Deep Learning in Medical Physics—LESSONS We Learned

Hui Lin
PhD candidate
Rensselaer Polytechnic Institute, Troy, NY
07/31/2017

Acknowledgements

- My PhD advisor – Dr. George Xu at RPI
- Dr. Chengyu Shi, Dr. Lily Tang at MSKCC and Dr. Tianyu Liu at RPI have made important contributions
- Nvidia for the donation of GPUs

2

Outline

- **<u>Part I – What's going on in Deep Learning that may matter to Medical Physics</u>**
  - ❑ConvNet
  - ❑RNN
  - ❑GAN
  - ❑Frameworks for Deep Learning
- **Part II – Practices of Deep Learning in Medical Physics – lessons we've learnt**
  - ❑ConvNet for Lung Cancer Detection
  - ❑ConvNet for Organ Segmentation
  - ❑RNN for EHR Mining
- **Part III – Concluding Remarks**

3

Deep Learning – Great Breakthroughs in AI

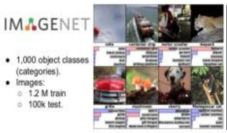

- Enormous data + Adequate computing power = Deep Learning Revolution!
- Some popular strategies of Deep Learning:
  - ❑Convolutional Neural Networks (CNN)
  - ❑Recurrent Neural Networks (RNN)
  - ❑Generative Adversarial Networks (GAN)

Images courtesy of [1-3]

4

Part I – What's going on in Deep Learning

- Convolutional Neural Networks (ConvNet) in a nutshell
  - ❑State-of-the-art for vision perception tasks
  - ❑Usually comprised as following stages:
    Input → Convolutional Layer → Pooling Layer → Convolutional Layer → Pooling Layer → Fully-connected Layer → Output
  - ❑Higher levels of the model can detect more abstract features that are useful for image recognition
  - ❑Useful to clinical applications involve patient anatomical images
    - ❖Automated anatomy classification
    - ❖Cancer/Nodule detection
    - ❖OAR segmentation

5

Part I – What's going on in Deep Learning

- ConvNet – a quick example
  - ❑Convolution: Try every possible feature and make one image into a stack of filtered images

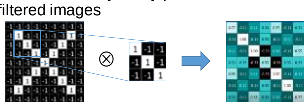

  - ❑Pooling: Shrink the image stack while preserving the important information

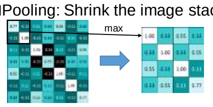

  - ❑Fully-connected: Output probability vector of N classes

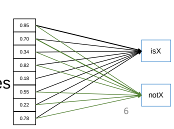

Images courtesy of [4]

6

Part I – What's going on in Deep Learning

- Recurrent Neural Networks(RNN) in a nutshell

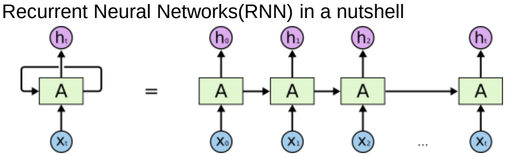

- ❑Current output depends on the previous inputs/outputs
- ❑Designed to solve problems involving temporal processing and sequential learning
- ❑Useful to clinical applications involve sequential data
  - ❖Electronic Health Records mining
  - ❖Respiration curve monitoring and prediction

Images courtesy of [5]

7

Part I – What's going on in Deep Learning

- RNN struggled with vanishing or exploding gradient problems
- Long-Short Term Memory (LSTM)
  - ❑Able to avoid the long-term dependency problem
  - ❑Memory cell contains:
    - ❖Input gate
    - ❖Forget gate
    - ❖Output gate

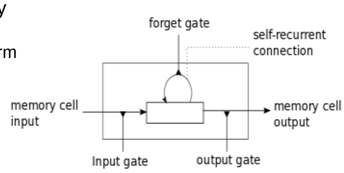

[6] http://deeplearning.net/tutorial/lstm.html

8

Part I – What's going on in Deep Learning

- Generative Adversarial Networks(GAN) – "the coolest idea in ML in last 10 years"

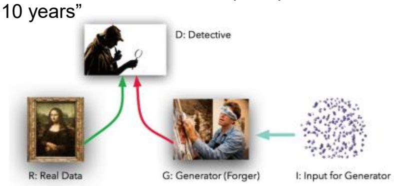

- ❑Generator (G): mimics examples from a training data that can fool D
- ❑Discriminator (D): predicts whether it is a data sample or a "forged" one
- ❑Both D and G improves over time until G can generate genuine sample and G is at loss unable to figure out the distribution differences

Image courtesy of [7]

9

Part I – What's going on in Deep Learning

- Popular Deep Learning frameworks
  - ❑Caffe and Caffe2
  - ❑Tensorflow
  - ❑Torch
  - ❑CNTK
  - ❑Theano

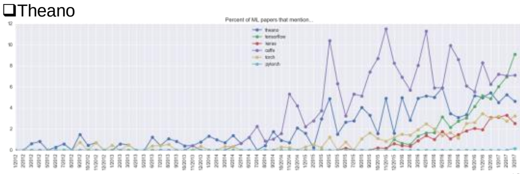

10

Image courtesy of [8]

Outline

- **Part I – What's going on in Deep Learning that may matter to Medical Physics**
  - ❑ConvNet
  - ❑RNN
  - ❑GAN
  - ❑Frameworks for Deep Learning
- **Part II – Practices of Deep Learning in Medical Physics – lessons we've learnt**
  - ❑ConvNet for Lung Cancer Detection
  - ❑ConvNet for Organ Segmentation
  - ❑RNN for EHR Mining
- **Part III – Challenges and Potential Trends of Deep Learning**

11

Part II – ConvNet for Lung Cancer Detection

- Objective: Look through a patient thoracic CT set and predict if is cancerous
- Data – Kaggle Data Science Bowl 2017[9]
  - ❑Training data 1397 sets, validation data ~300 sets

- Workflow

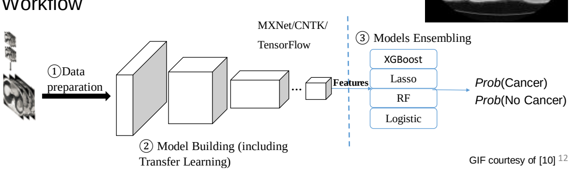

GIF courtesy of [10] 12

Part II – ConvNet for Lung Cancer Detection

- Data Preparation
  - ❑Data Augmentation – Rotation, small translation, zooming etc.
- Transfer Learning
  - ❑Pre-trained ConvNet as an initialization or a fixed feature extractor for our own task
  - ❑Only the last CNN blocks are fine-tuned to avoid overfitting
- Ensembling
  - ❑Model parameters need to be optimized via k-fold cross validation
- Performance Metrics
  - ❑AUROC
  - ❑Logloss

Part II – ConvNet for Lung Cancer Detection

- Performance

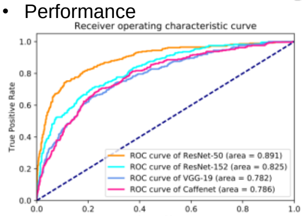

Procedure	Logloss
Baseline	0.62310
Data Augmentation	0.58278
Transfer Learning	0.47921
Model Ensembling	0.47137
All stacked	0.47137

	CaffeNet	ResNet-50	VGG-19	ResNet-152
Time (min)	56	75	78	172
Number of extracted features	9216	2048	25088	2048
Memory per patient (GB)	0.07	0.75	0.38	0.88

Part II – ConvNet for Medical Image Segmentation

- Objective: Perform pixel-wise classification to segment left ventricle from cardiac MR
- Data: Sunnybrook Cardiac MR dataset[11], 238 training data set

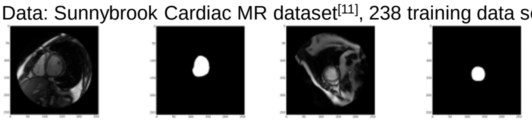

- Deep Learning strategies
  - ❑ Replace the fully-connected layers with deconvolution layers to output segmentation results

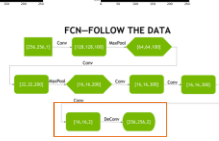

Part II – ConvNet for Medical Image Segmentation

- Deep Learning strategies
  - ❑ Transfer Learning: Convert a pre-trained complex ConvNet to FCN
  - ❑ Accuracy and loss are not enough: class imbalance → add a custom layer Dice metric to the ConvNet
  - ❑ Hyper parameters tuning
    - ❖ Adjustable learning rate etc.
- Output
  - ❑ Dice metric evaluation: 88.6%

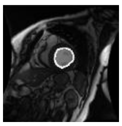
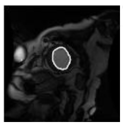
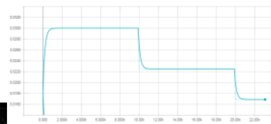

Segmentation plots were generated upon [12]

16

Part II – RNN for Electronic Health Records Mining

- Objective: early warnings of the severity of a patient's illness
- Data
  - ❑5000 ICU patients Electronic Health Record data in HDF5 format[13]
  - ❑Including statics, vitals, labs, interventions, drugs and outcome
  - ❑Heterogeneous, incomplete and redundant

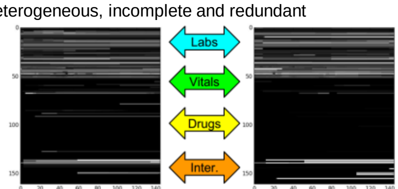

Image courtesy of [13]

17

Part II – RNN for Electronic Health Records Mining

- Data Preparation
  - ❑Data normalization: make sure small variables can be treated with the same emphasis as the large variables
  - ❑Data gaps filling
    - ❖Fill existing measurements forwardly for each patient
    - ❖Fill variable entries with no previous measurement to 0
  - ❑Data padding: force each patient record of dimension 500x265 and use zero padding to inflate the size if needed

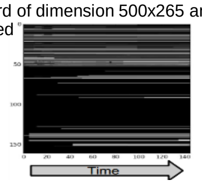

Image courtesy of [13]

18

Part II – RNN for Electronic Health Records Mining

- Deep Learning Strategies
  - ❑Model design

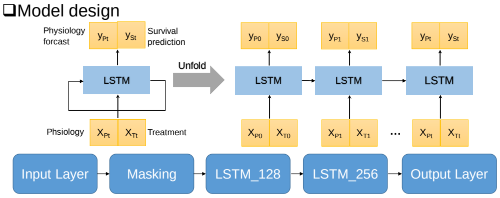

  - ❑Model Evaluation
    - ❑AUROC
    - ❑Comparisons against baseline models like PRISM3[14] and PIM2[15]

Part II – RNN for Electronic Health Records Mining

- Performance
  - ❑Able to output survivability prediction per patient
  - ❑Superior accuracy against classic models
    LSTM AUROC 0.920 Vs. PIM2 0.872

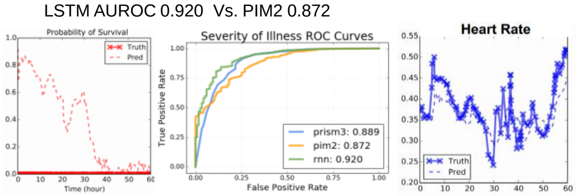

  - ❑Instantaneous prediction of survivability provides valuable feedback to assess the impact of treatment decisions

Images courtesy of [13]

Part III – Concluding Remarks

- Many exciting advancements empowered by Deep Learning are going on in Healthcare
- Deep Learning holds a great potential waiting to be exploited in the application of Medical Physics
- Challenges and Potential Solutions
  - ❑Lack of high-quality annotated medical data
    - ❖ ImageNet in Medical Physics
    - ❖ Text-image joint mining
  - ❑Interpretability of deep learning neural networks
    - ❖ Close collaborations between clinicians and data scientists

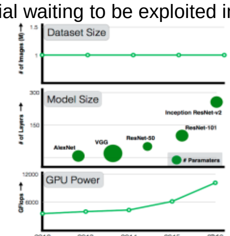

Image courtesy of [16]

References

[1] Deep Learning for Computer Vision: ImageNet Challenge (UPC 2016), accessed on May 10th, 2017. https://www.slideshare.net/xavigiro/deep-learning-for-computer-vision-imagenet-challenge-upc-2016

[2] Driving Innovation, accessed on May 10th, 2017. http://www.nvidia.com/object/drive-automotive-technology.html

[3] Can AlphaGo defeat Lee Sedol, accessed on May 10th, 2017. https://gogameguru.com/can-alphago-defeat-lee-sedol/

[4] How convolutional neural networks work, accessed on July 20th, 2017. https://brohrer.github.io/how_convolutional_neural_networks_work.html

[5] Understanding LSTM Networks, accessed on July 15th, 2017. http://colah.github.io/posts/2015-08-Understanding-LSTMs

[6] LSTM Networks for Sentiment Analysis, accessed on July 10th, 2017. http://deeplearning.net/tutorial/lstm.html

[7] Generative Adversarial Networks (GANs) in 50 lines of code (PyTorch), accessed on July 16th, 2017. Convolutional Neural Networks, accessed on 22nd April, 2017. https://medium.com/@devnag/generative-adversarial-networks-gans-in-50-lines-of-code-pytorch-e81b79659e3f

[8] A Peek at Trends in Machine Learning, accessed on July 19th, 2017. https://medium.com/@karpathy/a-peek-at-trends-in-machine-learning-ab8a1085a106

[9] Data Science Bowl 2017, accessed on July 20th, 2017. https://www.kaggle.com/c/data-science-bowl-2017

[10] Accessed on July 20th, 2017. http://blog.kaggle.com/2017/06/29/2017-data-science-bowl-predicting-lung-cancer-2nd-place-solution-write-up-daniel-hammack-and-julian-de-wit/

[11] Radau P, Lu Y, Connelly K, Paul G, Dick AJ, Wright GA. "Evaluation Framework for Algorithms Segmenting Short Axis Cardiac MRI." The MIDAS Journal – Cardiac MR Left Ventricle Segmentation Challenge, http://hdl.handle.net/10380/3070

[12] Medical Image Segmentation Using DIGITS, accessed on June 4th, 2017. https://nvidia.qwiklab.com/focuses/preview/2794

[13] David Ledbetter, Melissa Aczon, Randall Wetzel. "Deep learning Recommendation of Treatment from Electronic Data." GTC 2016

[14] PRISM3, accessed on July 20th, 2017. http://www.jpeds.com/article/S00223476(97)700659/abstract

[15] Slater, A., Shann, F., Pearson, G. and PIM Study Group, 2003. PIM2: a revised version of the Paediatric Index of Mortality. *Intensive care medicine*, *29*(2), pp.278-285.

[16] Revisiting the Unreasonable Effectiveness of Data, accessed on July 27th, 2017. https://research.googleblog.com/2017/07/revisiting-unreasonable-effectiveness.html

22

Thank you!

Rensselaer

Hui Lin, PhD candidate
Nuclear Engineering and Science
3021 Tibbits Ave, Troy, NY, 12180
linh7@rpi.edu

23

Backup

What's going on in Deep Learning

- Why GAN?
  - ❑Can represent and manipulate high-dimensional probabilitydistributions
  - ❑Can be trained with missing data – semi-supervised learning
  - ❑Allows multi-modal outputs

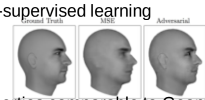

- Some interesting applications
  - ❑Generate CT images based on MRI[4]
  - ❑Reproduce electromagnetic shower properties comparable to Geant4[5]