Convolutional Neural Networks in Medical Imaging

Mitchell Finzel

Division of Science and Mathematics
University of Minnesota, Morris
Morris, Minnesota, USA

15 April 2017,
Senior Seminar Conference

Introduction

- Convolutional neural networks or CNNs, have seen a rise in popularity in image related fields.
- CNNs have been having great success in biological segmentation tasks.
- These tasks include:
  - The automated detection of lymph nodes
  - Segmentation of knee cartilage
  - Detection of Alzheimer's

Introduction

- We will be looking at two approaches to brain MRI segmentation
- The goal of this work is to provide unsegmented MRIs to the network and receive properly segmented MRIs as output
- Currently this requires time consuming labor from a skilled medical professional

Input Output Example

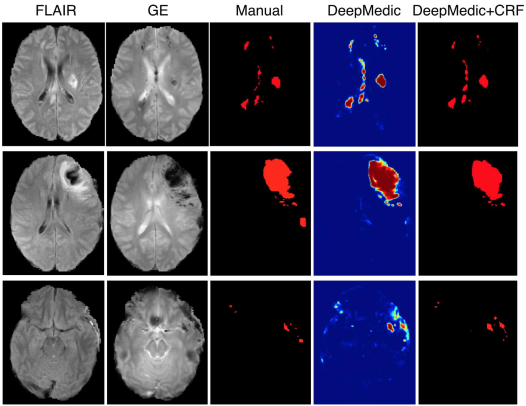

Taken from [KLN+17]

Outline

- Background - Information about basic structural concepts for CNNs
- A novel two pathway approach by Havaei, et al.
- 3D multi-scale approach by Kamnitsas, et al.
- Results
- Conclusions

Classification

- Classification is the process of identifying something
- In the case of images we might classify something as an image of a brain versus an image of a foot
- The name of these classifications is often referred to as ‘labels’

Biological Segmentation

- Segmentation is the process of identifying the boundaries of different structures and classifying them
- Segmentation is loosely defined and can have a wide range of granularities
  - Rough grained, such as identifying the different bones in a leg X-ray
  - Fine grained, such as determining the differing regions of a tumor

Neural Networks

- Neural Networks are a form of machine learning
- Neural Networks can be thought of as pattern recognizers.
- They are loosely based on the neuronal structure of the cerebral cortex, the part of the brain that takes in sensory data.

Neural Network Structure

- Comprised of layers of nodes
- Each node has an activation function that triggers when it recognizes something in the input
- These activations are then passed to neighboring nodes through weighted connections eventually leading to some type of output
- The network can ‘learn’ by altering the weights of its connections based on the accuracy of the output to the goal result

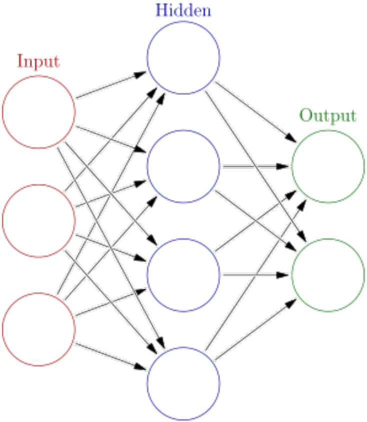
Hidden
Input
Output

Kernels

- Kernels, neurons and filters are interchangeable names
- Kernels are an array based representation of image features
- More kernels equals more recognizable features

Kernels

0	0	0	0	0	30	0
0	0	0	0	30	0	0
0	0	0	30	0	0	0
0	0	0	30	0	0	0
0	0	0	30	0	0	0
0	0	0	30	0	0	0
0	0	0	0	0	0	0

Pixel representation of filter

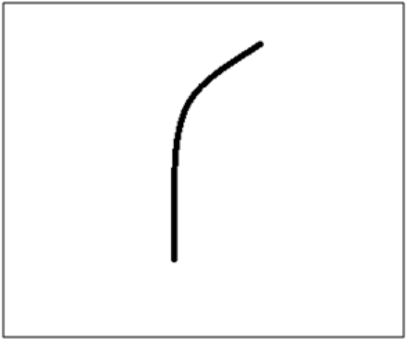

Visualization of a curve detector filter

Kernel
https://adeshpande3.github.io/

Convolutional Layers

- Convolutional Layers are where CNNs get their name
- Every CNN starts with a convolutional layer
- The kernel slides or 'convolves' around the input image
- The results of the convolutions are stored in the feature map

Convolutional Layers

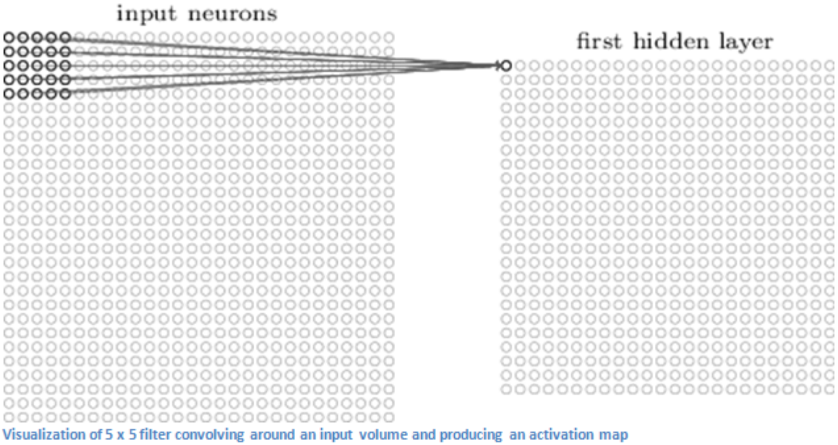

Feature Map
`https://adeshpande3.github.io/`

Fully Connected Layers

- Can be thought of as the final layers in the network
- Their job is to convert the feature maps from previous layers into label probabilities

Training

- Training is the crux that makes everything work
- Training requires data that has already been properly segmented
- Network is initialized with random kernel weights
- Training has four main steps:
  - The forward pass
  - The loss calculation
  - The backward pass
  - Weight update
- These four steps are performed on the entirety of the training data set multiple times

Overview of Havaei, et al.

- Havaei, et al. proposes a two pathway approach to the BRATS 2013 brain tumor segmentation challenge
- Their approach has three main components
  - Two pathways
  - Two CNNs concatenated together
  - A two phase approach to training

Two Pathways

- Havaei, et al. set up their network with two pathways
    - The local pathway with a smaller 7x7 pixel receptive field
    - The global pathway with a larger 13x13 pixel receptive field
- These two pathways allow the combination of fine detail with greater locational context

Two Pathways

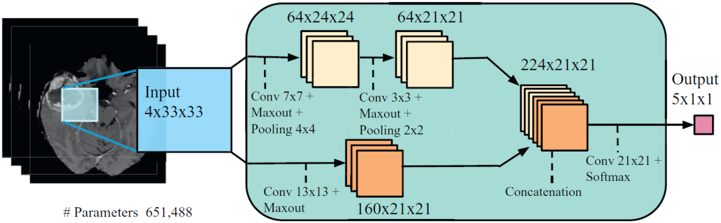

Taken from [HDWF^{+}17]

Two Phase Training

- The last approach implemented by Havaei, et al. is a two phase training system
- This is done to alleviate the relative abundance of healthy tissue versus the small quantity of tumor tissue in each image
- The two phases consist of:
    - First they train the network on image patches where the probability of each label being present is equal
    - Then they retrain the final layer with the relative probabilities of each label
- This allows for better label discrimination while maintaining proper output probabilities

Overview of Kamnitsas, et al.

- Kamnitsas, et al have five different architecture approaches
  - 3D kernels
  - Dense training
  - Two pathways
  - Deeper networks
  - 3D conditional random fields on the output
- These approaches lead to top performances in three different brain related segmentation challenges

Overview of Kamnitsas, et al.

- Kamnitsas, et al have five different architecture approaches
  - 3D kernels
  - Dense training
  - Two pathways
  - Deeper networks
  - 3D conditional random fields on the output
- These approaches lead to top performances in three different brain related segmentation challenges

3D Kernels

- 3D kernels can be thought of as 3 dimensional rectangular prisms
- Before the kernel convolved around a 2 dimensional space, but now it is convolving around a 3D space
- 3D kernels add to the computational costs
- Kamnitsas, et al. proposes a hybrid training scheme to resolve this

Two-Pathways

- Much like Havaei, et al. Kamnitsas, et al. use a two pathway approach
- These two pathways are meant to capture global and local context
- Unlike Havaei, et al. they downsample the input image for one of the pathways rather than change the size of the receptive field

Deeper Networks

- Kamnitsas, et al. also explore the use of deeper neural networks
- A deeper network has more consecutive layers
- Deeper networks increase the discriminative capability of CNNs
- A drawback is the increase in trainable parameters
- Kamnitsas, et al. address this by decreasing the size of the kernels, thus lowering the number of trainable parameters

Deeper Networks

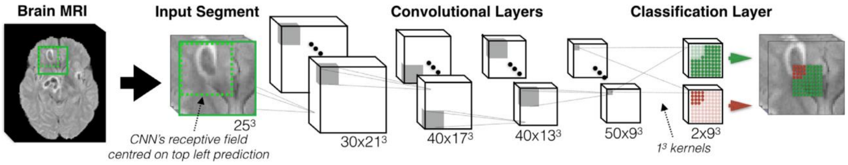

Taken from [KLN$^+$17]

Havaei, et al. Results

Name	Dice	Specificity	Sensitivity
InputCascadeCNN*	0.84	0.88	0.84
Tustison	0.79	0.83	0.81
Zhao	0.79	0.77	0.85
Meier	0.72	0.65	0.88
Reza	0.73	0.68	0.79
Cordier	0.75	0.79	0.78

Table: Comparison of Havaei, et al's. results on BRATS 2013 leaderboard set

Kamnitsas, et al. Results

Name	Dice	Precision	Sensitivity
Ensemble+CRF	90.1	91.9	89.1
Ensemble	90.0	90.3	90.4
DeepMedic+CRF	89.8	91.5	89.1
DeepMedic	89.7	89.7	90.5
bakas1	88	90	89
peres1	87	89	86
anon1	84	90	82
thirs1	80	84	79
peyrj	80	87	77

Table: Average performance of Kamnitsas, et al. on the training data from BRATS 2015 compared to other teams

Traumatic Brain Injury Example

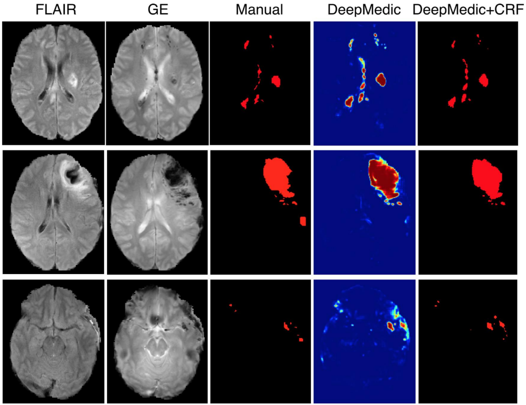

Taken from [KLN+17]

Acknowledgements

I would like to thank Nic McPhee and Elena Machkasova for their guidance throughout this process. I would also like to thank my parents for their support.

References

Mohammad Havaei, Axel Davy, David Warde-Farley, Antoine Biard, Aaron Courville, Yoshua Bengio, Chris Pal, Pierre-Marc Jodoin, and Hugo Larochelle, *Brain tumor segmentation with deep neural networks*, Medical Image Analysis **35** (2017), 18 – 31.

Konstantinos Kamnitsas, Christian Ledig, Virginia F.J. Newcombe, Joanna P. Simpson, Andrew D. Kane, David K. Menon, Daniel Rueckert, and Ben Glocker, *Efficient multi-scale 3d {CNN} with fully connected {CRF} for accurate brain lesion segmentation*, Medical Image Analysis **36** (2017), 61 – 78.