 degrees. Like the the random flip, this will help the model learn to recognize metastases regardless of angle.

5.2.3. Random Blur

Every image was blurred using a Gaussian filter with a randomly generated σ at most 5. This should help identify more and less pronounced metastases.

6. EXPANDING FILTER CONVOLUTIONAL NEURAL NETWORK

The model that was used to predict a label of macro, micro, itc, or negative given a single 256 x 256 crop was a convolutional neural network. The main innovation that is made to the standard convolutional network here is the use of expanding filters. After the convolution and pooling layers, there is a single hidden layer, a dropout layer for regularization, and the output layer with a softmax.

6.1. Convolutional Layers

The convolutional layers consist of 3 convolutional layers, each using relu activation, and each followed by max pooling. Specifically, we have the input image being fed into a convolutional layer with 32 filters, with filter size of 3. This is followed by relu activation and 2 x 2 max pooling. This is then fed into a convolutional layer with 16 filters and filter size of 7, again followed by relu activation and 2 x 2 max pooling. This is fed into a convolutional layer with 8 filters, and filter size 18, again followed by relu activation and 2 x 2 max pooling.

The idea of these expanding filters is that they will be able to identify distinguish smaller differences, and as the filter size grows can make larger conclusions about the existence of metastases. For example, in this architecture the idea is that the first two convolutional layers can learn to distinguish healthy cells from tumorous cells, and the layer with larger filter size can learn to draw conclusions from larger collections of healthy or infected cells.

6.2. Fully Connected Layers

The network consists a single fully connected hidden layer that comes after the convolutional layers. This layer contains 256 nodes, and uses relu activation. This is followed by a dropout layer, which sets node weights to 0 with probability 0.2.

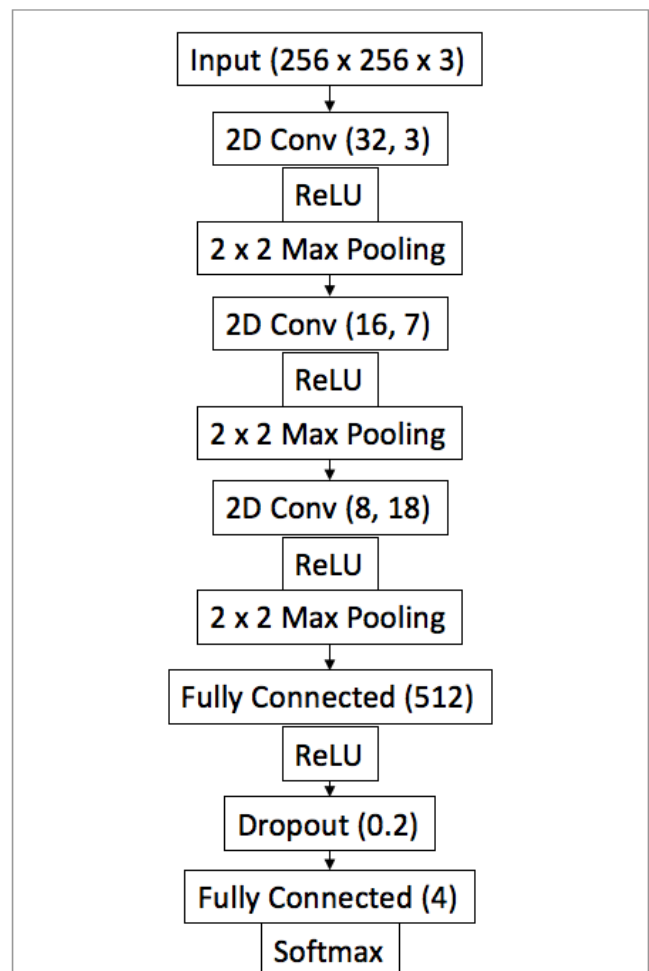


Fig. 2. Model Architecture

7. OUTPUT GENERATION

Once predictions were made on the 50 random crops per image, a majority vote across the crops was used to find the overall diagnosis for that image. Then, using the process outlined in the Background section, predictions for each node were used to create a diagnosis for the patient.

8. RESULTS

After only 100 steps, the model achieved close to 80% accuracy on the training data.




Fig. 3. Model Accuracy

Over the same period of time, the loss, categorical cross entropy, fell below 0.6.

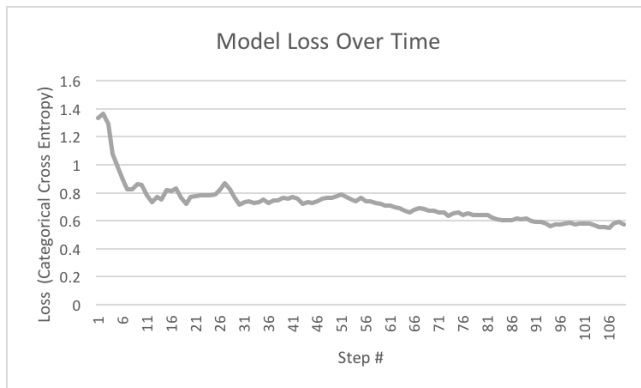


Fig. 4. Categorical Cross Entropy

9. REFERENCES

- [1] Consortium for Open Medical Image Computing, " in *CAMELYON17 Webpage*, 2017, pp. Background, Evaluation.