Can deep learning help in cancer diagnosis and treatment?

Lubomir Hadjiiski, PhD

University of Michigan
Medical School

The University of Michigan
Department of Radiology

Outline

- **Applications of Deep Learning to Medical Imaging**
- Transfer Learning
- **Some practical aspects of Deep Learning application to Medical Imaging**

Neural network (NN)

Input node

Output node

Convolution neural network (CNN)

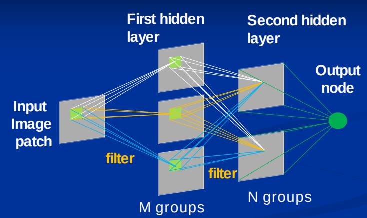

CNN in medical applications

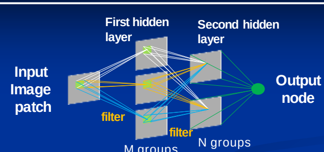

Chan H-P, Lo S-C B, Helvie MA, Goodsitt MM, Cheng S and Adler D D 1993 Recognition of mammographic microcalcifications with artificial neural network, RSNA Program Book 189(P) 318
Chan H-P, Lo S-C B, Sahiner B, Lam KL, Helvie MA, Computer-aided detection of mammographic microcalcifications: Pattern recognition with an artificial neural network, 1995, Medical Physics
Lo S-C B, Chan H-P, Lin J-S, Li H, Freedman MT and Mun SK 1995, Artificial convolution neural network for medical image pattern recognition, Neural Netw. 8 1201-14
Sahiner B, Chan H-P, Petrick N, Wei D, Helvie MA, Adler DD and Goodsitt MM 1996 Classification of mass and normal breast tissue: A convolution neural network classifier with spatial domain and texture images, IEEE Trans.on Medical Imaging 15
Samala RK, Chan H-P, Lu Y, Hadjiiski L, Wei J, Helvie MA, Digital breast tomosynthesis: computer-aided detection of clustered microcalcifications on planar projection images, 2014 Physics in Medicine and Biology

Deep Learning

- Deep-Learning Convolutional Neural Network (DL-CNN)

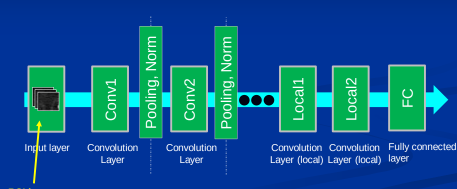

Krizhevsky et al, 2012, Adv. Neu. Info. Proc. Sys.

Deep Learning

- Task – to label voxels as inside or outside of bladder
- Large number of convolution kernels and weights
- Trained with 160,000 ROIs

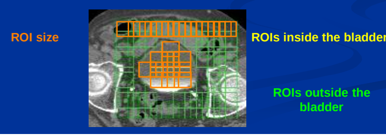

Deep Learning

Examples of training ROIs

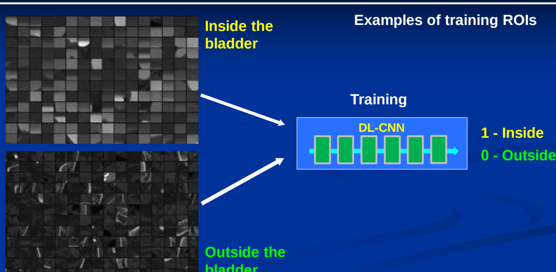

Deep Learning

Bladder Likelihood Map Generation –
Voxel-by-voxel labeling with trained DL-CNN

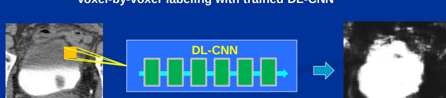

Original Image

Bladder Likelihood Map

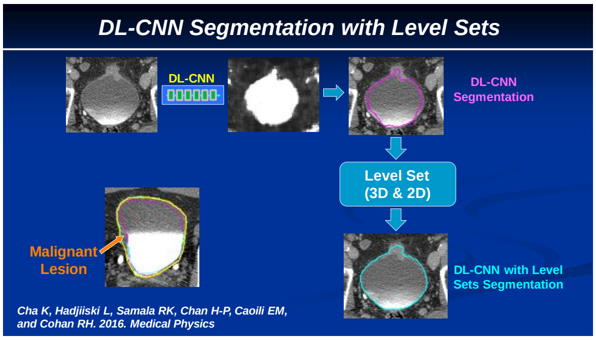
DL-CNN Segmentation with Level Sets
DL-CNN
DL-CNN Segmentation
Level Set (3D & 2D)
Malignant Lesion
DL-CNN with Level Sets Segmentation
Cha K, Hadjiiski L, Samala RK, Chan H-P, Caoili EM, and Cohan RH. 2016. Medical Physics

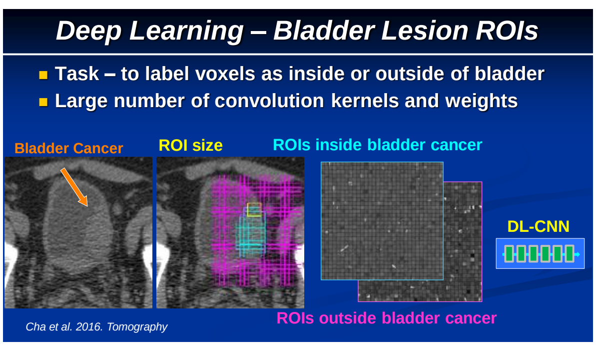
Deep Learning – Bladder Lesion ROIs
Task – to label voxels as inside or outside of bladder
Large number of convolution kernels and weights
Bladder Cancer
ROI size
ROIs inside bladder cancer
DL-CNN
ROIs outside bladder cancer
Cha et al. 2016. Tomography

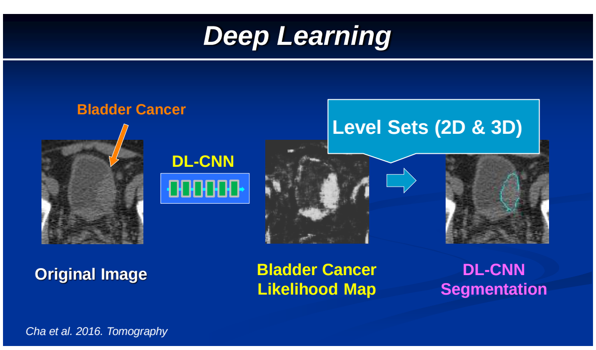
Deep Learning
Bladder Cancer
DL-CNN
Level Sets (2D & 3D)
Original Image
Bladder Cancer Likelihood Map
DL-CNN Segmentation
Cha et al. 2016. Tomography

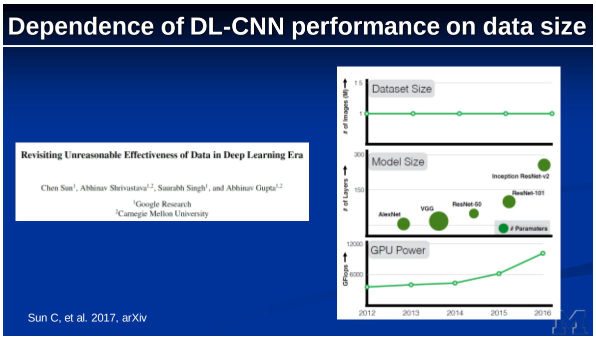
Dependence of DL-CNN performance on data size
Revisiting Unreasonable Effectiveness of Data in Deep Learning Era
Chen Sun1, Abhinav Shrivastava1,2, Saurabh Singh1, and Abhinav Gupta1,2
1Google Research
2Carnegie Mellon University
Dataset Size
Model Size
AlexNet
VGG
ResNet-50
ResNet-101
Inception ResNet-v2
Parameters
GPU Power
2012
2013
2014
2015
2016
Sun C, et al. 2017, arXiv

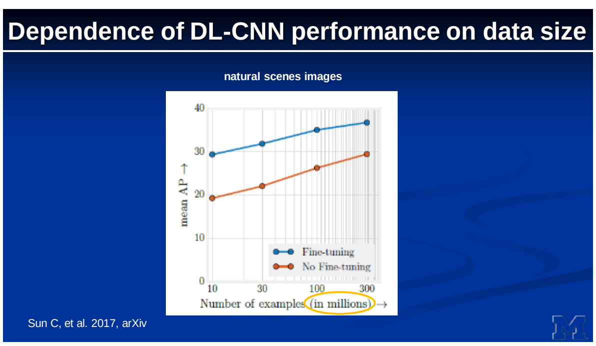
Dependence of DL-CNN performance on data size
natural scenes images
mean AP →
Fine-tuning
No Fine-tuning
Number of examples (in millions) →
Sun C, et al. 2017, arXiv

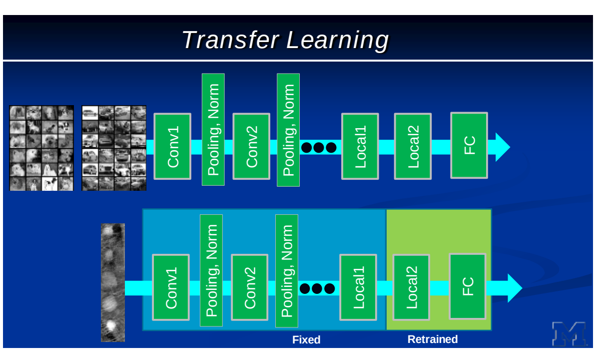
Transfer Learning
Conv1
Pooling, Norm
Conv2
Pooling, Norm
Local1
Local2
FC
Fixed
Retrained

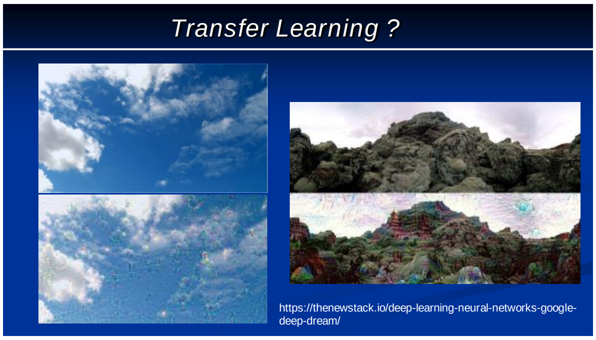
Transfer Learning ?
https://thenewstack.io/deep-learning-neural-networks-google-deep-dream/

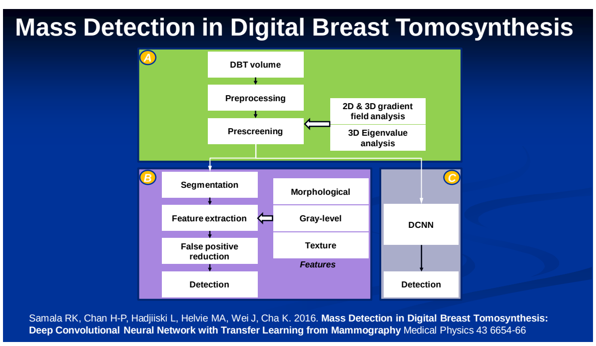
Mass Detection in Digital Breast Tomosynthesis
A
DBT volume
Preprocessing
Prescreening
2D & 3D gradient field analysis
3D Eigenvalue analysis
B
Segmentation
Feature extraction
False positive reduction
Detection
Morphological
Gray-level
Texture
Features
C
DCNN
Detection
Samala RK, Chan H-P, Hadjiiski L, Helvie MA, Wei J, Cha K. 2016. Mass Detection in Digital Breast Tomosynthesis: Deep Convolutional Neural Network with Transfer Learning from Mammography Medical Physics 43 6654-66

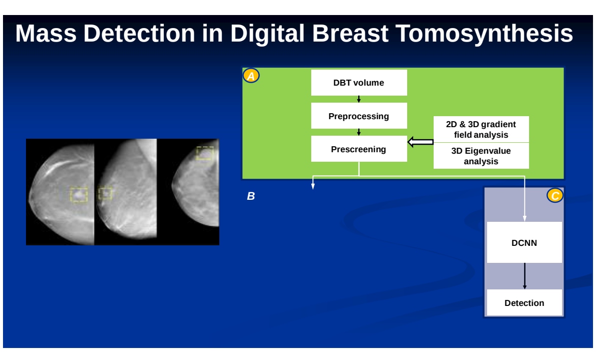
Mass Detection in Digital Breast Tomosynthesis
A
DBT volume
Preprocessing
Prescreening
2D & 3D gradient field analysis
3D Eigenvalue analysis
B
C
DCNN
Detection

Transfer Learning DL-CNN

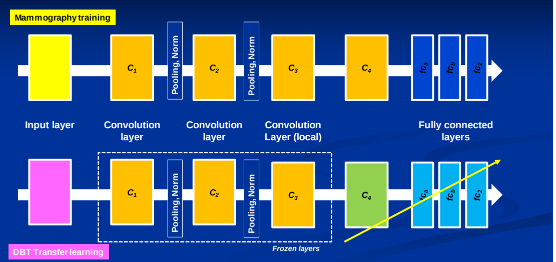

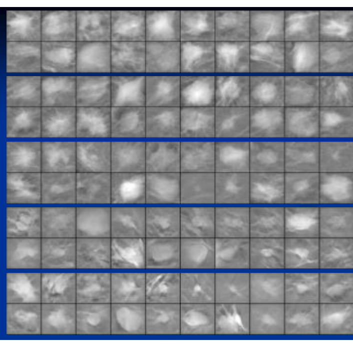

Transfer Learning DL-CNN

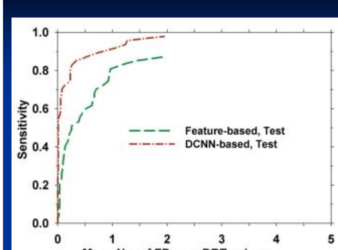

	Conventional CAD	DL-CNN Trained on mammograms	DL-CNN Transfer learning
Sensitivity	0.70	0.83	0.91
False Positive/ DBT Volume	1	1	1

($p < 0.05$)

Transfer Learning

Mammography		
Dhungel et al 2017[2]	Pre-training CNN on handcrafted features, feature extraction, random forest classifier for classification of masses	115 cases, 60% training, 20% validation and 20% test (5-fold cross-validation)
Kooi et al 2017[5]	TL from VGG network for classification of solitary cysts in mammography	30,184 digitized screen-film exams for first stage training, 956 digital exams for second stage train/test
Samala et al 2016[7]	Transfer learning from mammography to detection of masses in digital breast tomosynthesis	2282 mammography views for pretraining, 324 DBT volumes from 232 patients for training and independent test
Huynh et al 2016[10]	Feature extraction from pretrained AlexNet and SVM classifier for classification of masses on digital mammography	219 lesions from 190 cases from Li et al[11]
Lung		
Christodoulidis et al 2017[3]	Ensemble of CNNs, TL from six publicly available texture databases, classification of lung tissue patterns	135 CT scans from Interstitial lung disease cases
Paul et al 2016[6]	Three variations of VGG pretrained CNNs	Chest CT scans from 81 patients
Shin et al[13]	TL from CifarNet, AlexNet and GoogLeNet for thoraco-abdominal lymph node (LN) detection and interstitial lung disease classification	LN datasets: 90 patient CT scans ILD: 120 patients
Colon		
Zhang et al 2017[4]	Two public non-medical image databases for TL	291 polyps cases
Brain, Abdomen		
Han 2017[1]	TL from VGG network trained on natural scene image, Synthetic CT generation from MR, U-net architecture	18 patient cases, 6-fold cross-validation
Cheng et al 2016[8]	Multi-domain transfer from joint learning of multiple auxiliary tasks for diagnosis of Alzheimer's disease	807 subjects from public database
Cheng et al 2016[9]	Transfer learning from CaffeNet and VGG trained on natural scene images for classification of abdominal ultrasound	185 cases

Transfer Learning

Pulmonary Embolism, Other		
Tajbakhsh et al 2016[12]	TL from AlexNet for four medical imaging tasks	Polyp detection: 40 short colonoscopy videos Image quality of colonoscopy videos: 6 complete colonoscopy videos Pulmonary embolism detection in CT: 121 cases Segmentation in ultrasonographic images: 92 Carotid intima-media thickness videos

DL-CNN Classification of Breast Masses

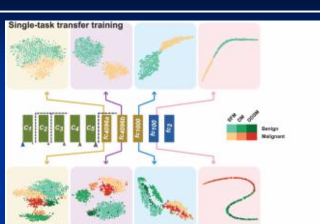

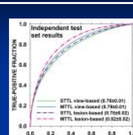

STTL: Single-task transfer learning
MTTL: Multi-task transfer learning

Samala RK, Chan HP, Hadjiiski L, Helvie MA, Cha K, Richter C. Physics in Medicine and Biology. (Submitted)
Samala RK, Chan HP, Hadjiiski L, Cha K, Helvie MA, Richter C. Accepted to the RSNA meeting, 2017.

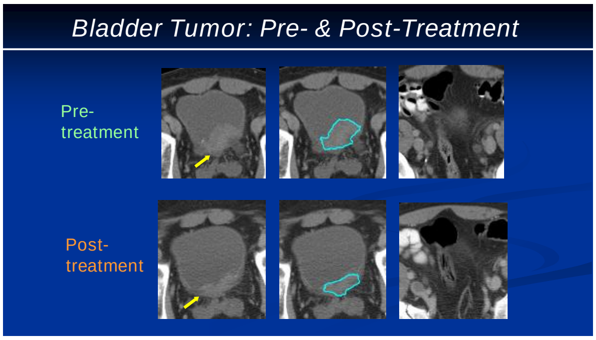
Bladder Tumor: Pre- & Post-Treatment
Pre-
treatment
Post-
treatment

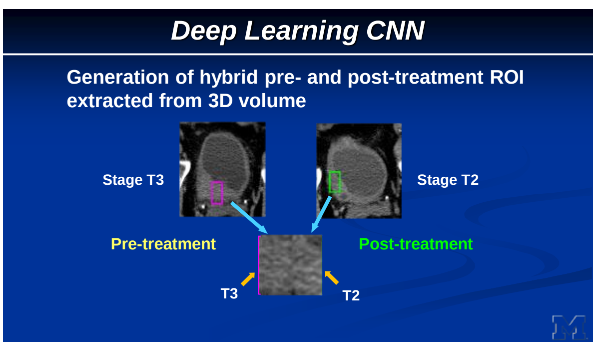
Deep Learning CNN
Generation of hybrid pre- and post-treatment ROI
extracted from 3D volume
Stage T3
Stage T2
Pre-treatment
Post-treatment
T3
T2

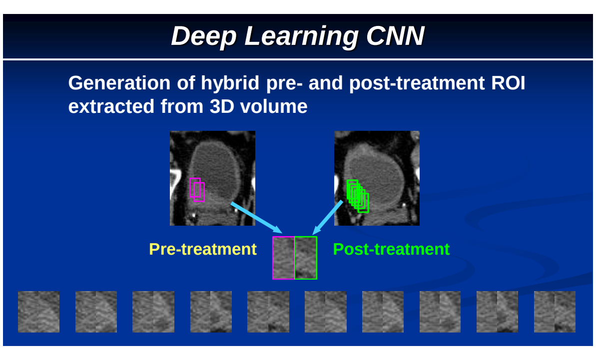
Deep Learning CNN
Generation of hybrid pre- and post-treatment ROI
extracted from 3D volume
Pre-treatment
Post-treatment

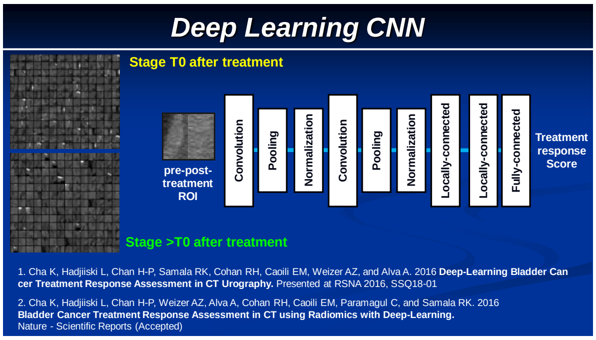
Deep Learning CNN
Stage T0 after treatment
pre-post-treatment ROI
Convolution
Pooling
Normalization
Convolution
Pooling
Normalization
Locally-connected
Locally-connected
Fully-connected
Treatment response Score
Stage >T0 after treatment
1. Cha K, Hadjiiski L, Chan H-P, Samala RK, Cohan RH, Caoili EM, Weizer AZ, and Alva A. 2016 Deep-Learning Bladder Cancer Treatment Response Assessment in CT Urography. Presented at RSNA 2016, SSQ18-01
2. Cha K, Hadjiiski L, Chan H-P, Weizer AZ, Alva A, Cohan RH, Caoili EM, Paramagul C, and Samala RK. 2016 Bladder Cancer Treatment Response Assessment in CT using Radiomics with Deep-Learning. Nature - Scientific Reports (Accepted)

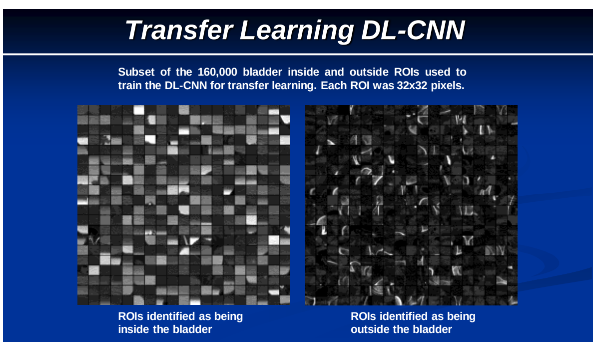
Transfer Learning DL-CNN
Subset of the 160,000 bladder inside and outside ROIs used to train the DL-CNN for transfer learning. Each ROI was 32x32 pixels.
ROIs identified as being inside the bladder
ROIs identified as being outside the bladder

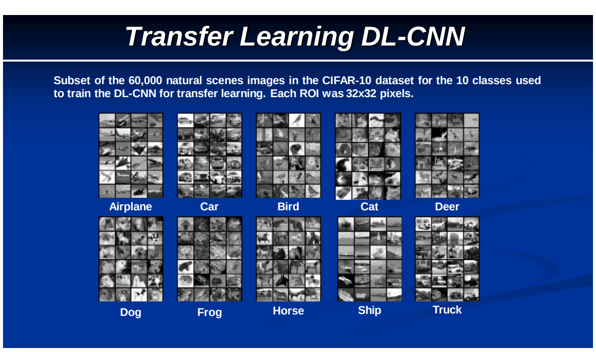
Transfer Learning DL-CNN
Subset of the 60,000 natural scenes images in the CIFAR-10 dataset for the 10 classes used to train the DL-CNN for transfer learning. Each ROI was 32x32 pixels.
Airplane
Car
Bird
Cat
Deer
Dog
Frog
Horse
Ship
Truck

Transfer Learning DL-CNN

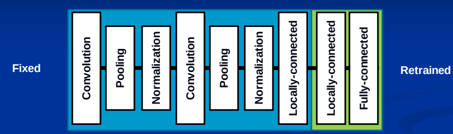

Cha K, Hadjiiski L, Chan H-P, Samala RK, Cohan RH, Caoili EM, Paramagul C, Alva A, and Weizer AZ. 2017. Proceedings of SPIE Medical Imaging 10134 1013404

Results – Leave-one-case-out

	AUC
DL-CNN	0.75 ± 0.05
DL Bladder Transfer Learning	0.72 ± 0.05
DL Natural Scene Transfer Learning	0.68 ± 0.06
Radiologist 1	0.70 ± 0.06
Radiologist 2	0.75 ± 0.05

Convolution activation layer output

Bladder In/Out Learned Kernel

Bladder Pre-Post Learned Kernel

CIFAR Learned Kernel

Convolution activation layer output

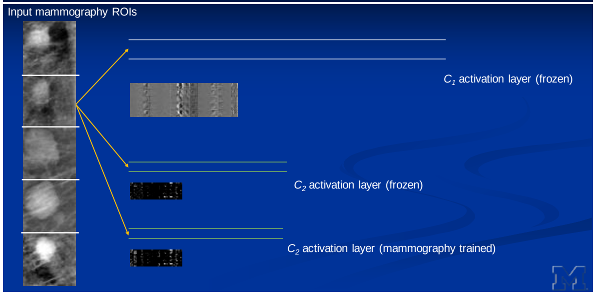

Dependence of DL-CNN performance on data size

Dependence of DL-CNN performance of breast masses on mammography data (100% - 2282 views)

C_1 frozen

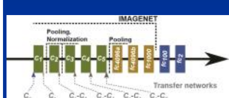

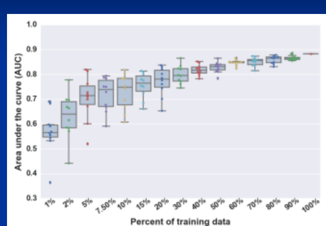

Samala RK, Chan H-P, Hadjiiski L, et al

Transfer Learning

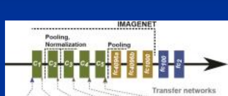

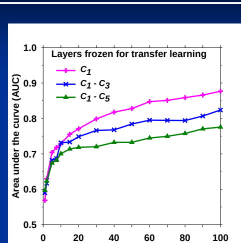

Samala RK, Chan H-P, Hadjiiski L, et al

Transfer Learning

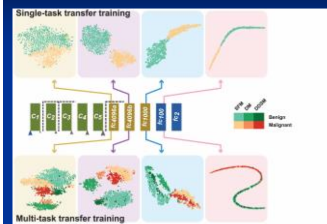

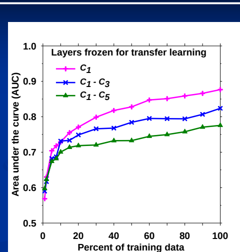

Samala RK, Chan H-P, Hadjiiski L, et al

DL-CNN Packages

Software	Creator	Software license[a]	Open source	Platform	Written in	Interface	OpenMP support	OpenCL support	CUDA support	Automatic differentiation	Has pretrained models	Recurrent nets	Convolutional nets	RBM/DBNs	Parallel execution (multi node)
Apache SINGA	Apache Incubator	Apache 2.0	Yes	Linux, Mac OS X, Windows	C++	Python, C++, Java	No	Yes	Yes	?	Yes	Yes	Yes	Yes	Yes
Caffe	Berkeley Vision and Learning Center	BSD license	Yes	Linux, Mac OS X, Windows[2]	C++	Python, MATLAB	Yes	Under development	Yes	Yes	Yes	Yes	Yes	No	?
Deeplearning4j	Skymind engineering team; Deeplearning4j community; originally Adam Gibson	Apache 2.0	Yes	Linux, Mac OS X, Windows, Android (Cross-platform)	java	Java, Scala, Clojure, Python (Keras)	Yes	On roadmap	Yes	Computational Graph	Yes	Yes	Yes	Yes	Yes
Dlib	Davis King	Boost Software License	Yes	Cross-Platform	C++	C++	Yes	No	Yes	Yes	Yes	No	Yes	Yes	Yes
Keras	François Chollet	MIT license	Yes	Linux, Mac OS X, Windows	Python	Python, R	Only if using Theano as backend	Under development for the Theano backend (and on roadmap for the TensorFlow backend)	Yes	Yes	Yes	Yes	Yes	Yes	Yes
MatConvNet	Andrea Vedaldi,Karel Lenc	BSD license	Yes	Windows, Linux[11] (OSX via Docker on roadmap)	C++	MATLAB, C++,	No	No	Yes	Yes	Yes	Yes	Yes	No	Yes

https://en.wikipedia.org/wiki/Comparison_of_deep_learning_software

DL-CNN Packages

Software	Creator	Software license	Open source	Platform	Written in	Interface	OpenMP support	OpenCL support	CUDA support	Automatic differentiation[1]	Has pretrained models	Recurrent nets	Convolutional nets	RBM/DBNs	Parallel execution (multi node)
Microsoft Cognitive Toolkit	Microsoft Research	MIT license[12]	Yes	Windows, Linux[13] (OSX via Docker on roadmap)	C++	Python, C++, Command line,[14] BrainScript[15] (.NET on roadmap[16])	Yes	No	Yes	Yes	Yes	Yes	Yes	No	Yes
MXNet	Distributed (Deep) Machine Learning Community	Apache 2.0	Yes	Linux, Mac OS X, Windows,[22][23] AWS, Android,[24] iOS, JavaScript[25]	Small C++ core library	C++, Python, Julia, Matlab, JavaScript, Go, R, Scala, Perl	Yes	On roadmap	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Neural Designer	Artelnics	Proprietary	No	Linux, Mac OS X, Windows	C++	Graphical user interface	Yes	No	No	?	?	No	No	No	?
OpenNN	Artelnics	GNU LGPL	Yes	Cross-platform	C++	C++	Yes	No	No	?	?	No	No	No	?
TensorFlow	Google Brain team	Apache 2.0	Yes	Linux, Mac OS X, Windows[30]	C++, Python	Python (Keras), C/C++, Java, Go, R[31]	No	On roadmap	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Theano	Université de Montréal	BSD license	Yes	Cross-platform	Python	Python	Yes	Under development	Yes	Yes[37][38]	Through Lasagne's model zoo	Yes	Yes	Yes	Yes
Torch	Ronan Collobert, Koray Kavukcuoglu, Clement Farabet	BSD license	Yes	Linux, Mac OS X, Windows,[41] Android,[42] iOS	C, Lua	Lua, LuaJIT,[43] C, utility library for C++/OpenCL[44]	Yes	Third party implementations[45][46]	Yes[47][48]	Through Twitter's Autograd[49]	Yes	Yes	Yes	Yes	Yes
Wolfram Mathematica	Wolfram Research	Proprietary	No	Windows, Mac OS X, Linux, Cloud computing	C++	Wolfram Language	No	No	Yes	Yes	Yes	Yes	Yes	Yes	Yes

https://en.wikipedia.org/wiki/Comparison_of_deep_learning_software

DL-CNN Packages

TensorFlow ***Caffe***

Software	Creator	Software license	Open source	Platform	Written in	Interface	CUDA support	Automatic differentiation[1]	Has pretrained models	Recurrent nets	Convolutional nets	Parallel execution (multi node)
Caffe	Berkeley Vision and Learning Center	BSD license	Yes	Linux, Mac OS X, Windows	C++	Python, MATLAB	Yes	Yes	Yes	Yes	Yes	?
TensorFlow	Google Brain team	Apache 2.0	Yes	Linux, Mac OS X, Windows	C++, Python	Python (Keras), C/C++, Java, Go, R	Yes	Yes	Yes	Yes	Yes	Yes

https://en.wikipedia.org/wiki/Comparison_of_deep_learning_software

Deep-Learning

- Hardware requirements:
  - Graphics Processing Unit (GPU) with specific CUDA (programming language) compatibility
  - Enough system memory (RAM) to hold image data
- Software requirements:
  - Operating system – a specific Linux distribution
  - CUDA software and GPU drivers
  - Additional dependencies as required by DL software

Deep-Learning

- Input Image Preparation:
  - Size of the images – dependent on the task
  - Distribution of image values – system may train better with a specific image distribution (min, max, peak)
- Input Data File Preparation:
  - Create data format required by DL package
    - Cuda-convnet – pkl
    - Caffe – HDF5
    - Tensorflow – binary
  - Contains both image and reference truth data
  - Can use existing code available online (ie Python libraries)

Deep-Learning

- **Running the network:**
  - Determine network parameters by changing a parameter file or through GUI
  - Tutorials and examples online
- **Training parameters:**
  - Structure – number of layers and nodes
  - Learning rate and bias
- **Testing – "Deployment"**
  - Apply trained model to test data set to get results

Deep-Learning

- **Visualization:**
  - Packages contain general functions to visualize parts of the network
  - May need to write a program to extract specific parts
- **Results:**
  - Validation and interpretation of the results very important to develop useful models

Conclusions

- **Deep Learning is promising approach for Medical Imaging applications**
- Transfer Learning is important technique for applications with small datasets
- Transfer Learning still needs sufficient data for robust training

Acknowledgments

CAD Research Lab :

Heang-Ping Chan, PhD
Ravi Samala, PhD
Kenny Cha, MS

Clinical Faculty

Richard Cohan, MD
Elaine Caoili, MD
Anon Weizer, MD
Ajjai Alva, MD
Mark Helvie, MD
Marilyn Roubidoux, MD
Chintana Paramagul, MD