Informatics Tools for Optimized Imaging Biomarkers for Cancer Research & Discovery

Ken Chang, Jayashree Kalpathy-Cramer, PhD

Athinoula A. Martinos Center for Biomedical Imaging

MGH/Harvard Medical School

MICCAI 2018

Why “challenges”?

- Reproducibility is an issue in all aspects of medicine
- Algorithm performance often not replicated by other sites
- Access to clinical data of sufficient variety can be a challenge for (computational) scientists developing algorithms
- Can evaluate the performance of techniques on real, noisy clinical data
- Test data (sequestered) can provide indication of algorithm generalizability to unseen data
- Allows for cross-pollination of methods from other domains
- Best algorithms can be translated into commercial products

C-BIBOP challenge infrastructure

- Developed (in part) with funding from contracts and NIH grants
- Infrastructure to host challenges in the community
- Built on open-source CodaLab platform
- Integrates with ePAD and Cornerstone web viewers (radiology)
    - For ground truth generation
    - Results viewing
- Integrates with caMicrosope (digital pathology)
- System for assigning annotation tasks to experts for ground truth generation
- Been used for ~ two dozen challenges , over 2000 participants have signed up in the system
- Recent extensions to support code uploads through Docker

Past challenges

- MICCAI 2015, 2016, 2017
    - Brain tumor segmentation
    - Digital pathology (segmentation and classification)
    - Joint radiology/pathology challenge
    - Digital mammography
    - Liver metastases survival prediction
- QIN
    - BMMR (breast MR response prediction)
        - CT feature challenge
- QIBA lung nodule challenge
- SPIE ProstateX
- SPIE/AAPM/NCI ProstateX-2

Recent Challenges

Athinoula A.
Martinos Center
For Biomedical Imaging

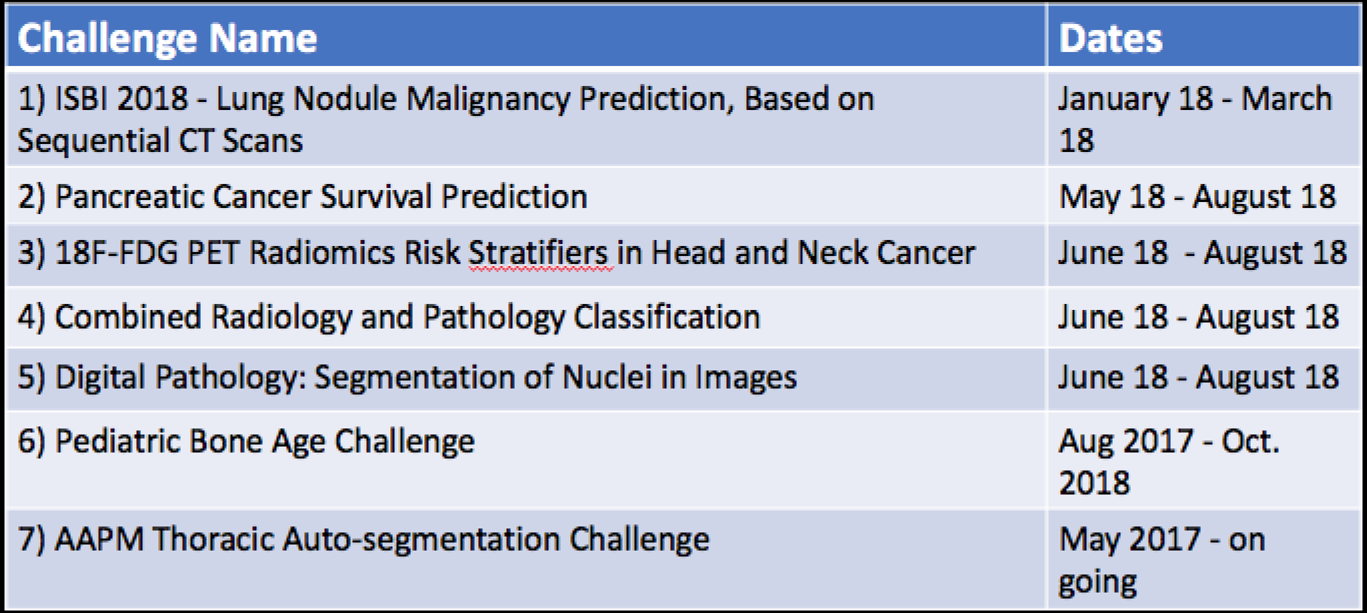

Challenge Name	Dates
1) ISBI 2018 - Lung Nodule Malignancy Prediction, Based on Sequential CT Scans	January 18 - March 18
2) Pancreatic Cancer Survival Prediction	May 18 - August 18
3) 18F-FDG PET Radiomics Risk Stratifiers in Head and Neck Cancer	June 18 - August 18
4) Combined Radiology and Pathology Classification	June 18 - August 18
5) Digital Pathology: Segmentation of Nuclei in Images	June 18 - August 18
6) Pediatric Bone Age Challenge	Aug 2017 - Oct. 2018
7) AAPM Thoracic Auto-segmentation Challenge	May 2017 - on going

Recent Challenges

Athinoula A.
Martinos Center
For Biomedical Imaging

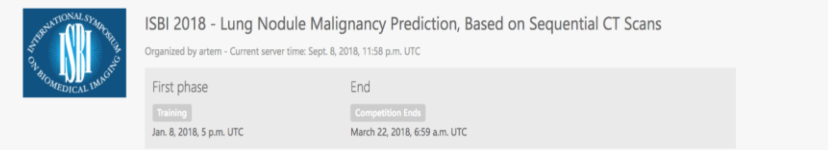

ISBI 2018 - Lung Nodule Malignancy Prediction, Based on Sequential CT Scans

Organized by artem - Current server time: Sept. 8, 2018, 11:58 p.m. UTC

First phase	End
Training Jan. 8, 2018, 5 p.m. UTC	Competition Ends March 22, 2018, 6:59 a.m. UTC

Combined Radiology and Pathology Classification

Organized by cpm.organizing.committee - Current server time: Sept. 9, 2018, 12:18 a.m. UTC

First phase	End
Training June 29, 2018, 11:59 p.m. UTC	Competition Ends Aug. 17, 2018, 11:59 p.m. UTC

Digital Pathology: Segmentation of Nuclei in Images

Organized by cpm.organizing.committee - Current server time: Sept. 9, 2018, 12:23 a.m. UTC

First phase	End
Training June 15, 2018, 11:59 p.m. UTC	Competition Ends Aug. 17, 2018, 11:59 p.m. UTC

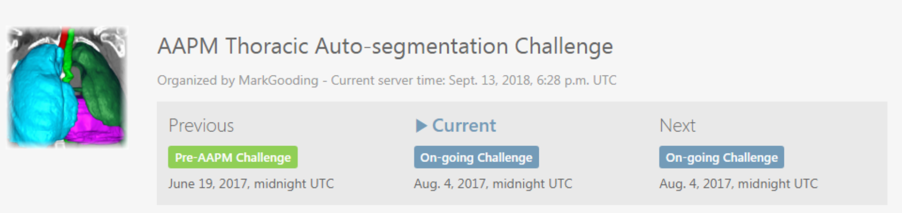

AAPM Thoracic Auto-segmentation Challenge

Organized by MarkGooding - Current server time: Sept. 13, 2018, 6:28 p.m. UTC

Previous	▶ Current	Next
Pre-AAPM Challenge June 19, 2017, midnight UTC	On-going Challenge Aug. 4, 2017, midnight UTC	On-going Challenge Aug. 4, 2017, midnight UTC

Challenge at RSNA 2017

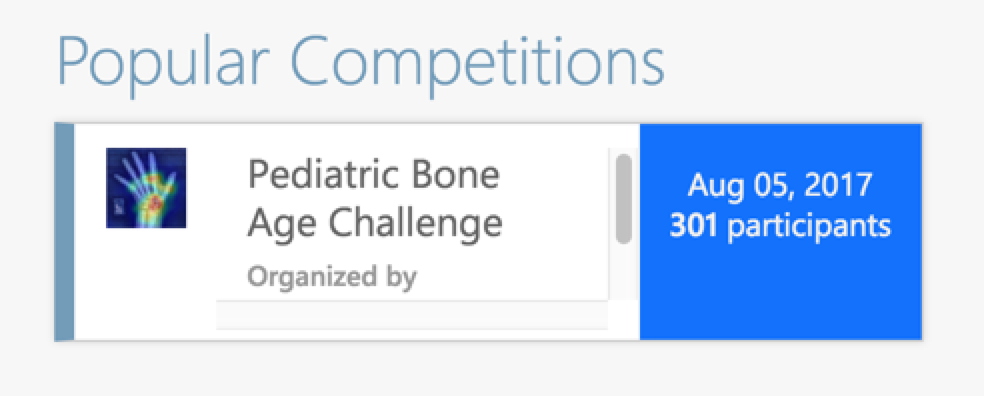

Popular Competitions

Pediatric Bone Age Challenge

Organized by

Aug 05, 2017
301 participants

Over 400 participants, still open
Official results presented at RNSA
Best results improved on recently published results from the authors who released dataset (Stanford)
All but one of top 10 used deep learning
Publication in review
Dataset still being used in research/publications

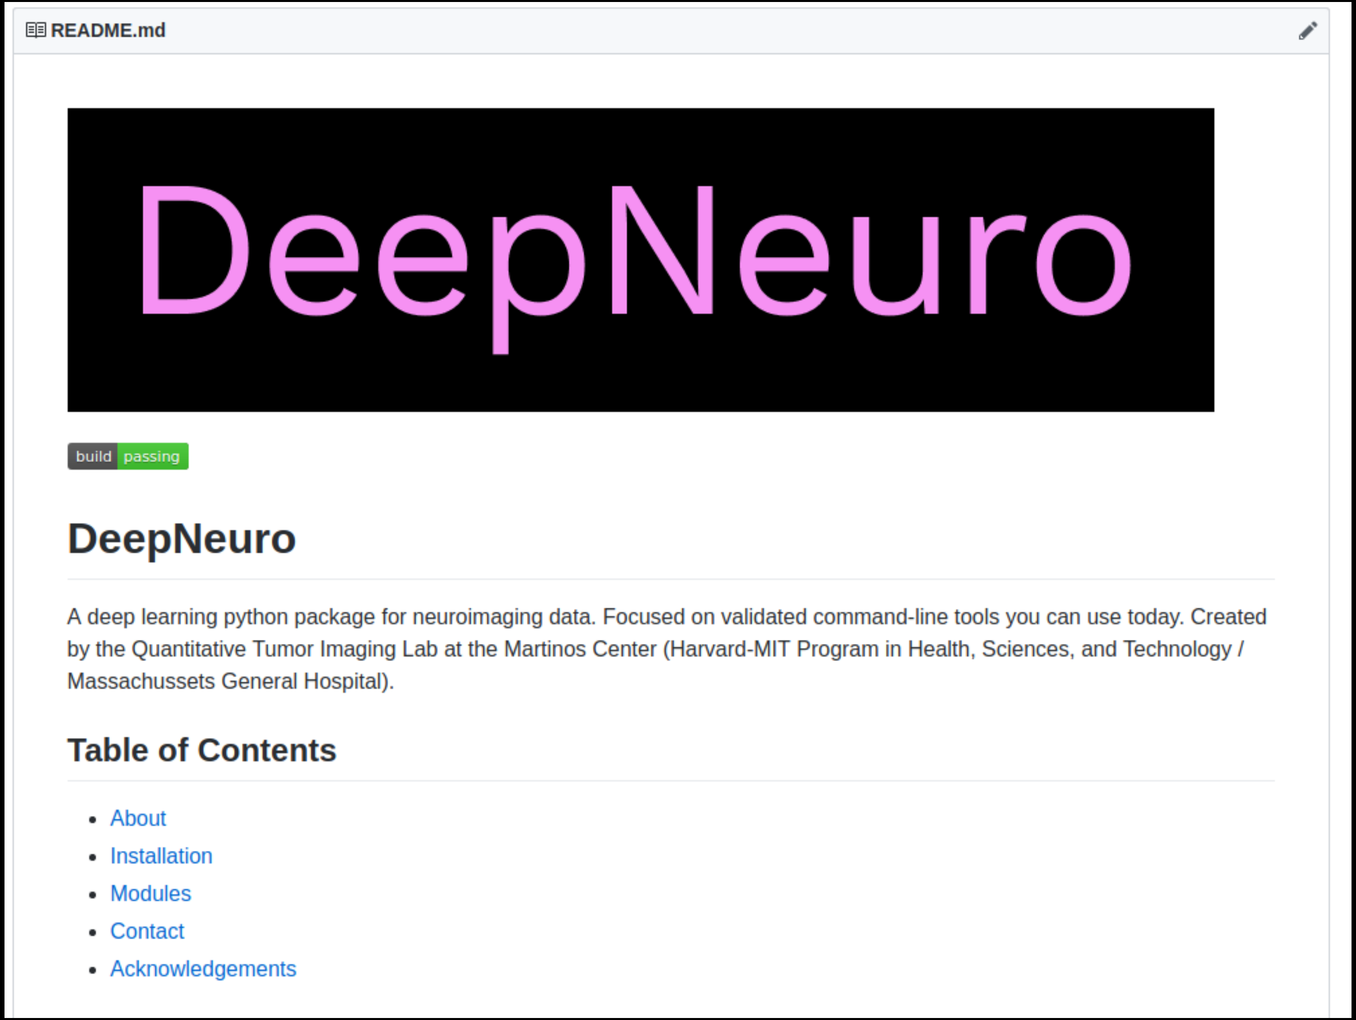

DeepNeuro

A deep learning python package for neuroimaging data. Focused on validated command-line tools you can use today. Created by the Quantitative Tumor Imaging Lab at the Martinos Center (Harvard-MIT Program in Health, Sciences, and Technology / Massachussets General Hospital).

Table of Contents

Deep Learning Software

What are the challenges of developing software for deep learning and neuroimaging?

- Complex hardware requirements (GPUs, Linux, libraries, etc.)
- Fast-moving field. Popular packages (Tensorflow, Theano, Keras, etc.) change monthly.
- Low interoperability with standard neuroimaging tools.
- Difficulty sharing and validating trained models.

DeepNeuro

Deep Learning Software

DeepNeuro

How to we get around these challenges?

- *Docker Containers*. Reduce the hassle of extensive and complex system requirements.

- *Open Source Code*. Deep learning algorithms can be sensitive to over-fitting, and validation and retraining on new datasets will be essential.

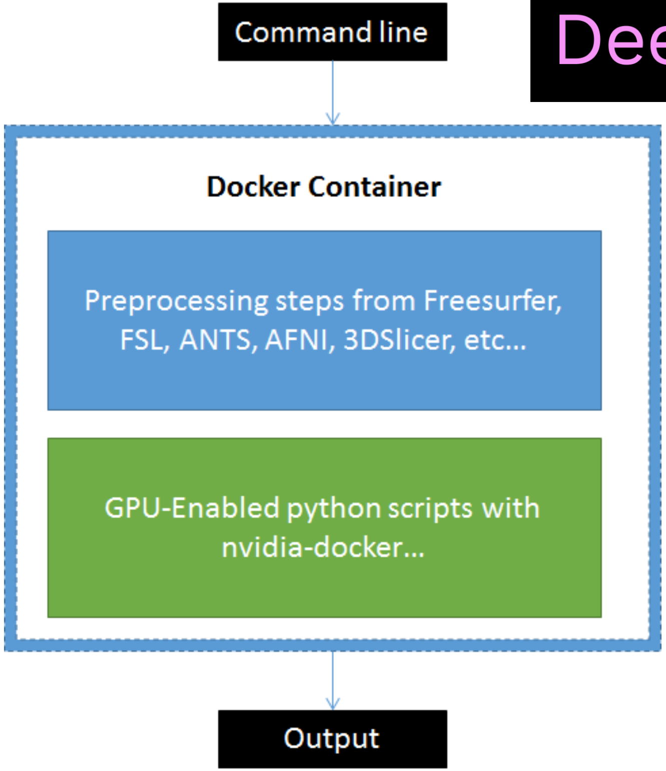

Command line

Docker Container

Preprocessing steps from Freesurfer, FSL, ANTS, AFNI, 3DSlicer, etc...

GPU-Enabled python scripts with nvidia-docker...

Output

Easy Configuration!

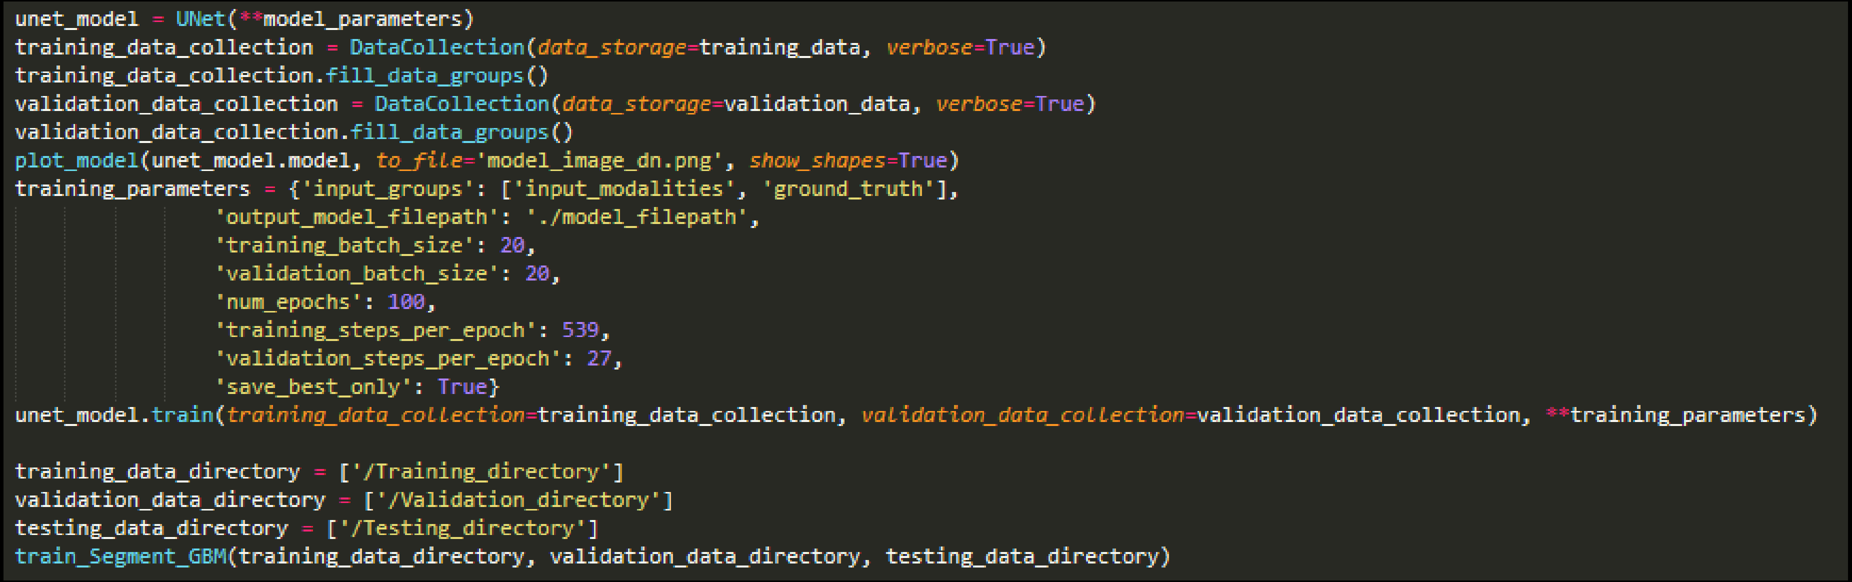

```
unet_model = UNet(**model_parameters)
training_data_collection = DataCollection(data_storage=training_data, verbose=True)
training_data_collection.fill_data_groups()
validation_data_collection = DataCollection(data_storage=validation_data, verbose=True)
validation_data_collection.fill_data_groups()
plot_model(unet_model.model, to_file='model_image_dn.png', show_shapes=True)
training_parameters = {'input_groups': ['input_modalities', 'ground_truth'],
                'output_model_filepath': './model_filepath',
                'training_batch_size': 20,
                'validation_batch_size': 20,
                'num_epochs': 100,
                'training_steps_per_epoch': 539,
                'validation_steps_per_epoch': 27,
                'save_best_only': True}
unet_model.train(training_data_collection=training_data_collection, validation_data_collection=validation_data_collection, **training_parameters)

training_data_directory = ['/Training_directory']
validation_data_directory = ['/Validation_directory']
testing_data_directory = ['/Testing_directory']
train_Segment_GBM(training_data_directory, validation_data_directory, testing_data_directory)
```

Just specify:

- Parameters of your neural network architecture
- Parameters for training
- Directory of your data

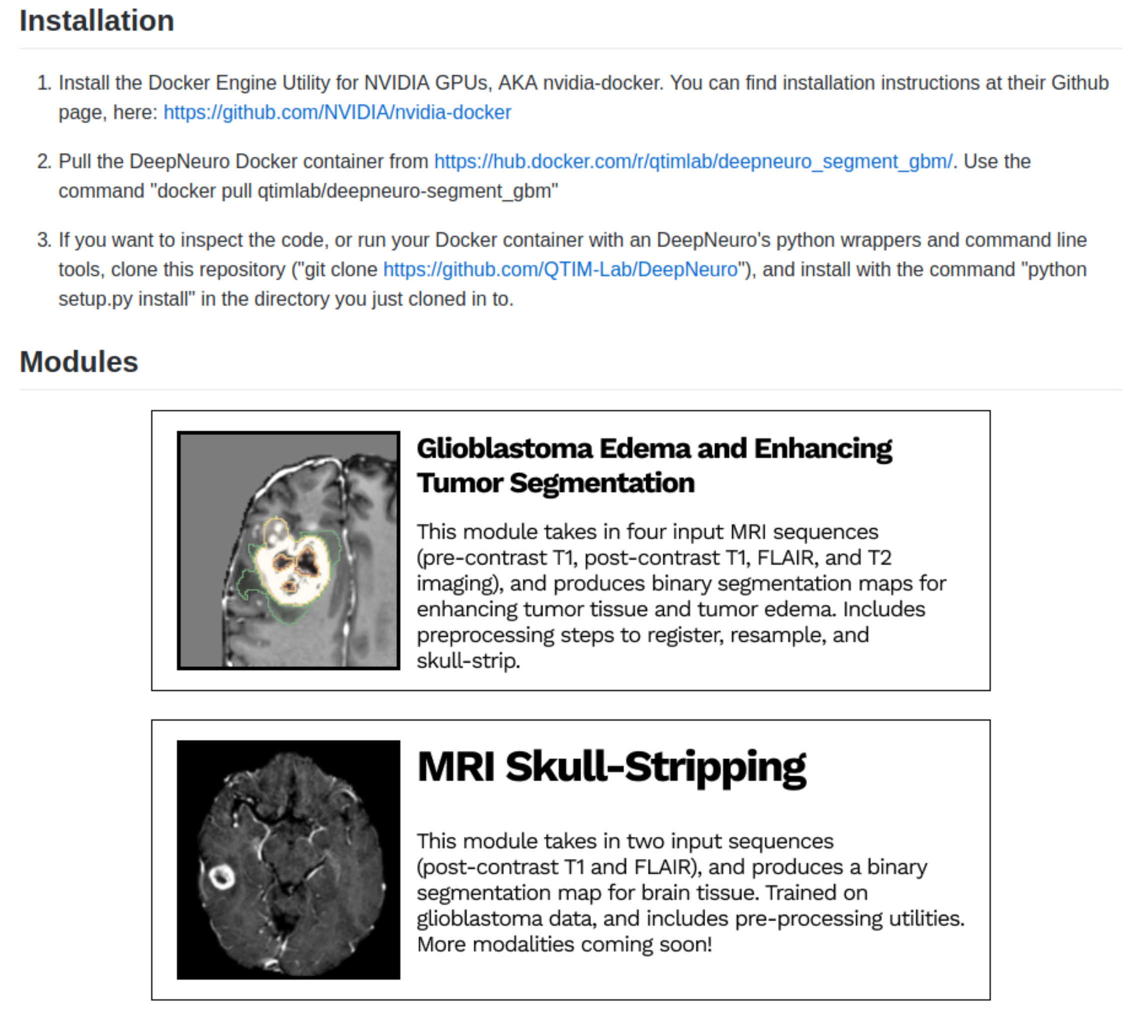

Installation

1. Install the Docker Engine Utility for NVIDIA GPUs, AKA nvidia-docker. You can find installation instructions at their Github page, here: https://github.com/NVIDIA/nvidia-docker
2. Pull the DeepNeuro Docker container from https://hub.docker.com/r/qtimlab/deepneuro_segment_gbm/. Use the command "docker pull qtimlab/deepneuro-segment_gbm"
3. If you want to inspect the code, or run your Docker container with an DeepNeuro's python wrappers and command line tools, clone this repository ("git clone https://github.com/QTIM-Lab/DeepNeuro"), and install with the command "python setup.py install" in the directory you just cloned in to.

Modules

Glioblastoma Edema and Enhancing Tumor Segmentation

This module takes in four input MRI sequences (pre-contrast T1, post-contrast T1, FLAIR, and T2 imaging), and produces binary segmentation maps for enhancing tumor tissue and tumor edema. Includes preprocessing steps to register, resample, and skull-strip.

MRI Skull-Stripping

This module takes in two input sequences (post-contrast T1 and FLAIR), and produces a binary segmentation map for brain tissue. Trained on glioblastoma data, and includes pre-processing utilities. More modalities coming soon!

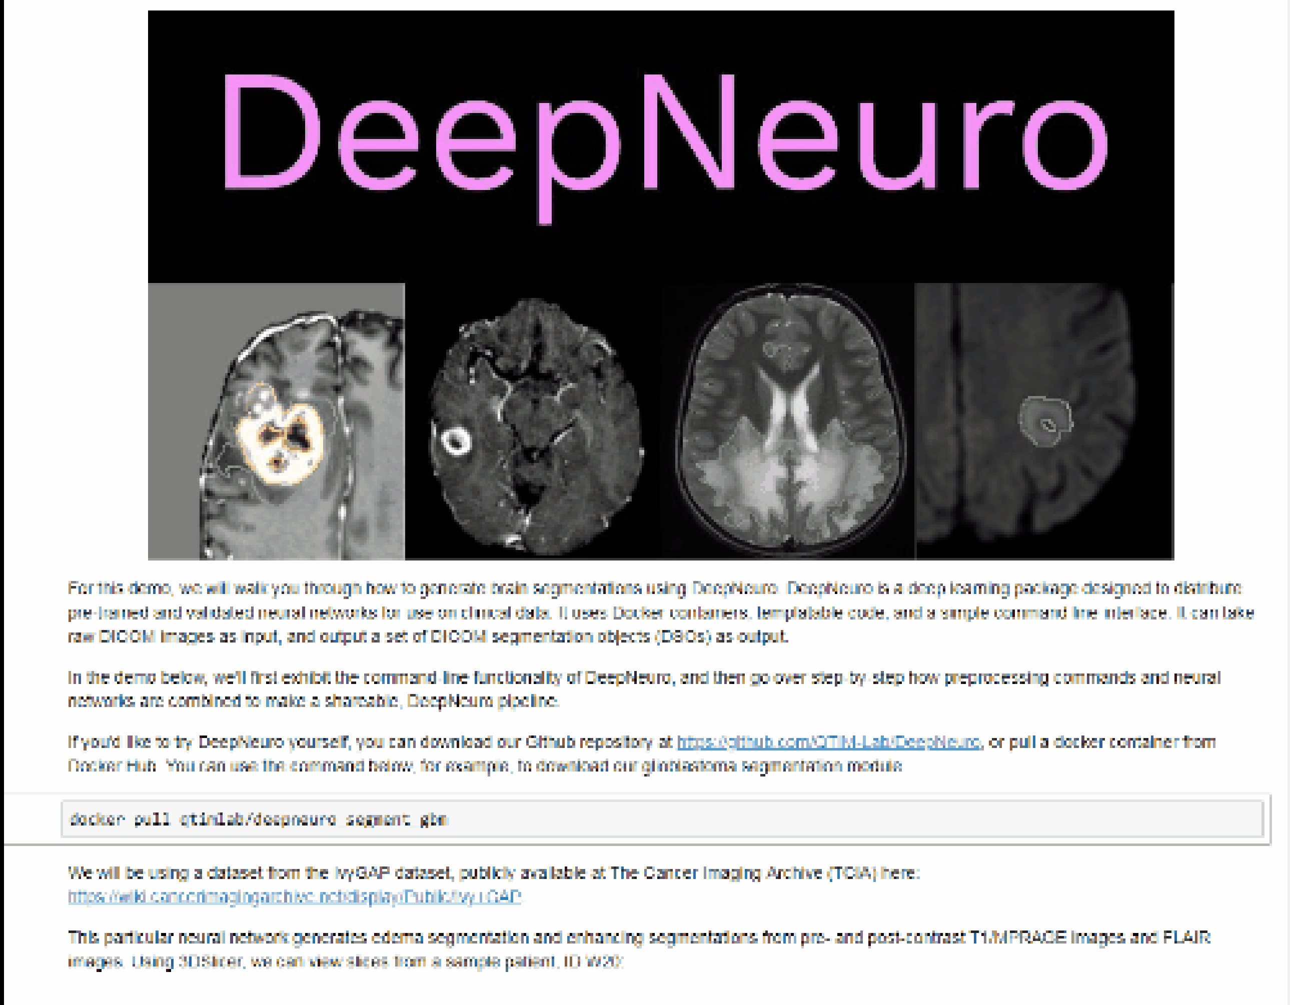

For this demo, we will walk you through how to generate brain segmentations using DeepNeuro. DeepNeuro is a deep learning package designed to distribute pre-trained and validated neural networks for use on clinical data. It uses Docker containers, templatable code, and a simple command line interface. It can take raw DICOM images as input, and output a set of DICOM segmentation objects (DSOs) as output.

In the demo below, we'll first exhibit the command-line functionality of DeepNeuro, and then go over step-by-step how preprocessing commands and neural networks are combined to make a sharable, DeepNeuro pipeline.

If you'd like to try DeepNeuro yourself, you can download our Github repository at https://github.com/QTIM-Lab/DeepNeuro, or pull a docker container from Docker Hub. You can use the command below, for example, to download our glioblastoma segmentation module.

```
docker pull qtimlab/deepneuro_segment_gbm
```

We will be using a dataset from the IvyGAP dataset, publicly available at The Cancer Imaging Archive (TCIA) here:
https://wiki.cancerimagingarchive.net/display/Public/Ivy+GAP

This particular neural network generates edema segmentation and enhancing segmentations from pre- and post-contrast T1/MPRAGE images and FLAIR images. Using 3DSlicer, we can view slices from a sample patient, ID W20:

DeepNeuro

Athinoula A.
Martinos
Center
For Biomedical Imaging

Tumor segmentation

- Manual segmentation of tumor used to take hours of clinician and technician time as they segmented MRI volumes slice-by-slice.
- Inter-operator variability can be high between oncologists, especially in low-resolution scans.

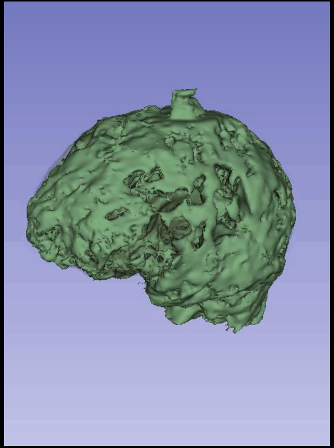

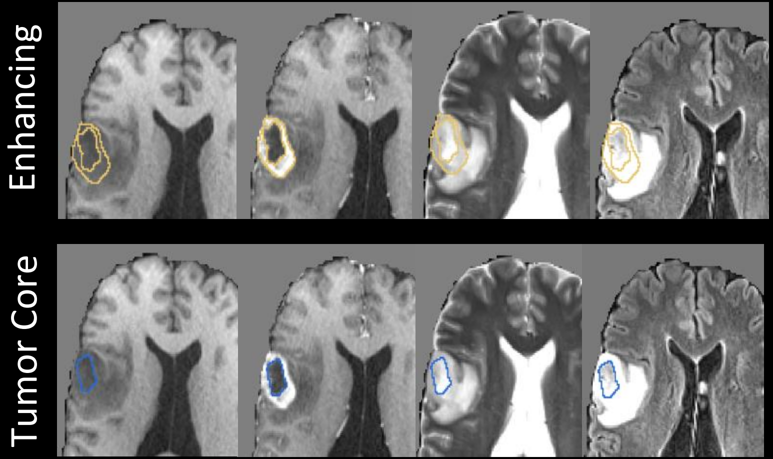

Beers et al., arXiv, 2018

DeepNeuro

Athinoula A.
Martinos Center
For Biomedical Imaging

IDH prediction

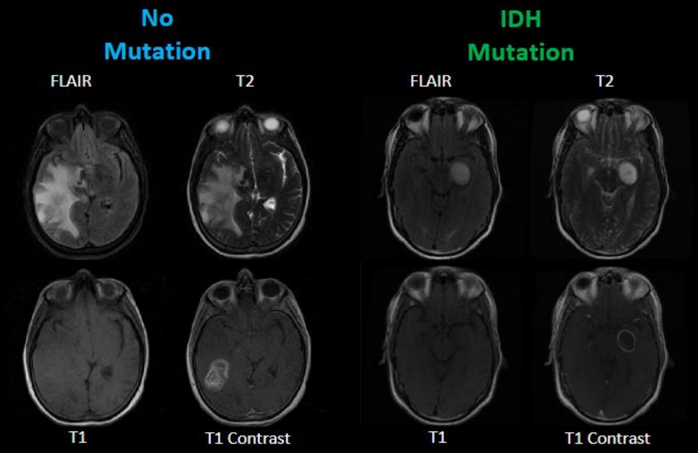

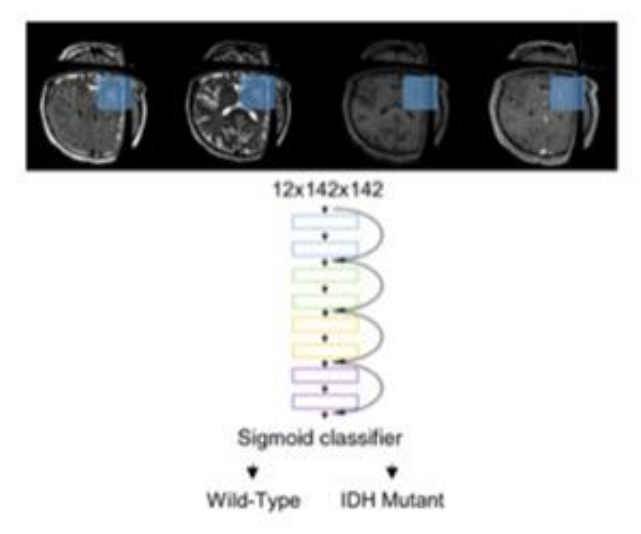

	Accuracy
Training	93%
Validation	94%
Testing	88%

Chang et al., Clinical Cancer Research 2018

DeepNeuro

Athinoula A.
Martinos Center
For Biomedical Imaging

Stroke

- Major cause of long-term disability + death
- Cerebral ischemia causes 80% of cases
- Diffusion-weighted MR Imaging (DWI) allows for early-stage diagnosis
- Automatic detection and quantification may aid clinical-decision making

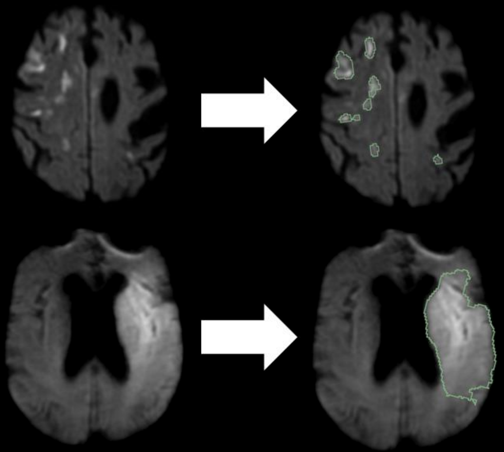

Disability & Survival Prediction

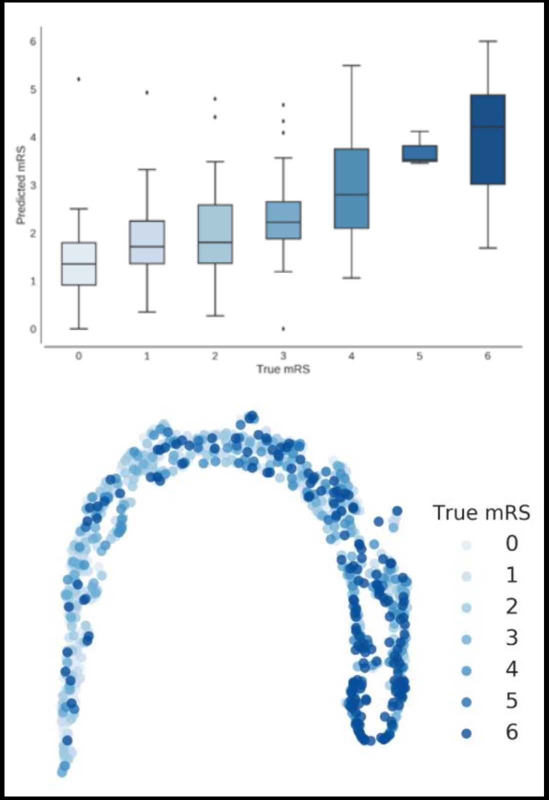

DeepNeuro

Vessel Segmentation

Athinoula A.
Martinos
Center
For Biomedical Imaging

Magnetic Resonance Angiography

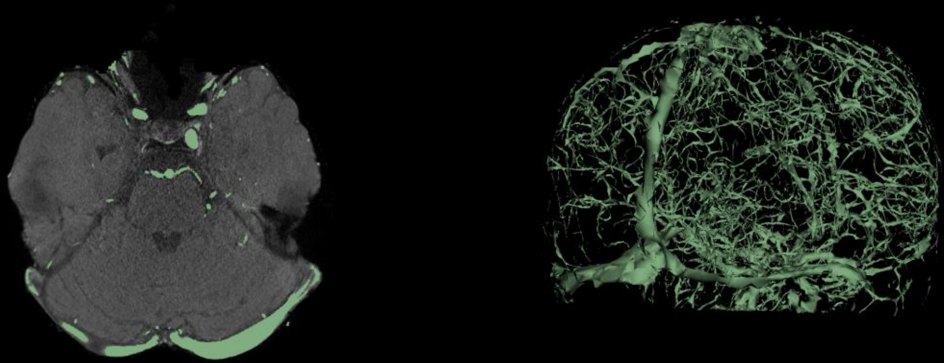

CLARITY

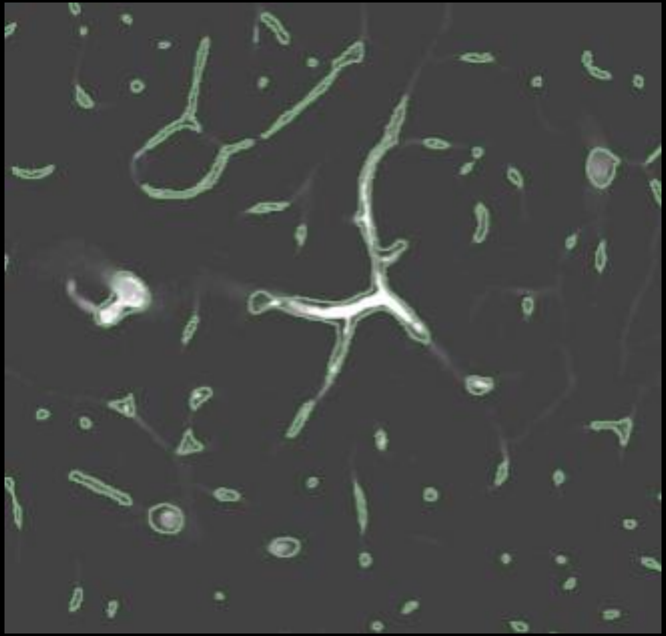

Athinoula A.
Martinos
Center
For Biomedical Imaging

Challenges in deep learning

Most deep learning methods need lots of data

Patient images are hard to share

Annotations are difficult to acquire

Annotations are noisy/biased

Deep learning models (typically) don’t provide uncertainty

Models are considered to be “black boxes”

Models can be “fooled”

Opportunities

Patient data is hard to share

- Distributed learning where data remains within institutions but models are shared

Distributed Deep Learning Techniques for Collaborative Multi-Institutional Studies

Currently, researchers host all patient data to a central location…

but this is often times not possible due to:

- Limited data storage
- Patient privacy concerns
- Institutions prefer not to share patient data because it is valuable

MGH 1811

OHSU

BWH

Neural Network Model

Distributed learning

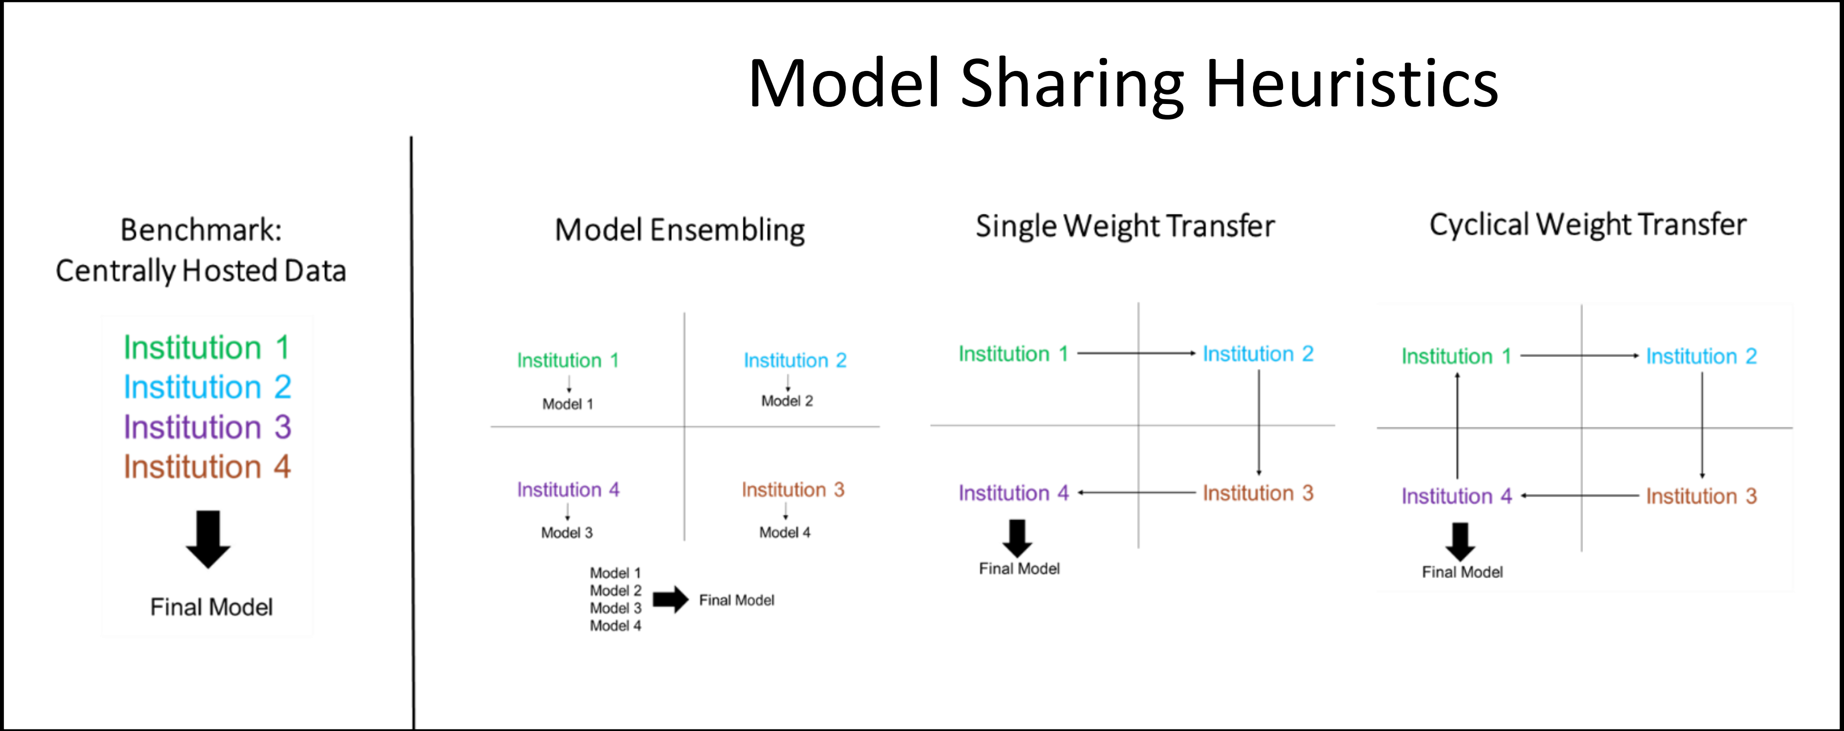

Model Sharing Heuristics

Benchmark:
Centrally Hosted Data

Institution 1
Institution 2
Institution 3
Institution 4

Final Model

Model Ensembling

Institution 1 → Model 1
Institution 2 → Model 2
Institution 4 → Model 3
Institution 3 → Model 4

Model 1
Model 2
Model 3
Model 4
→ Final Model

Single Weight Transfer

Institution 1 → Institution 2 → Institution 3 → Institution 4 → Final Model

Cyclical Weight Transfer

Institution 1 → Institution 2 → Institution 3 → Institution 4 → Institution 1

Institution 4 → Final Model

Chang et al., JAMIA 2018

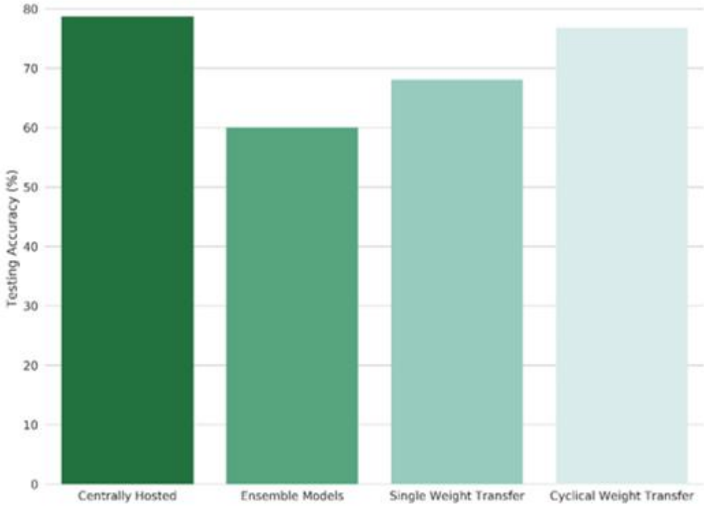

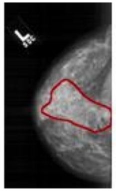

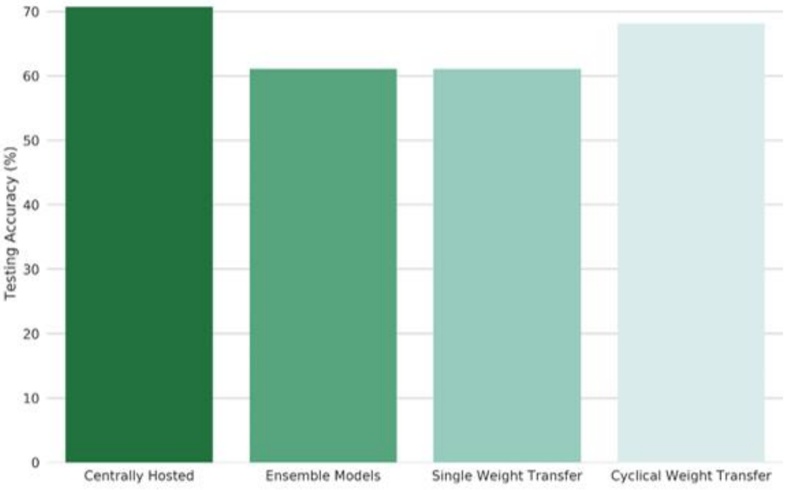

Chang et al., JAMIA 2018

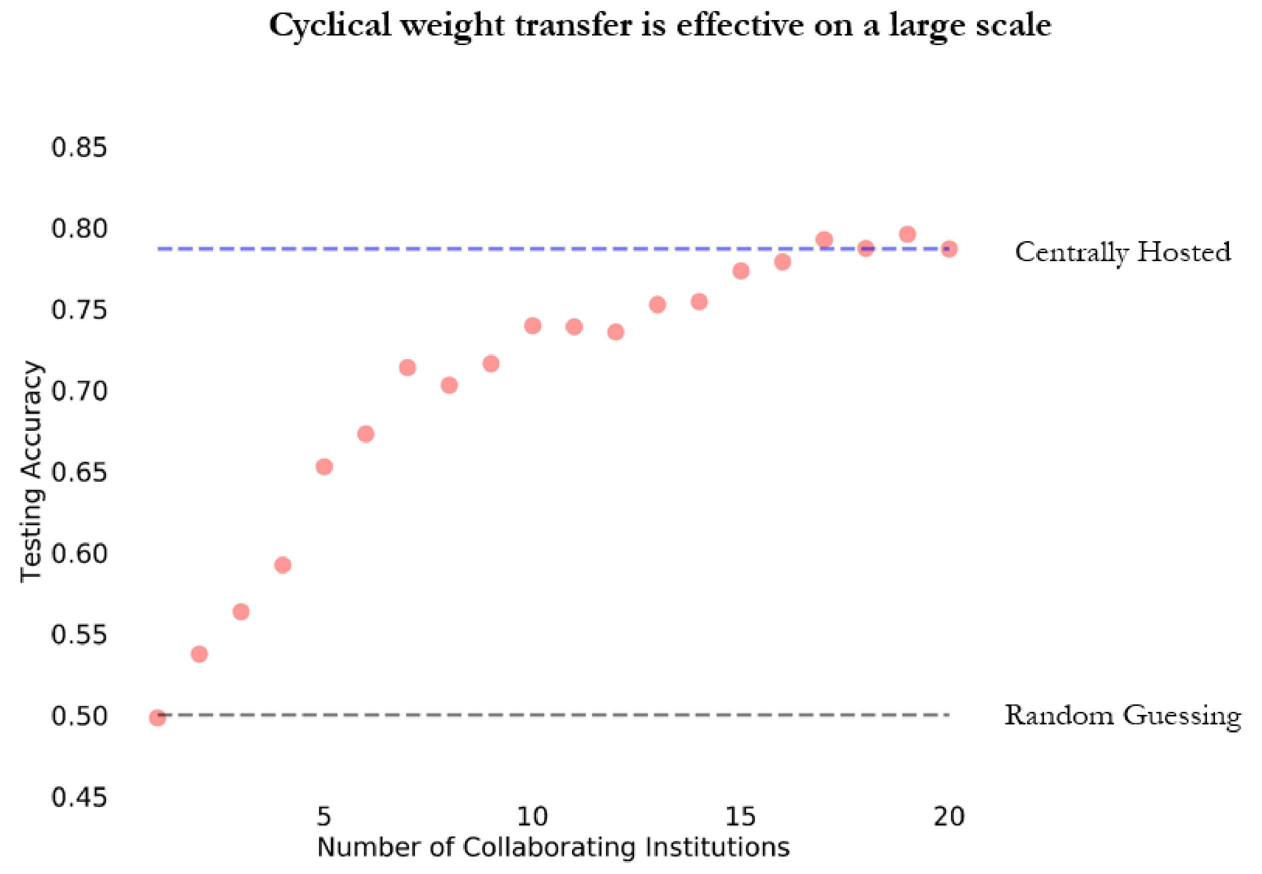

Cyclical weight transfer is effective on a large scale

Chang et al., JAMIA 2018

Opportunities

Annotations are difficult to acquire and can be very time-consuming

- Crowdsource

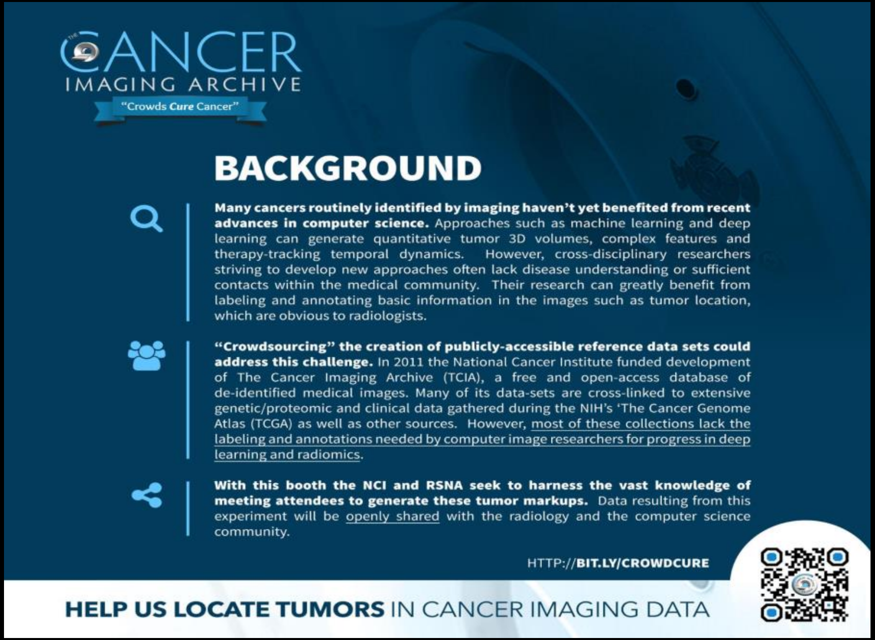

THE CANCER IMAGING ARCHIVE

"Crowds ***Cure*** Cancer"

BACKGROUND

Many cancers routinely identified by imaging haven't yet benefited from recent advances in computer science. Approaches such as machine learning and deep learning can generate quantitative tumor 3D volumes, complex features and therapy-tracking temporal dynamics. However, cross-disciplinary researchers striving to develop new approaches often lack disease understanding or sufficient contacts within the medical community. Their research can greatly benefit from labeling and annotating basic information in the images such as tumor location, which are obvious to radiologists.

"Crowdsourcing" the creation of publicly-accessible reference data sets could address this challenge. In 2011 the National Cancer Institute funded development of The Cancer Imaging Archive (TCIA), a free and open-access database of de-identified medical images. Many of its data-sets are cross-linked to extensive genetic/proteomic and clinical data gathered during the NIH's 'The Cancer Genome Atlas (TCGA) as well as other sources. However, most of these collections lack the labeling and annotations needed by computer image researchers for progress in deep learning and radiomics.

With this booth the NCI and RSNA seek to harness the vast knowledge of meeting attendees to generate these tumor markups. Data resulting from this experiment will be openly shared with the radiology and the computer science community.

HTTP://**BIT.LY/CROWDCURE**

HELP US LOCATE TUMORS IN CANCER IMAGING DATA

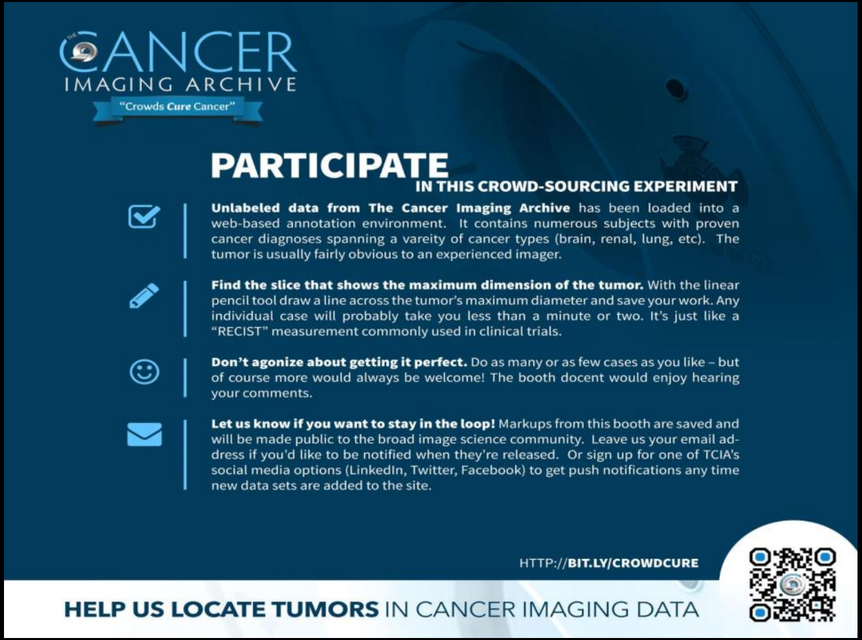

THE CANCER IMAGING ARCHIVE

"Crowds ***Cure*** Cancer"

PARTICIPATE

IN THIS CROWD-SOURCING EXPERIMENT

Unlabeled data from The Cancer Imaging Archive has been loaded into a web-based annotation environment. It contains numerous subjects with proven cancer diagnoses spanning a vareity of cancer types (brain, renal, lung, etc). The tumor is usually fairly obvious to an experienced imager.

Find the slice that shows the maximum dimension of the tumor. With the linear pencil tool draw a line across the tumor's maximum diameter and save your work. Any individual case will probably take you less than a minute or two. It's just like a "RECIST" measurement commonly used in clinical trials.

Don't agonize about getting it perfect. Do as many or as few cases as you like – but of course more would always be welcome! The booth docent would enjoy hearing your comments.

Let us know if you want to stay in the loop! Markups from this booth are saved and will be made public to the broad image science community. Leave us your email address if you'd like to be notified when they're released. Or sign up for one of TCIA's social media options (LinkedIn, Twitter, Facebook) to get push notifications any time new data sets are added to the site.

HTTP://**BIT.LY/CROWDCURE**

HELP US LOCATE TUMORS IN CANCER IMAGING DATA

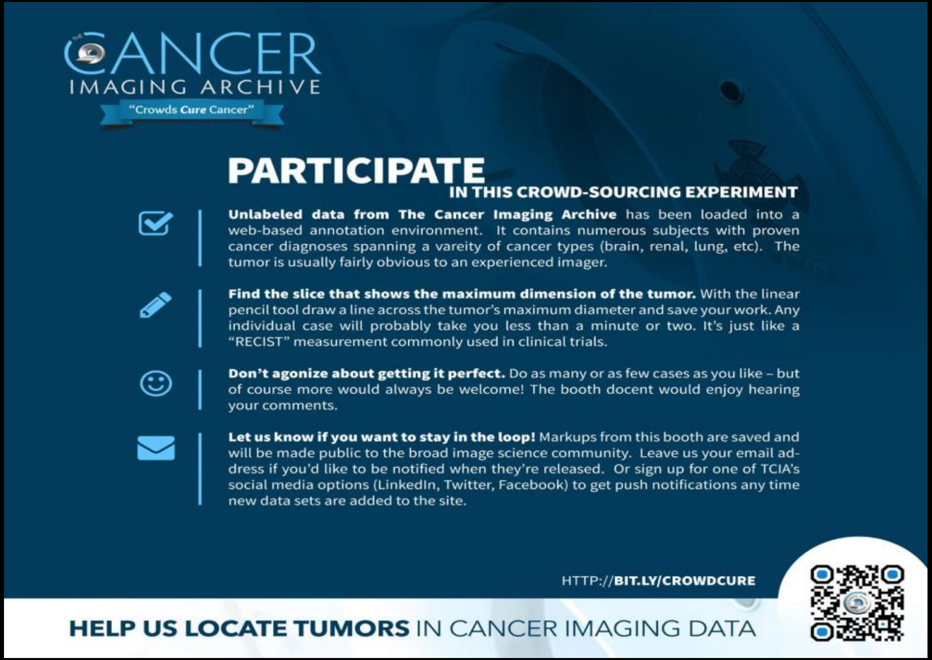

THE CANCER
IMAGING ARCHIVE
"Crowds *Cure* Cancer"

PARTICIPATE
IN THIS CROWD-SOURCING EXPERIMENT

☑ **Unlabeled data from The Cancer Imaging Archive** has been loaded into a web-based annotation environment. It contains numerous subjects with proven cancer diagnoses spanning a vareity of cancer types (brain, renal, lung, etc). The tumor is usually fairly obvious to an experienced imager.

✎ **Find the slice that shows the maximum dimension of the tumor.** With the linear pencil tool draw a line across the tumor's maximum diameter and save your work. Any individual case will probably take you less than a minute or two. It's just like a "RECIST" measurement commonly used in clinical trials.

☺ **Don't agonize about getting it perfect.** Do as many or as few cases as you like – but of course more would always be welcome! The booth docent would enjoy hearing your comments.

✉ **Let us know if you want to stay in the loop!** Markups from this booth are saved and will be made public to the broad image science community. Leave us your email address if you'd like to be notified when they're released. Or sign up for one of TCIA's social media options (LinkedIn, Twitter, Facebook) to get push notifications any time new data sets are added to the site.

HTTP://**BIT.LY/CROWDCURE**

HELP US LOCATE TUMORS IN CANCER IMAGING DATA

Annotators

- Over 250 signed up
  - 211 had at least 1 annotation
  - 112 were radiologists
- Selected
  - Liver: 189
  - Lung: 211
  - Renal: 165
  - Ovarian: 133
- Selected
  - 4: 120
  - 3: 27
  - 2: 29
  - 1: 78

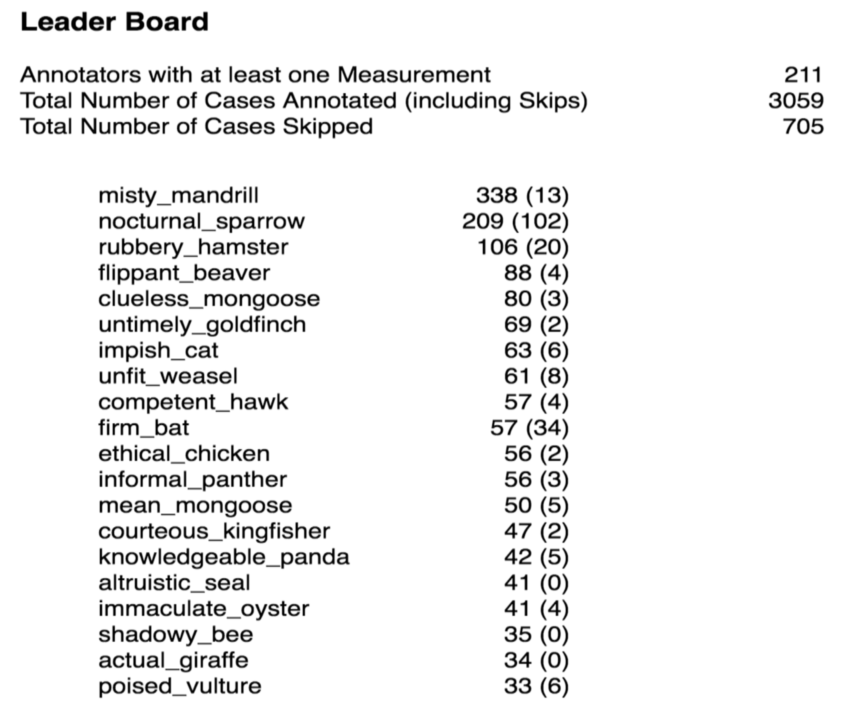

Leader Board

Annotators with at least one Measurement	211
Total Number of Cases Annotated (including Skips)	3059
Total Number of Cases Skipped	705

misty_mandrill	338 (13)
nocturnal_sparrow	209 (102)
rubbery_hamster	106 (20)
flippant_beaver	88 (4)
clueless_mongoose	80 (3)
untimely_goldfinch	69 (2)
impish_cat	63 (6)
unfit_weasel	61 (8)
competent_hawk	57 (4)
firm_bat	57 (34)
ethical_chicken	56 (2)
informal_panther	56 (3)
mean_mongoose	50 (5)
courteous_kingfisher	47 (2)
knowledgeable_panda	42 (5)
altruistic_seal	41 (0)
immaculate_oyster	41 (4)
shadowy_bee	35 (0)
actual_giraffe	34 (0)
poised_vulture	33 (6)

Results Disseminated

Athinoula A.
Martinos Center
For Biomedical Imaging

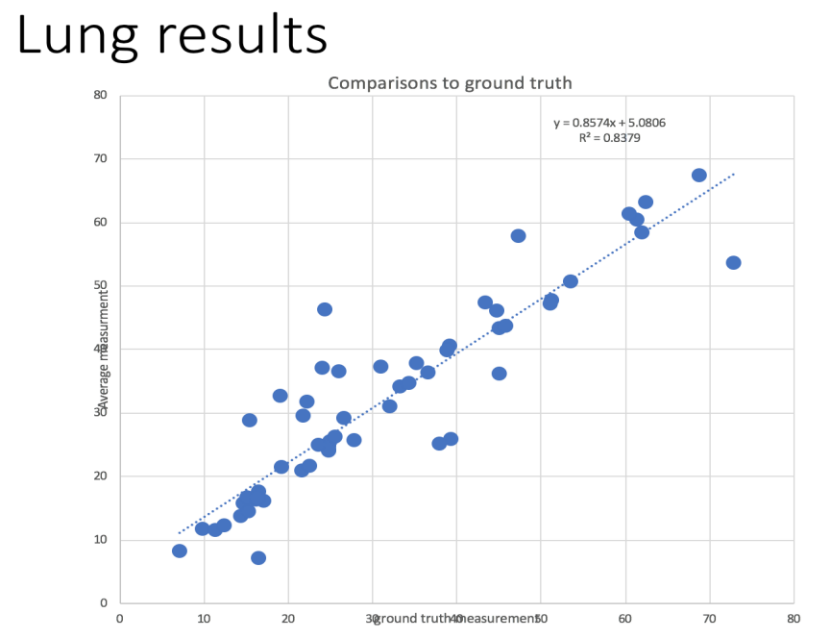

Lung results

Comparisons to ground truth

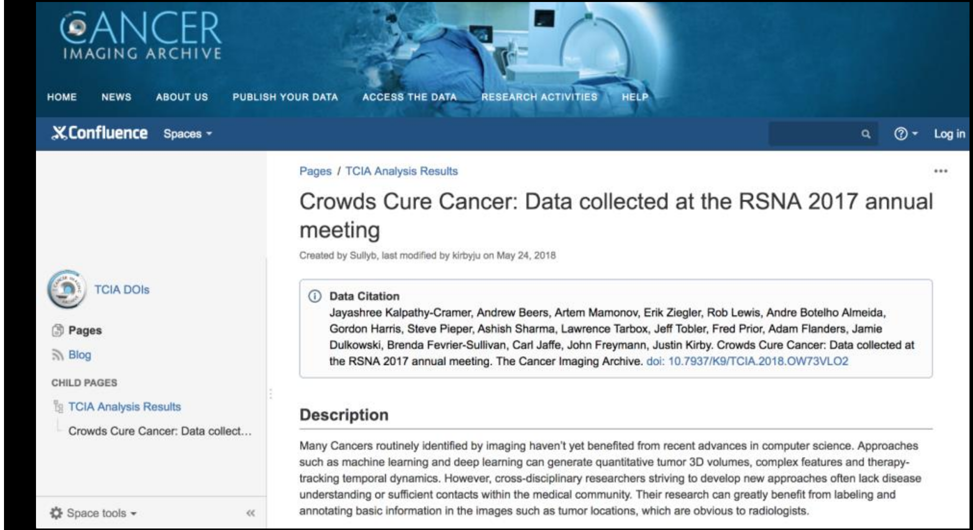

Pages / TCIA Analysis Results

Crowds Cure Cancer: Data collected at the RSNA 2017 annual meeting

Created by Sullyb, last modified by kirbyju on May 24, 2018

Data Citation
Jayashree Kalpathy-Cramer, Andrew Beers, Artem Mamonov, Erik Ziegler, Rob Lewis, Andre Botelho Almeida, Gordon Harris, Steve Pieper, Ashish Sharma, Lawrence Tarbox, Jeff Tobler, Fred Prior, Adam Flanders, Jamie Dulkowski, Brenda Fevrier-Sullivan, Carl Jaffe, John Freymann, Justin Kirby. Crowds Cure Cancer: Data collected at the RSNA 2017 annual meeting. The Cancer Imaging Archive. doi: 10.7937/K9/TCIA.2018.OW73VLO2

Description

Many Cancers routinely identified by imaging haven't yet benefited from recent advances in computer science. Approaches such as machine learning and deep learning can generate quantitative tumor 3D volumes, complex features and therapy-tracking temporal dynamics. However, cross-disciplinary researchers striving to develop new approaches often lack disease understanding or sufficient contacts within the medical community. Their research can greatly benefit from labeling and annotating basic information in the images such as tumor locations, which are obvious to radiologists.

Opportunities

Deep Learning models are black boxes

Model explainability

Tumor Segmentation Network

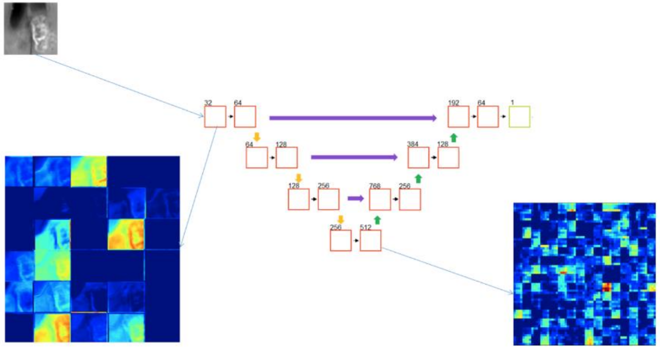

Dimensionality Reduction of Intermediate Layer

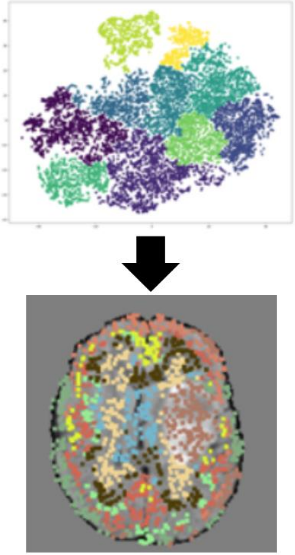

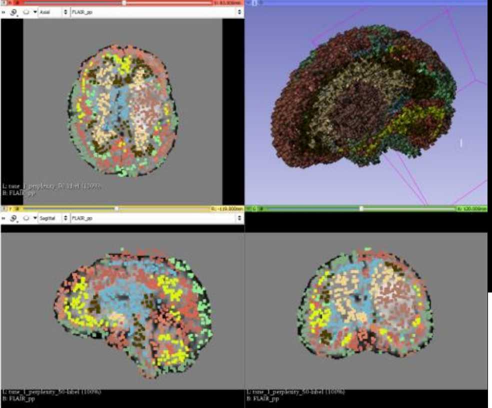

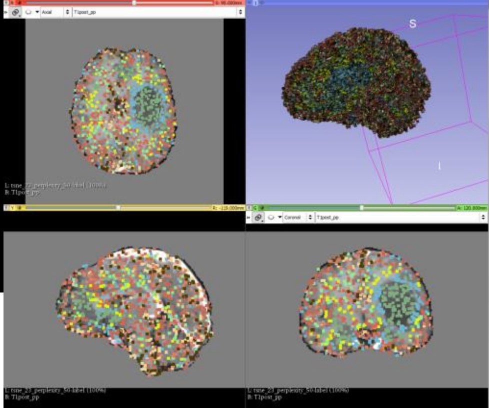

Neural Network Architecture: 3D U-Net

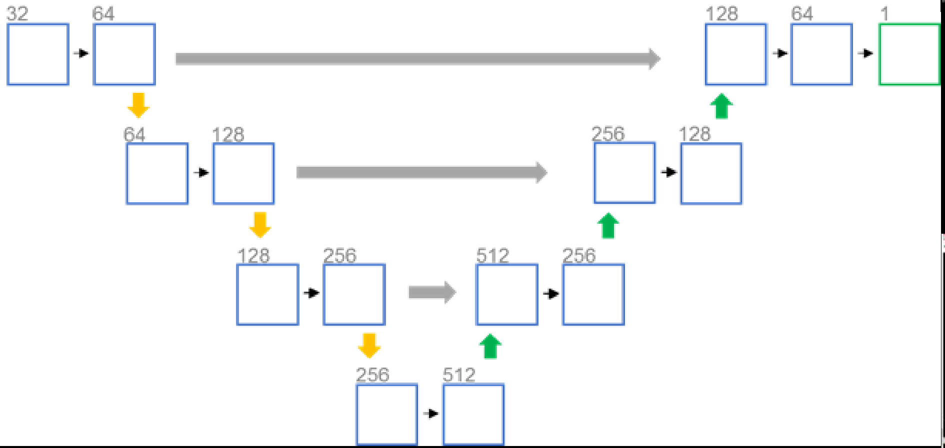

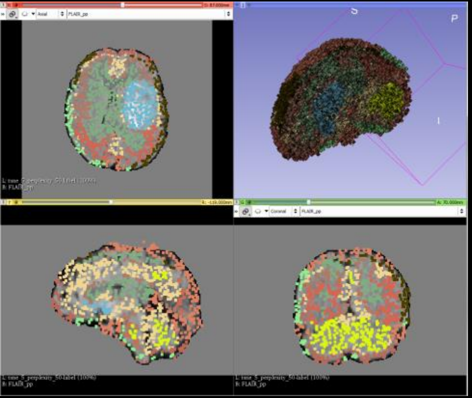

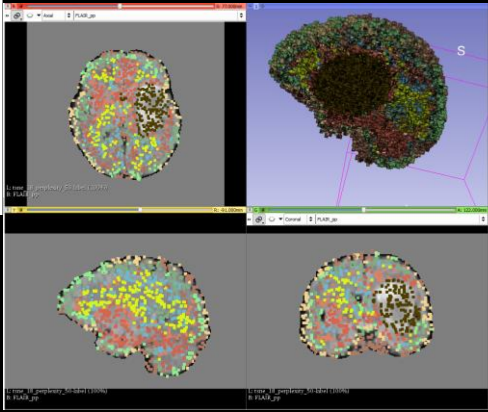

JAMA Ophthalmology | Original Investigation

Automated Diagnosis of Plus Disease in Retinopathy of Prematurity Using Deep Convolutional Neural Networks

James M. Brown, PhD; J. Peter Campbell, MD, MPH; Andrew Beers, BA; Ken Chang, MSE; Susan Ostmo, MS; R. V. Paul Chan, MD;
Jennifer Dy, PhD; Deniz Erdogmus, PhD; Stratis Ioannidis, PhD; Jayashree Kalpathy-Cramer, PhD; Michael F. Chiang, MD;
for the Imaging and Informatics in Retinopathy of Prematurity (i-ROP) Research Consortium

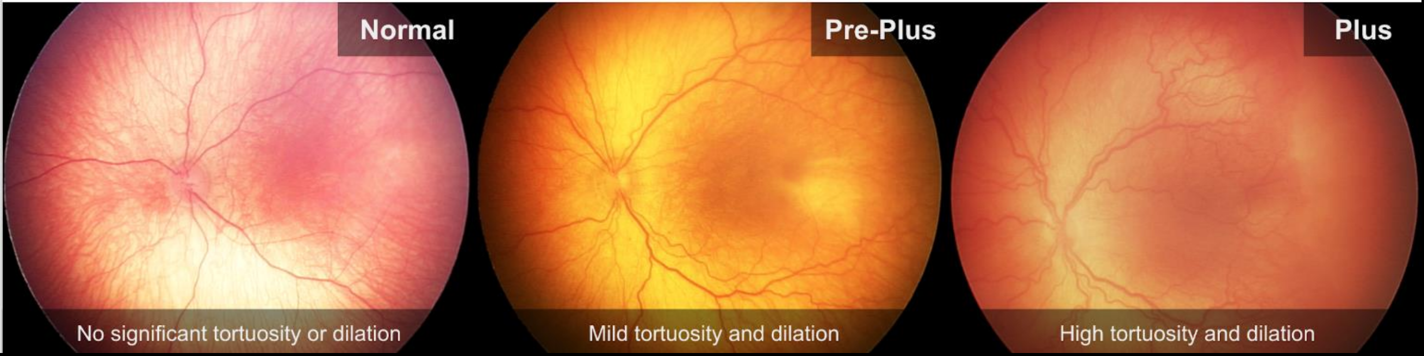

No significant tortuosity or dilation

Mild tortuosity and dilation

High tortuosity and dilation

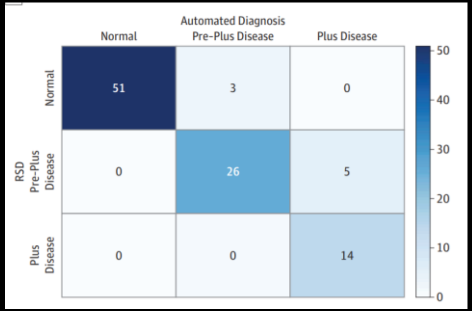

RSD	Normal	Pre-Plus Disease	Plus Disease
Normal	51	3	0
Pre-Plus Disease	0	26	5
Plus Disease	0	0	14

Pre-processing workflow

Masking, vessel segmentation², cropping

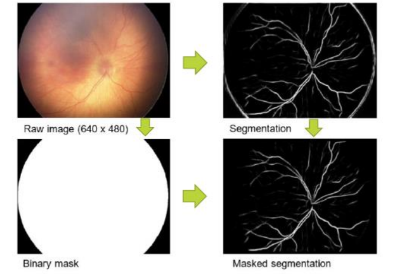

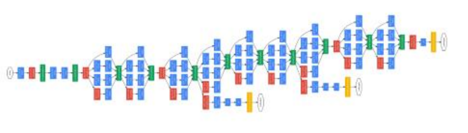

	Reader 1	Reader 3	Reader 2	Reader 8	Reader 7	Reader 6	i-ROP DL	Reader 5	Reader 4	RSD	Consensus
Reader 1	1.00	0.72	0.74	0.59	0.64	0.63	0.73	0.69	0.73	0.80	0.76
Reader 3	0.72	1.00	0.72	0.65	0.71	0.71	0.76	0.73	0.76	0.81	0.81
Reader 2	0.74	0.72	1.00	0.67	0.67	0.73	0.77	0.79	0.81	0.82	0.82
Reader 8	0.59	0.65	0.67	1.00	0.83	0.86	0.83	0.85	0.82	0.81	0.84
Reader 7	0.64	0.71	0.67	0.83	1.00	0.88	0.86	0.87	0.87	0.83	0.86
Reader 6	0.63	0.71	0.73	0.86	0.88	1.00	0.87	0.91	0.87	0.86	0.89
i-ROP DL	0.73	0.76	0.77	0.83	0.86	0.87	1.00	0.93	0.90	0.92	0.93
Reader 5	0.69	0.73	0.79	0.85	0.87	0.91	0.93	1.00	0.92	0.91	0.94
Reader 4	0.73	0.76	0.81	0.82	0.87	0.87	0.90	0.92	1.00	0.95	0.96
RSD	0.00	0.81	0.82	0.81	0.83	0.86	0.92	0.91	0.95	1.00	0.97
Consensus	0.76	0.81	0.82	0.84	0.86	0.89	0.93	0.94	0.96	0.97	1.00

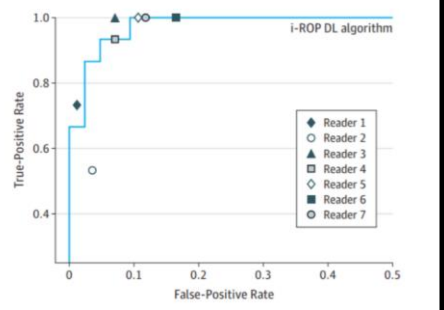

The Quantitative Tumor Imaging Lab at the Martinos Center (QTIM)

MGH & BWH CENTER FOR
CLINICAL DATA SCIENCE

Athinoula A.
Martinos
Center
For Biomedical Imaging

HST Harvard–MIT
Health Sciences & Technology

Athinoula A.
Martinos
Center
For Biomedical Imaging

Thank you!

QTIM

Andrew Beers

James Brown

Katharina Hoebel

Ken Chang

Samarth Nandekar

Ina Ly

Yi-Fen Yen

Elizabeth Gerstner

Martinos Center

Bo Zhu

Matt Rosen

Bruce Rosen

Hakan Ay

CCDS

Kathy Andriole

Mark Michalski

OHSU

Michael Chiang

J. Peter Campbell

Susan Ostmo

NEU

Deniz Erdogmus

Stratis Ioannidis

BWH

Steve Pieper

Andrey Fedorov

Stanford

Daniel Rubin

Niranjan Balachandar

I2R/A*STAR

Pavitra Krishnaswamy

Vijay Chandrasheka

Athinoula A.
Martinos
Center
For Biomedical Imaging

Acknowledgements

- U01 CA154601 Quantitative MRI Of Glioblastoma Response
- U24 CA180927 Informatics Tools For Optimized Imaging Biomarkers For Cancer Research & Discovery
- NSF 1622542 SCH: INT: Collaborative Research: Assistive Integrative Support Tool for Retinopathy of Prematurity
- R01 EY019474 Clinical And Genetıc Analysıs Of Retınopathy Of Prematurıty
- U24 CA180918 Quantitative Image Informatics For Cancer Research (QIICR)
- CCDS for computing resources