Deep learning approaches to medical applications

Joseph Paul Cohen

Postdoctoral Fellow at Montreal Institute for Learning Algorithms
Friend of the Farlow Fellow at Harvard University
U.S. National Science Foundation Graduate Fellow

Mila Medical

1. What is Deep Learning?
2. Image and Volume
   a. Segmentation
   b. Counting
   c. Interpreting
3. Genomics
   a. Representation

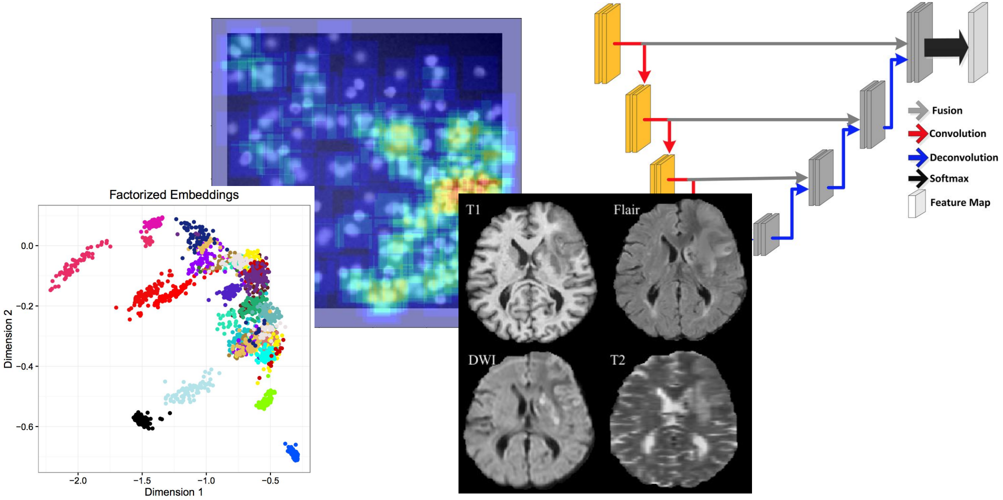

What is Deep Learning?

Logistic Regression

[200, 45, -4, 80]

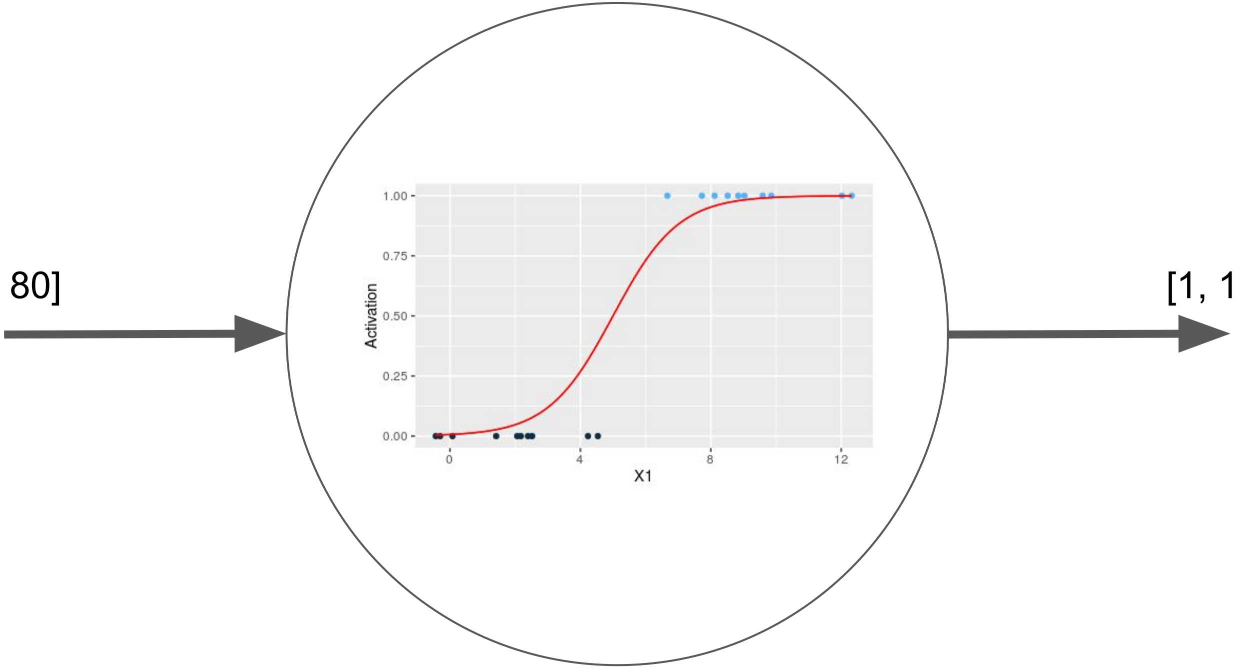

[1, 1, 0, 1]

What is Deep Learning?

Logistic Regression

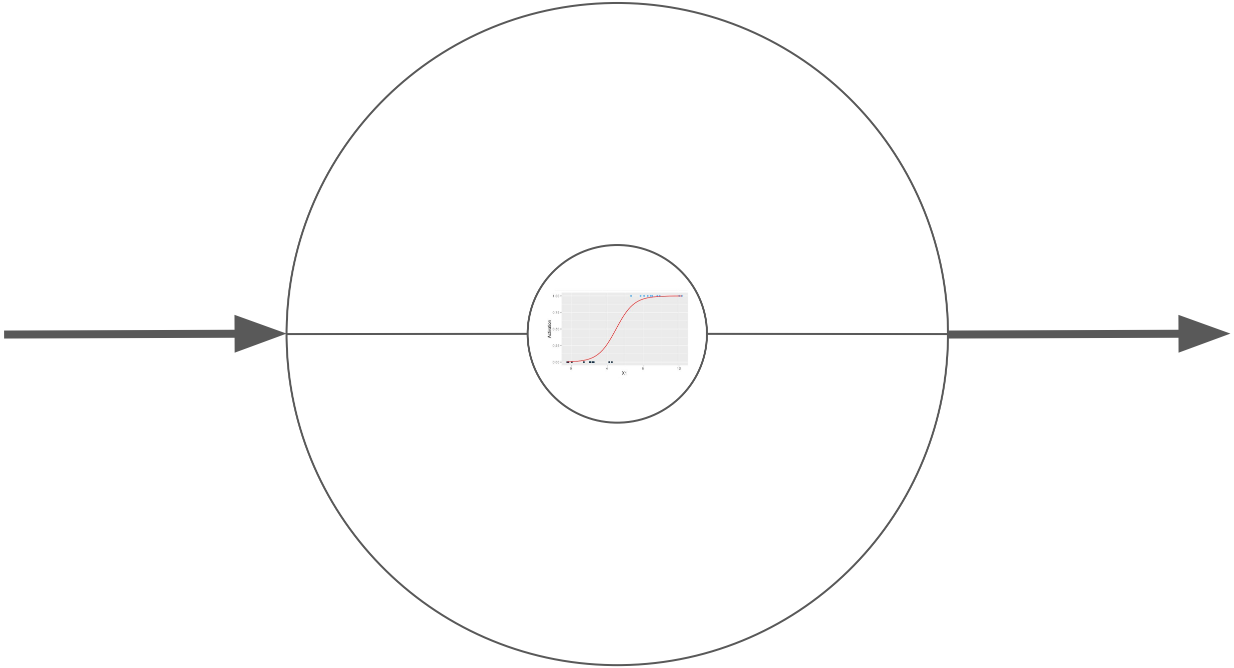

What is Deep Learning?

Logistic Regressions

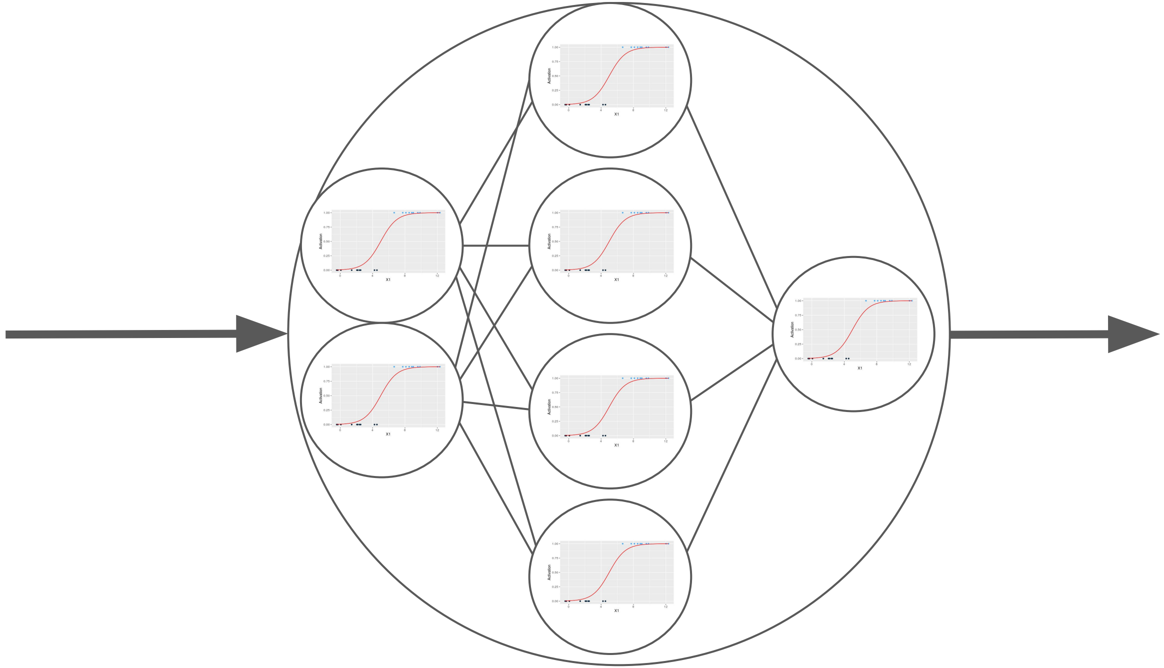

Image credit Pascal Vincent

What is Deep Learning?

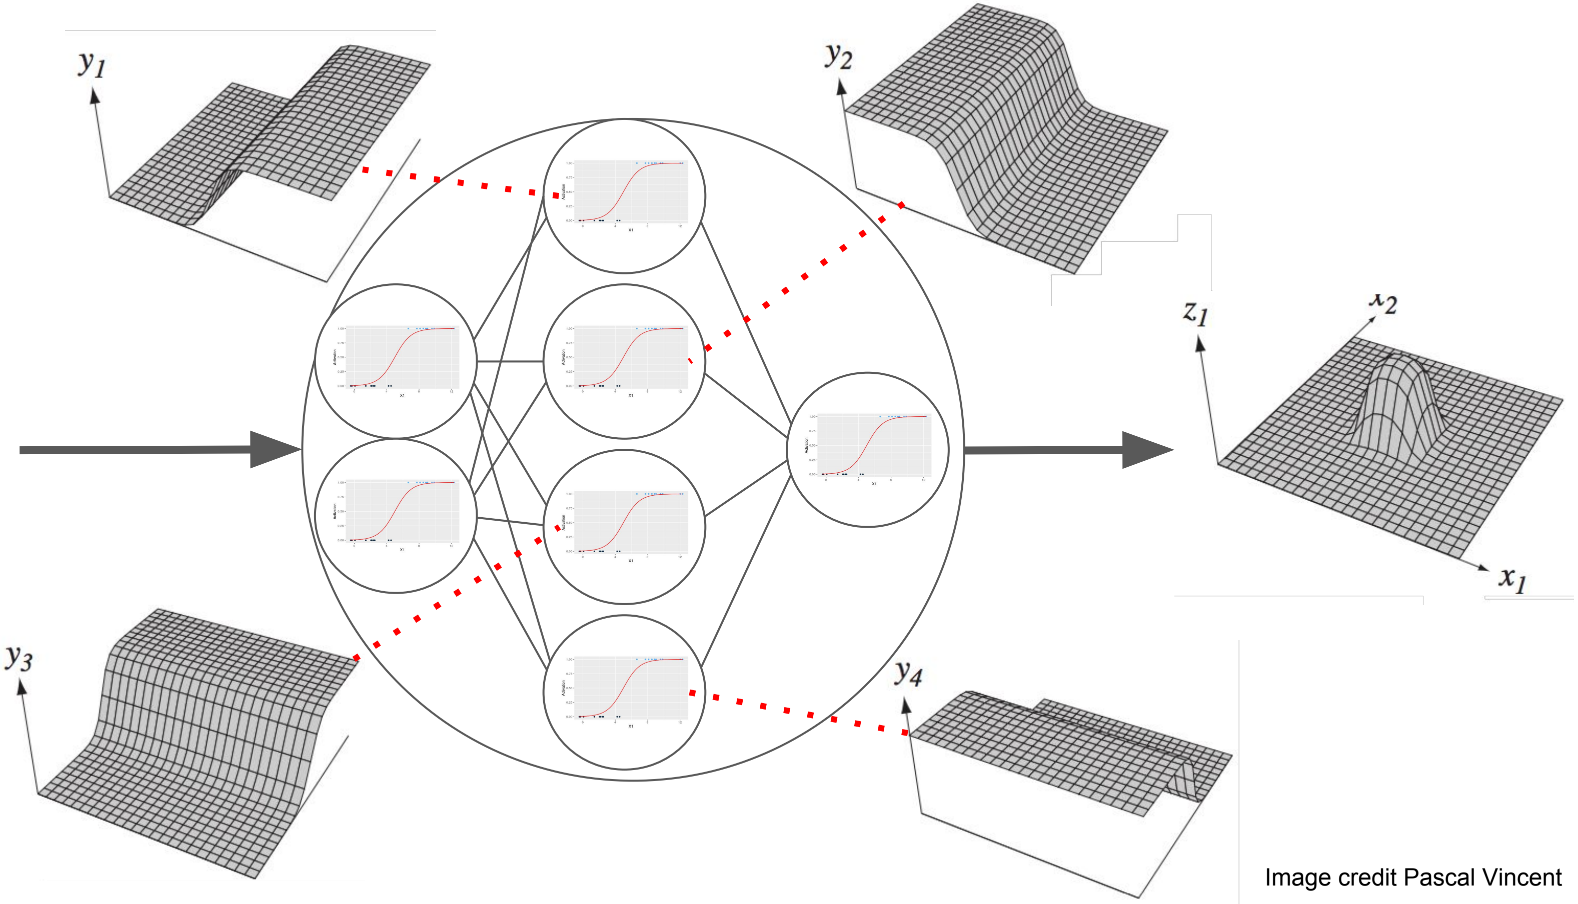

Image credit Pascal Vincent

What is Deep Learning?

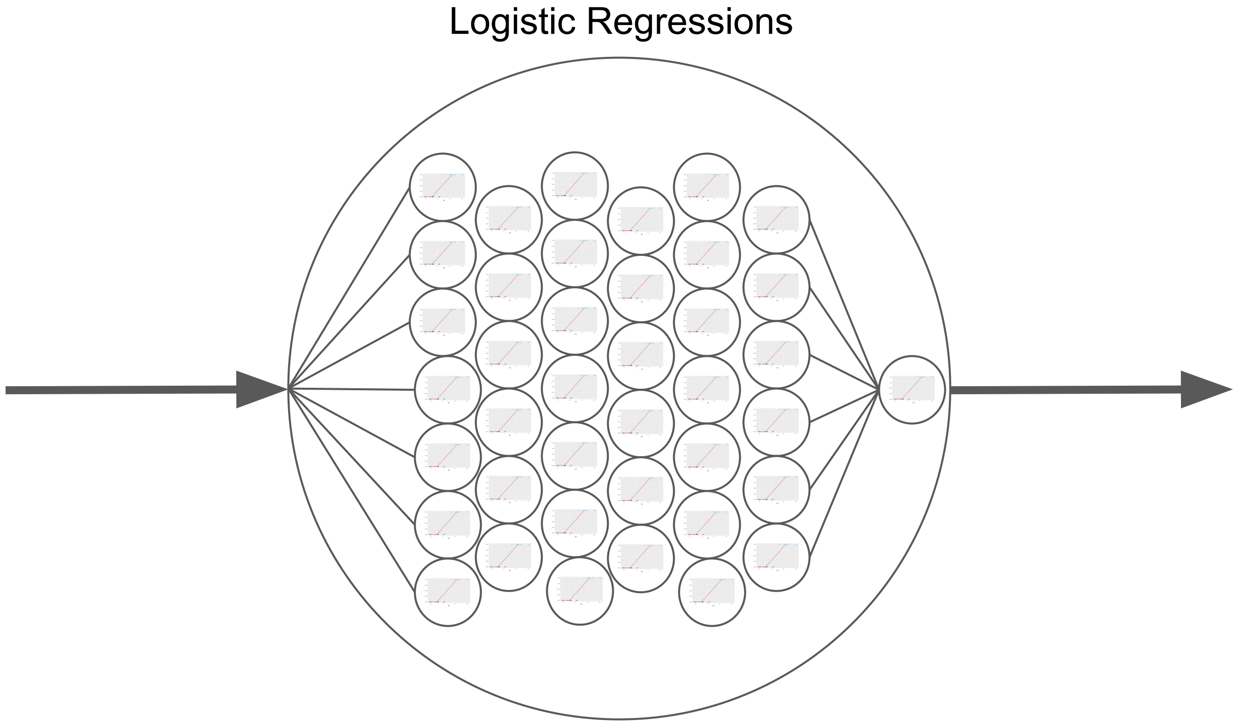

Automated pathology

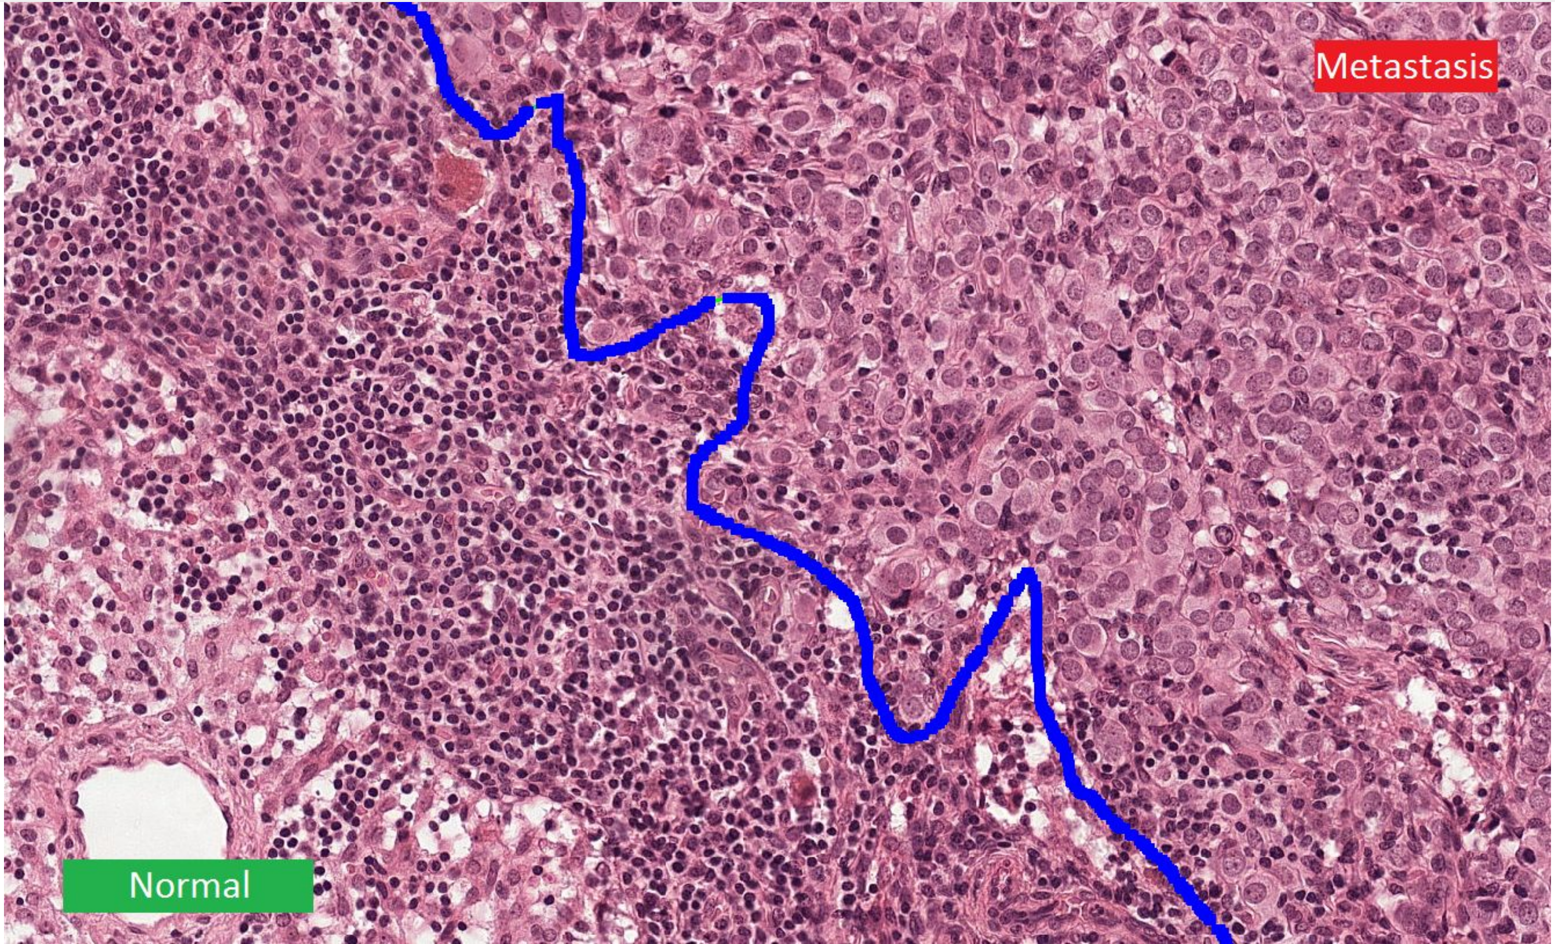

Tumor Segmentation

Automated detection of metatases in hematoxylin and eosin (H&E) stained whole-slide images of lymph node sections.

Reduce the workload of the pathologists and the subjectivity in diagnosis!

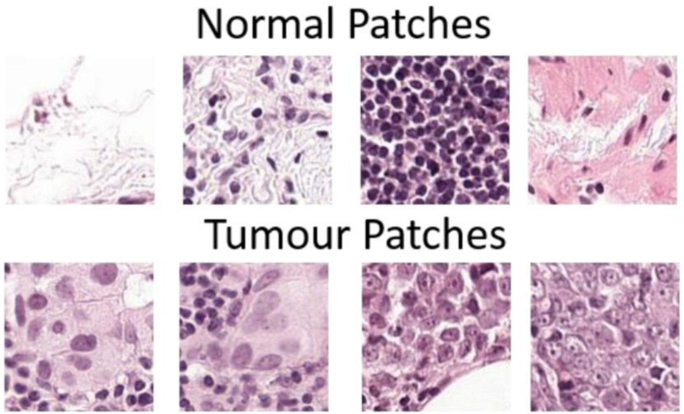

Determining the probability that an image patch is a tumour can be solved using deep learning models.

Image credit: Harvard Medical School, MIT, and EXB Research

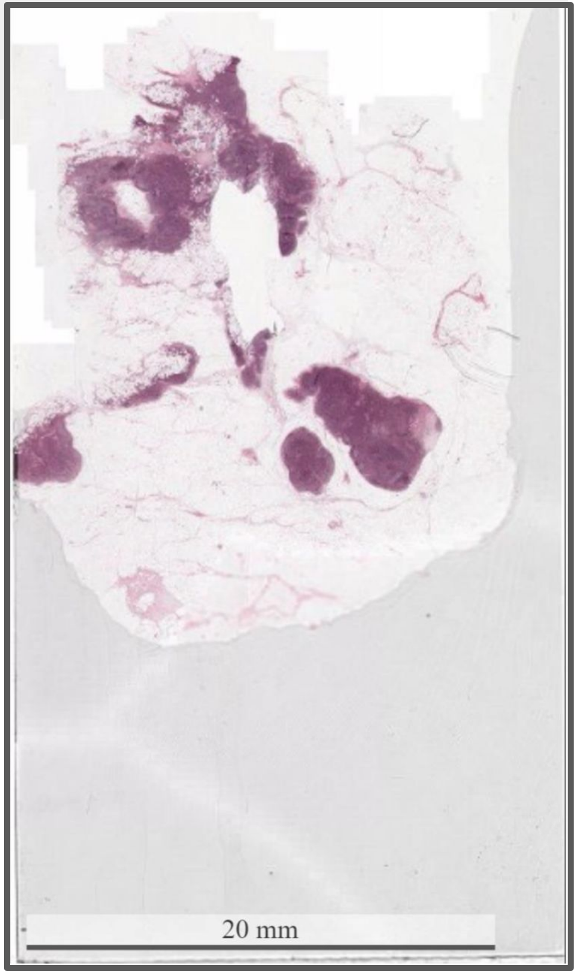

https://camelyon16.grand-challenge.org/

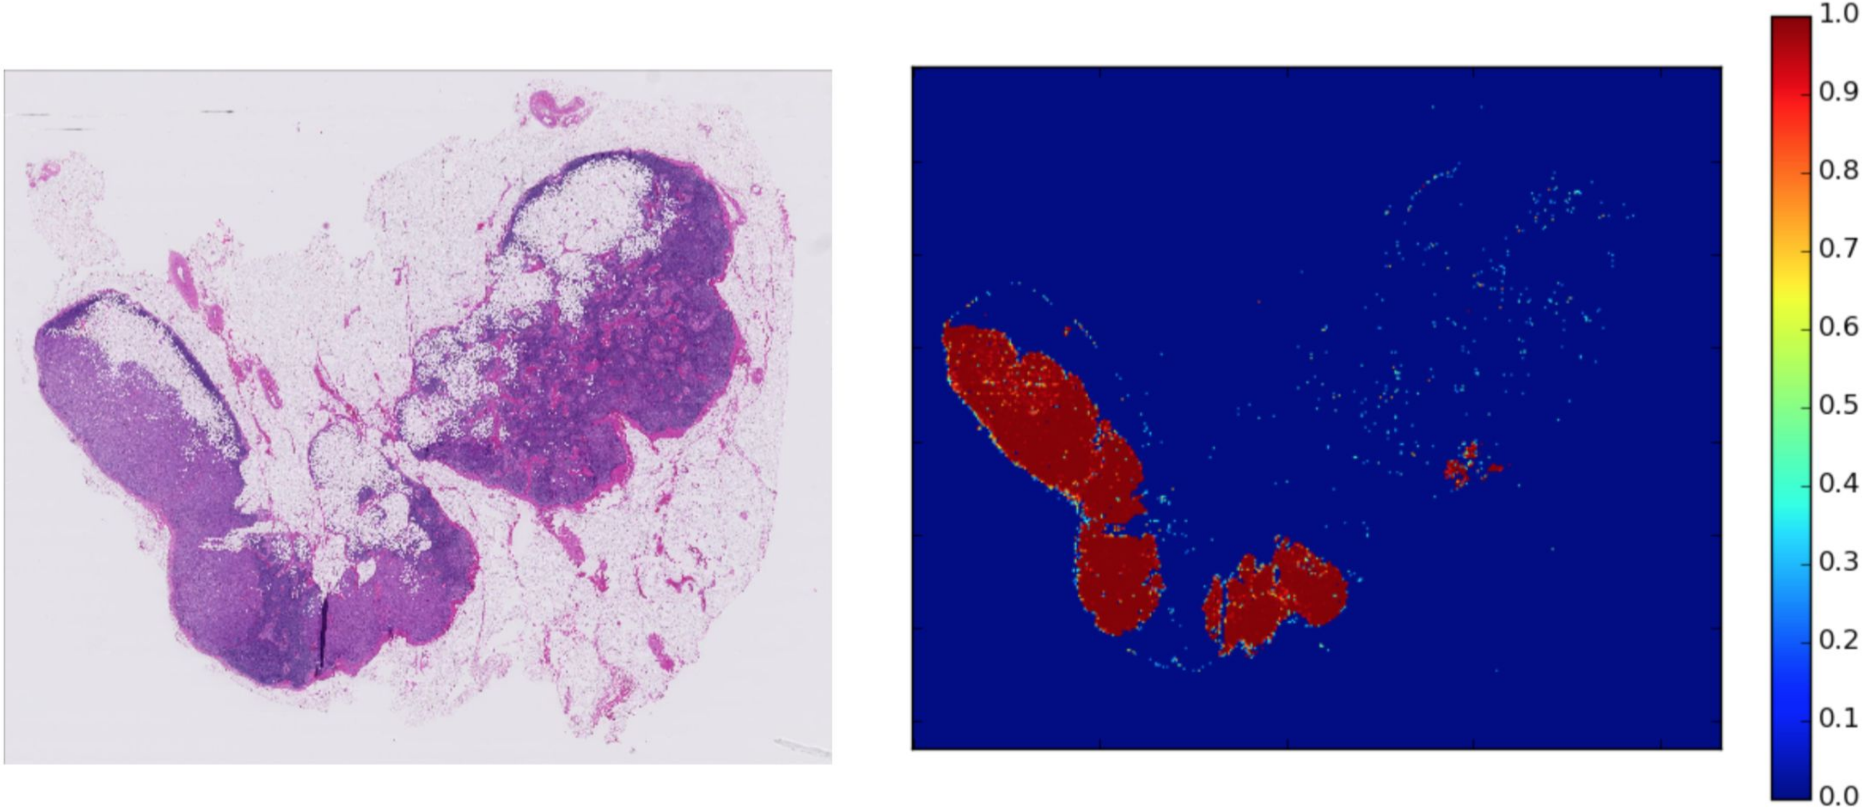

https://camelyon16.grand-challenge.org/
Image credit: Harvard Medical School, MIT, and EXB Research

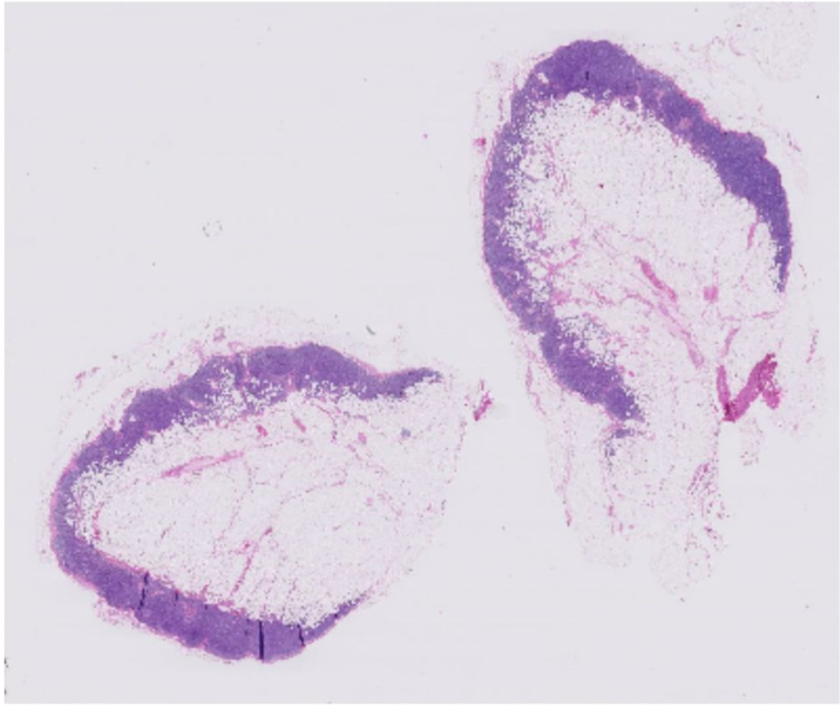

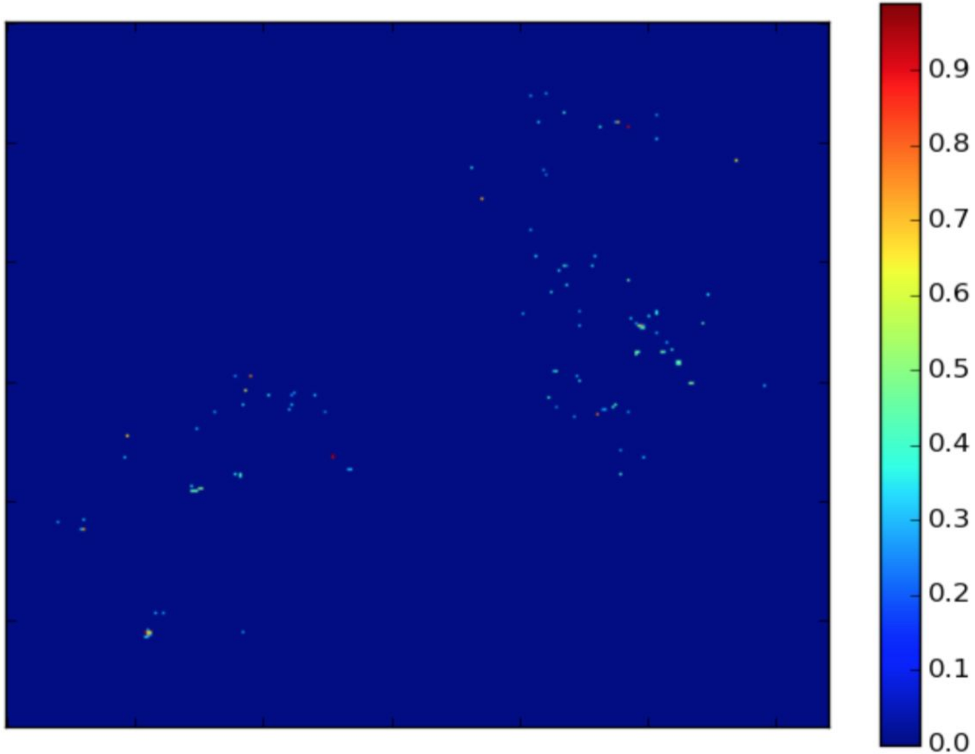

https://camelyon16.grand-challenge.org/
Image credit: Harvard Medical School, MIT, and EXB Research

High accuracy screening and very reproducible!

Team	AUC
Harvard Medical School and MIT, Method 1	0.9234
EXB Research and Development co., Germany	0.9156
Independent participant, Germany	0.8654
Middle East Technical University, Departments of EEE, NSNT and HS, Turkey	0.8642
NLP LOGIX co., USA	0.8298

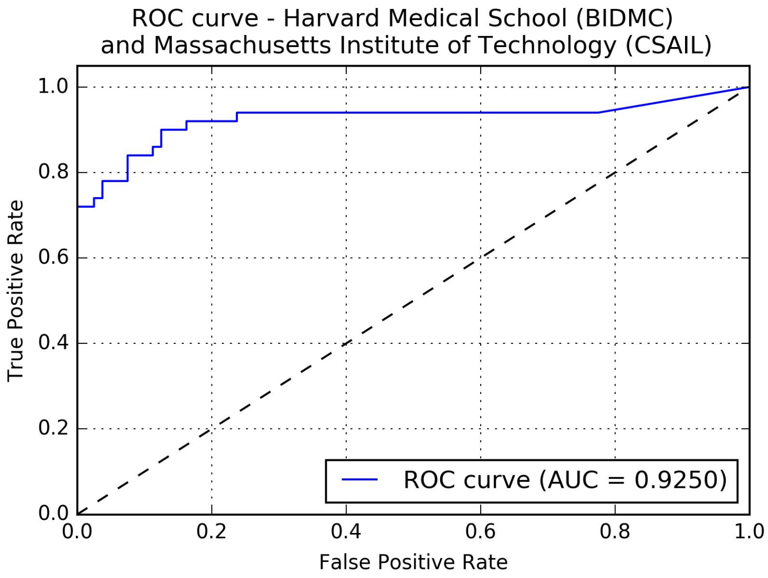

https://camelyon16.grand-challenge.org/

Image credit: Harvard Medical School, MIT, and EXB Research

How does it work?

Segmentation with Deep Learning

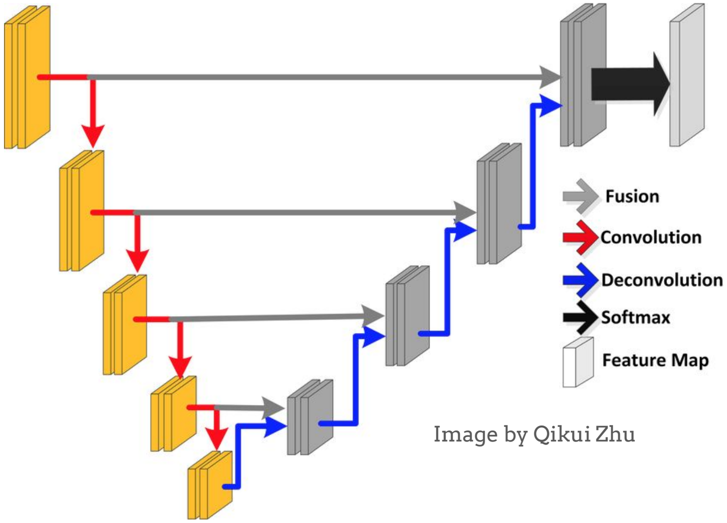

Image by Qikui Zhu

[Ronneberger et al., U-Net: Convolutional Networks for Biomedical Image Segmentation 2015]

[Honari et al., Recombinator Networks: Learning Coarse-to-Fine Feature Aggregation 2015]

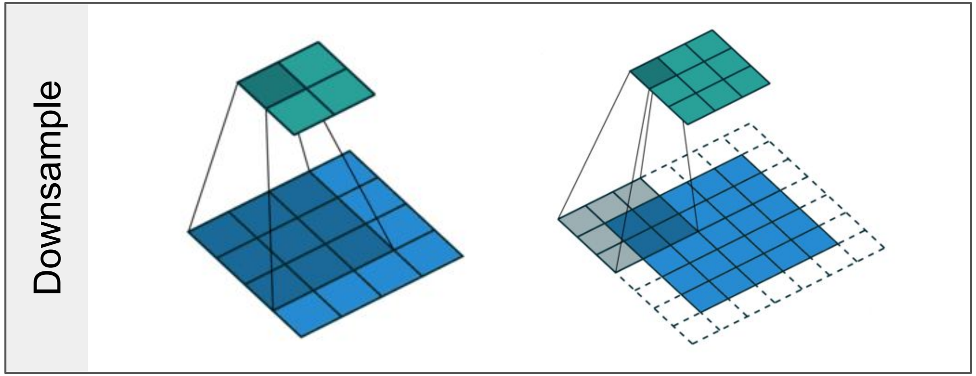

Images by Vincent Dumoulin (UdeM)

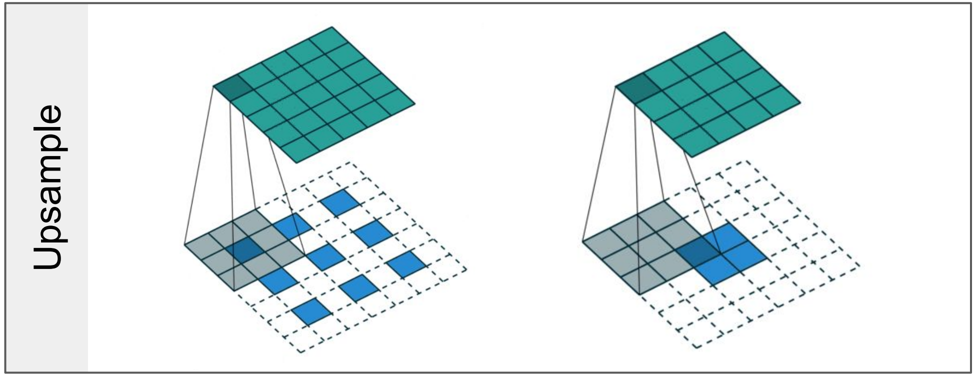

Segmentation training in progress

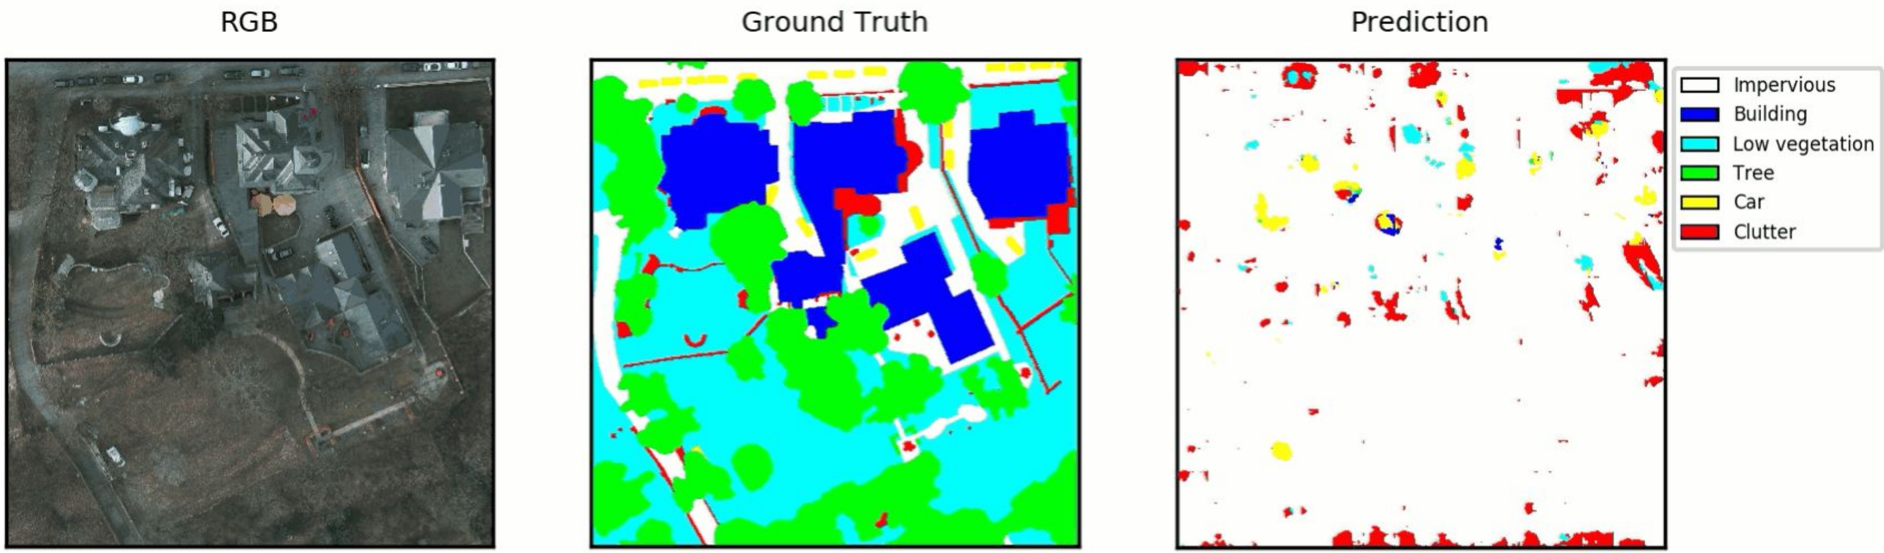

Image credit: Lewis Fishgold and Rob Emanuele

How can the network learn?

Example

stretch pixels into single column

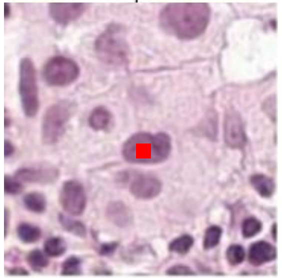

Image Patch

$$\underbrace{\begin{bmatrix} 0.2 & -0.5 & 0.1 & 2.0 \\ 1.5 & 1.3 & 2.1 & 0.0 \\ 0 & 0.25 & 0.2 & -0.3 \end{bmatrix}}_{W} \underbrace{\begin{bmatrix} 56 \\ 231 \\ 24 \\ 2 \end{bmatrix}}_{x_i} + \underbrace{\begin{bmatrix} 1.1 \\ 3.2 \\ -1.2 \end{bmatrix}}_{b} \longrightarrow \underbrace{\begin{bmatrix} -96.8 \\ 437.9 \\ 61.95 \end{bmatrix}}_{f(x_i; W, b)} \begin{array}{l} \text{No Tumor} \\ \text{Tumor} \\ \text{Necrosis} \end{array}$$

Given a patch predict the class of the center pixel

Image credits: http://cs231n.github.io/linear-classify/

Softmax

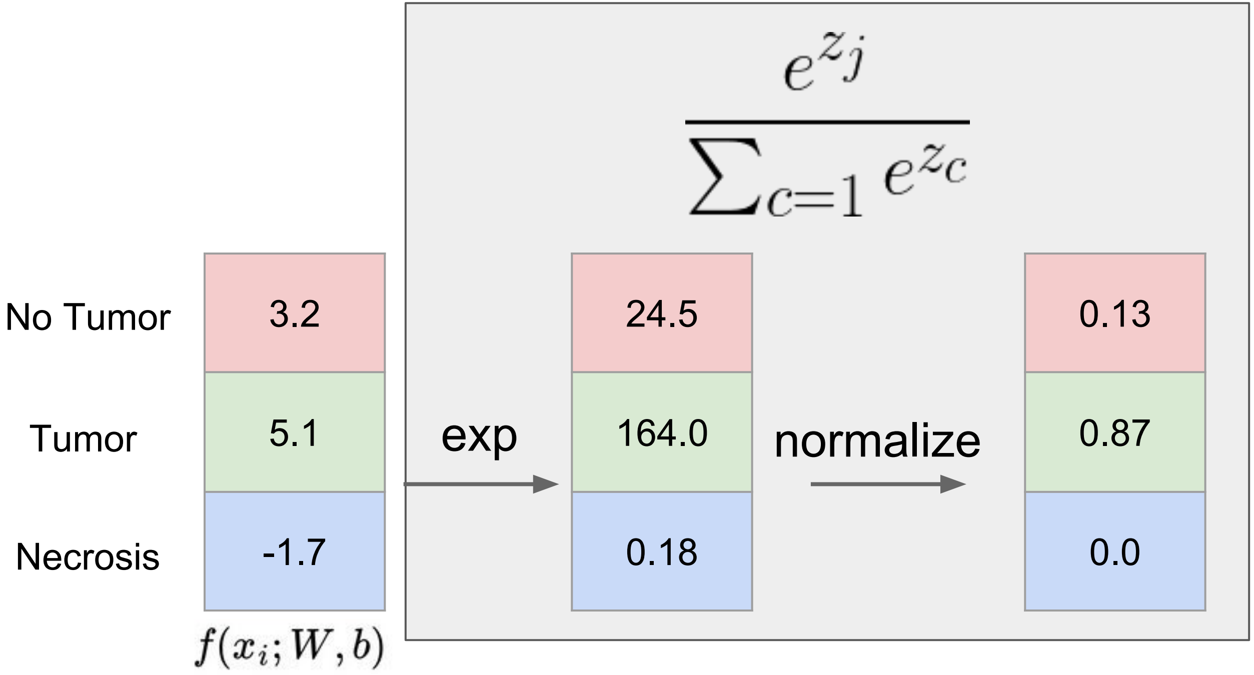

To predict multiple classes we project to a probability distribution

Simplex

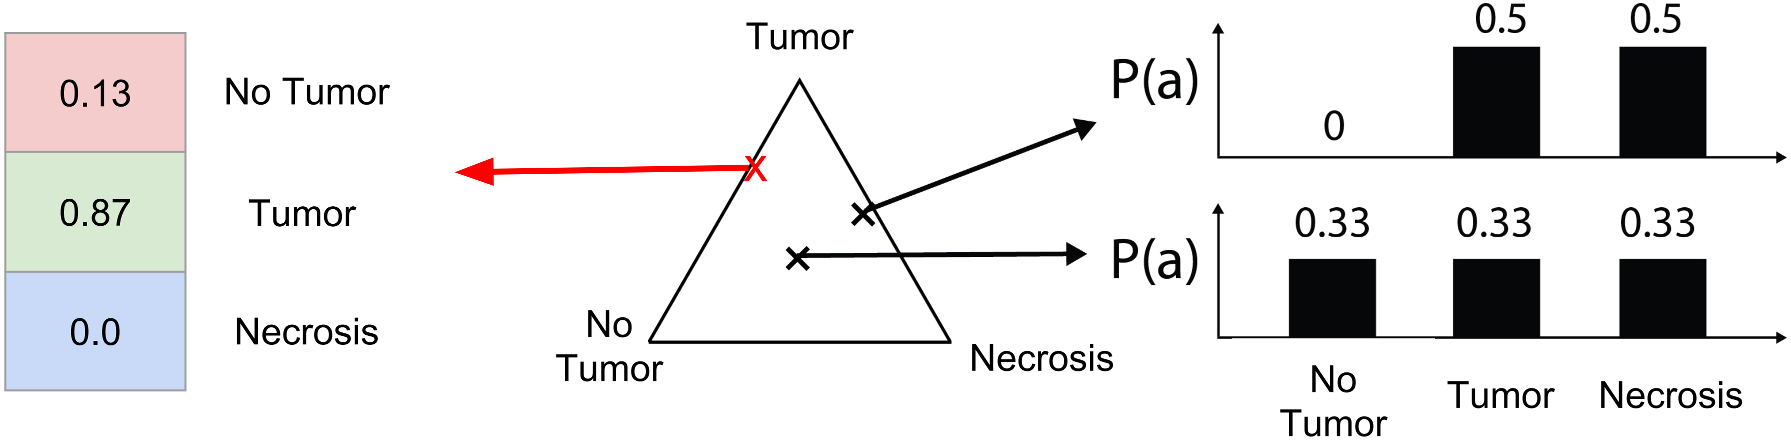

Because it is on a simplex; the correction of one term impacts all

Image credits: http://gureckislab.org/

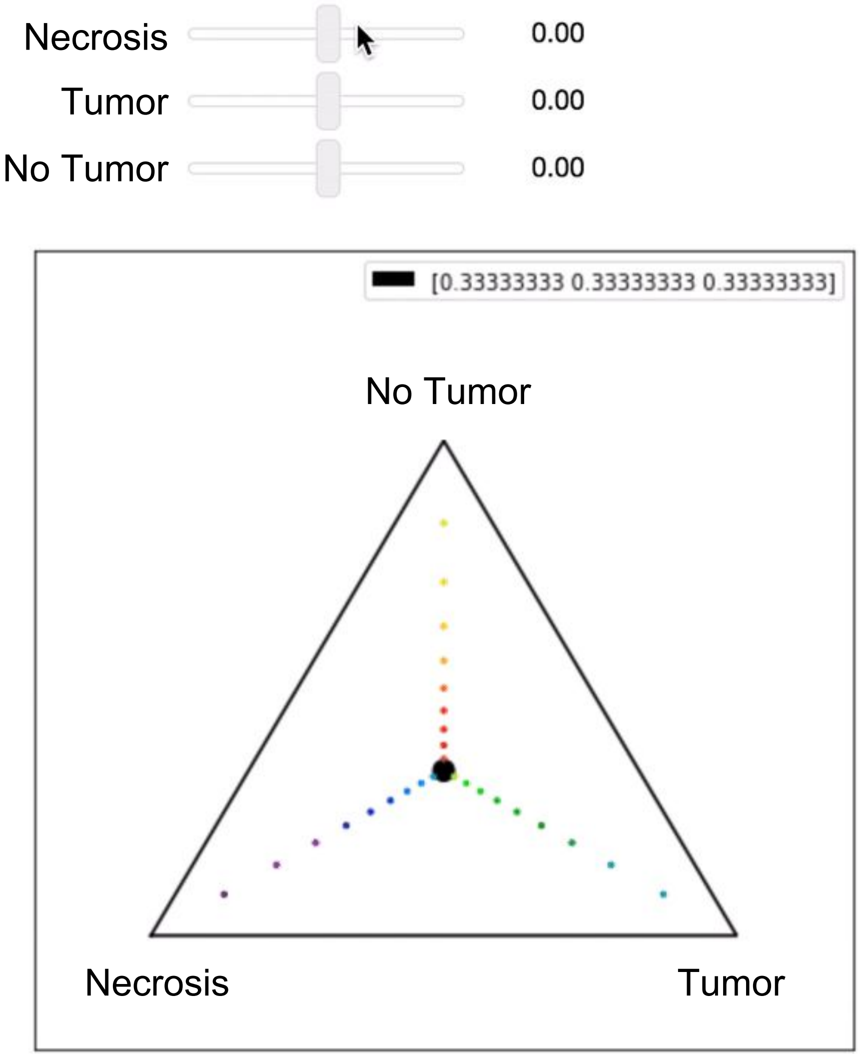

Infinite ways to generate the same output.

A correction of one sends gradients to others

We can learn unseen classes through a process of elimination.

https://github.com/ieee8023/NeuralNetwork-Examples/blob/master/general/simplex-softmax.ipynb

Softmax and Cross-entropy loss

$$\frac{e^{z_j}}{\sum_{c=1} e^{z_c}}$$

$$-\sum_{c} y_i log(p_c)$$

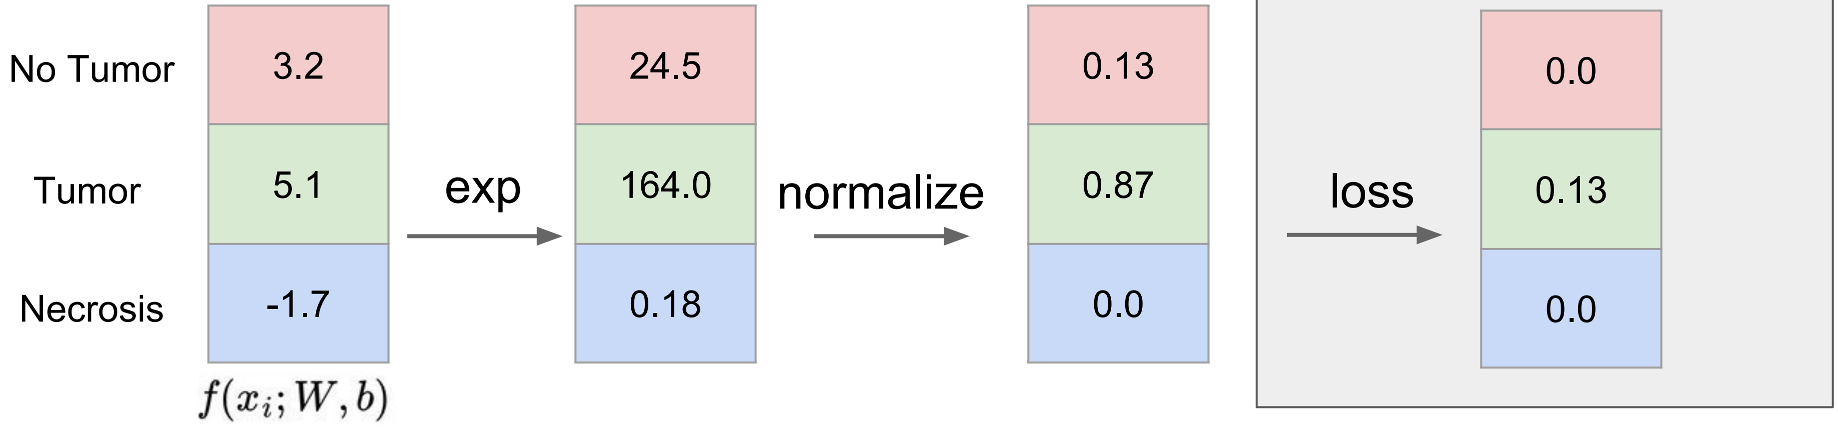

To predict multiple class we can project the output onto a simplex and compute the loss there.

More segmentation tasks!

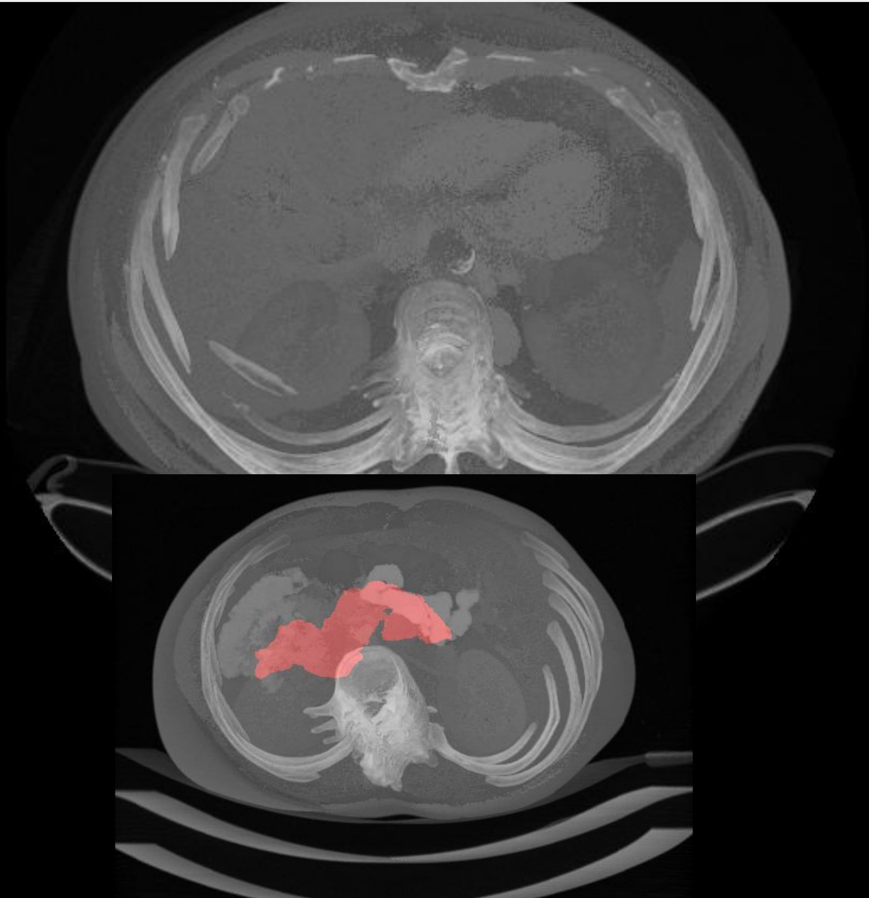

CT Pancreas Segmentation

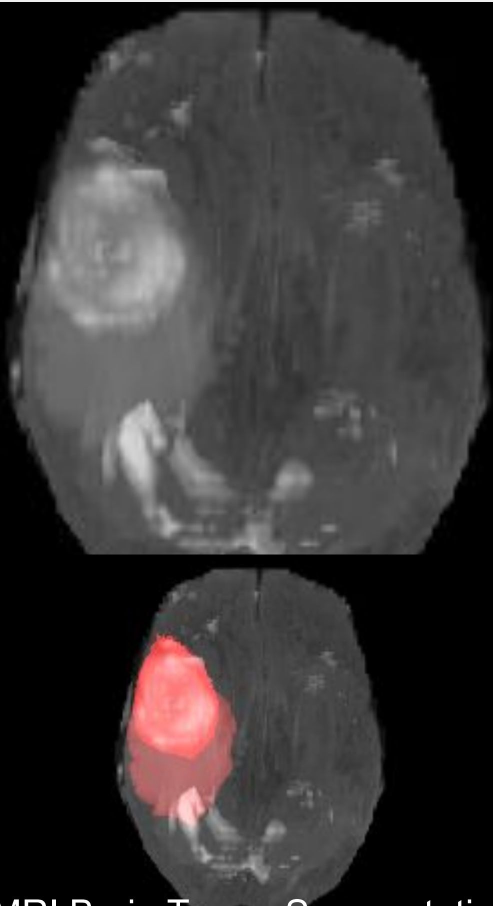

MRI Brain Tumor Segmentation

More segmentation tasks!

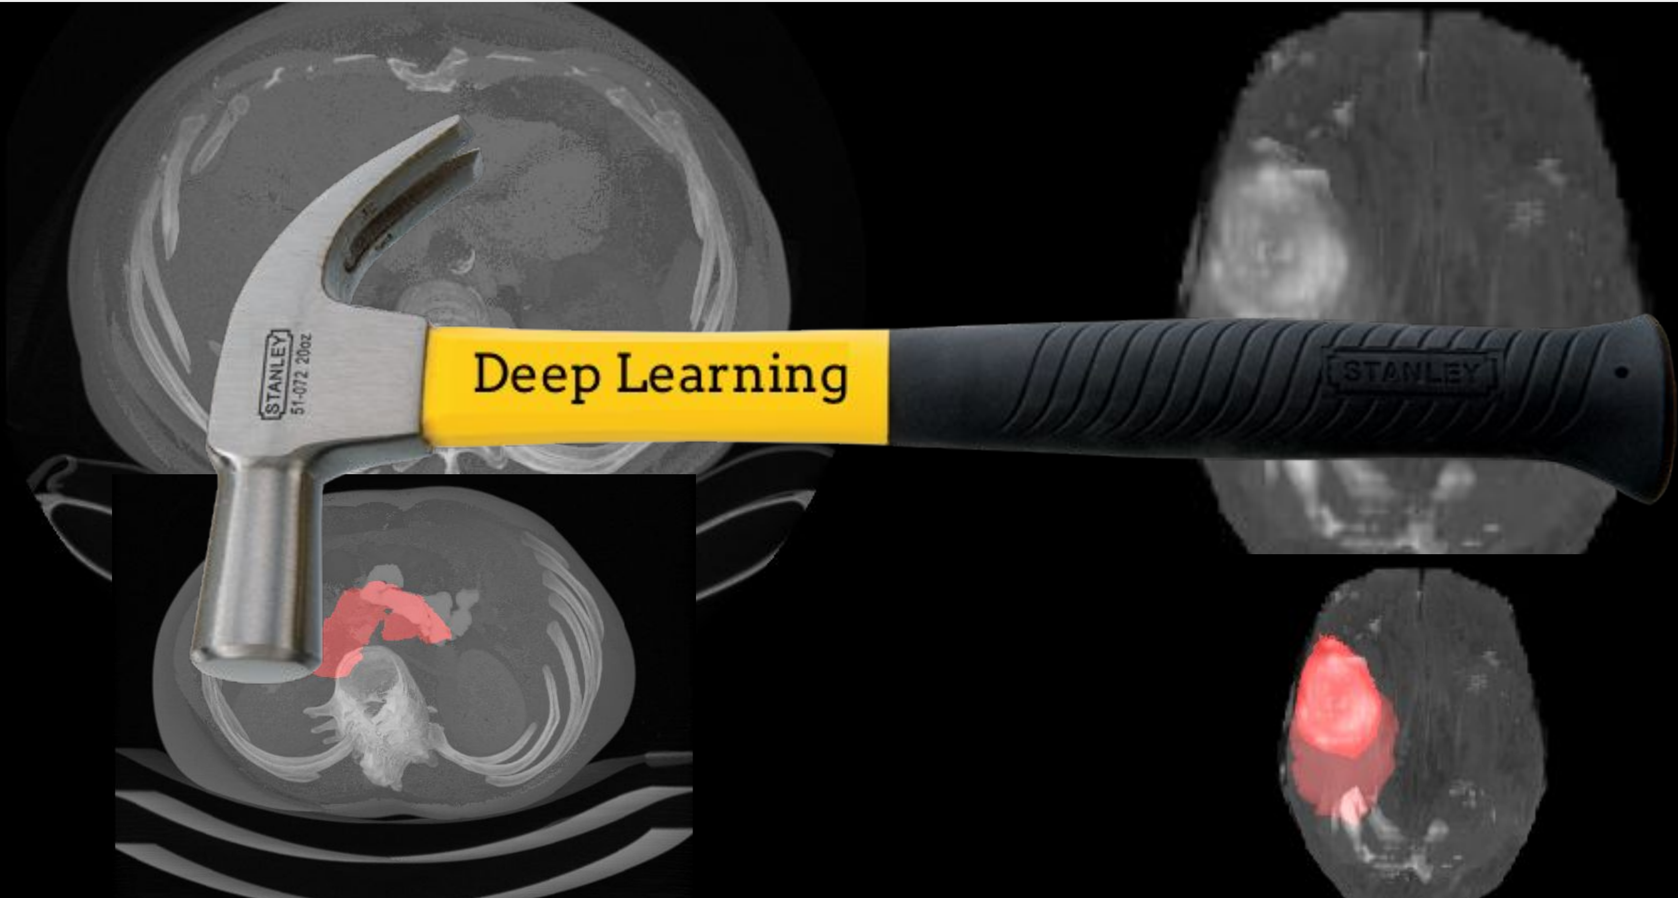

CT Pancreas Segmentation

MRI Brain Tumor Segmentation

More segmentation tasks!

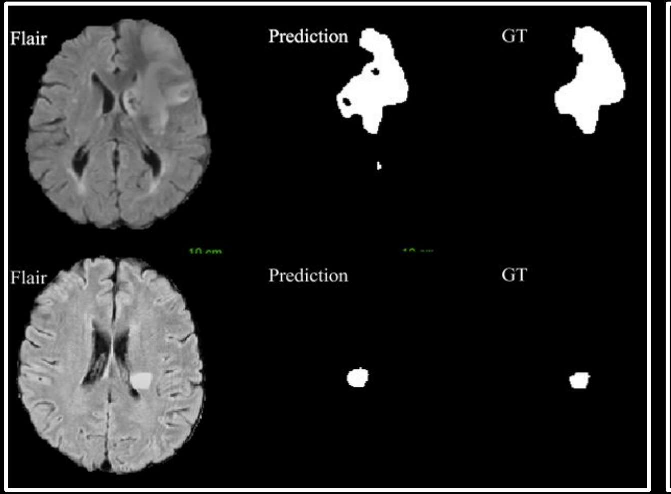

Sub-acute ischemic stroke lesion segmentation

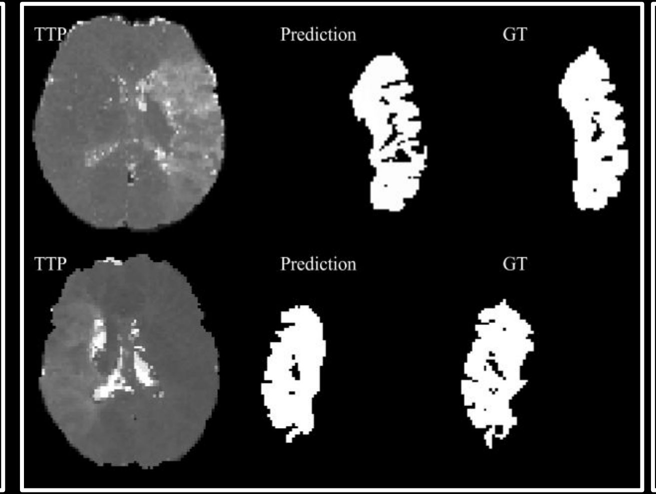

Stroke Perfusion Estimation

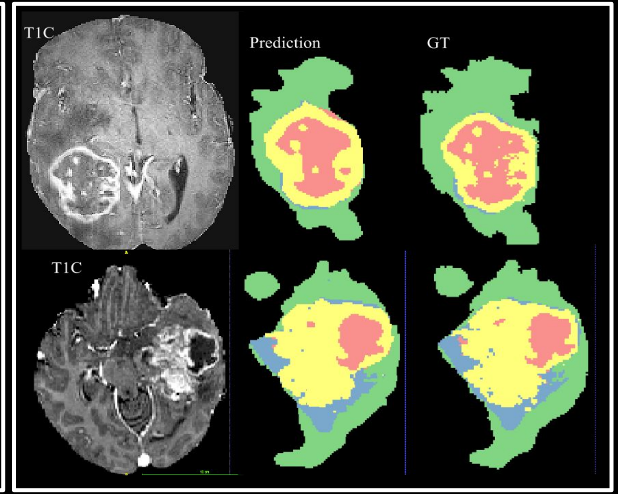

Brain tumor segmentation

Image Credit: Mohammad Havaei

Multi-modal MRI acquisition

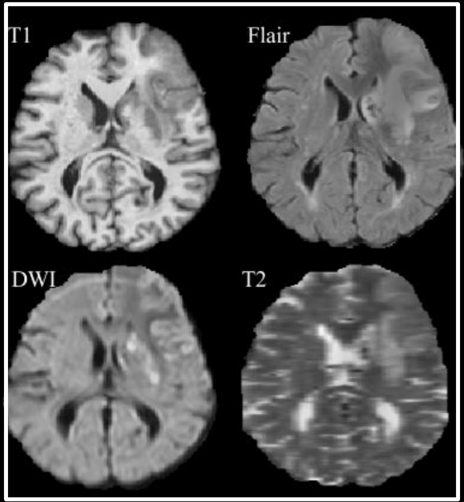

Sub-acute ischemic stroke lesion segmentation

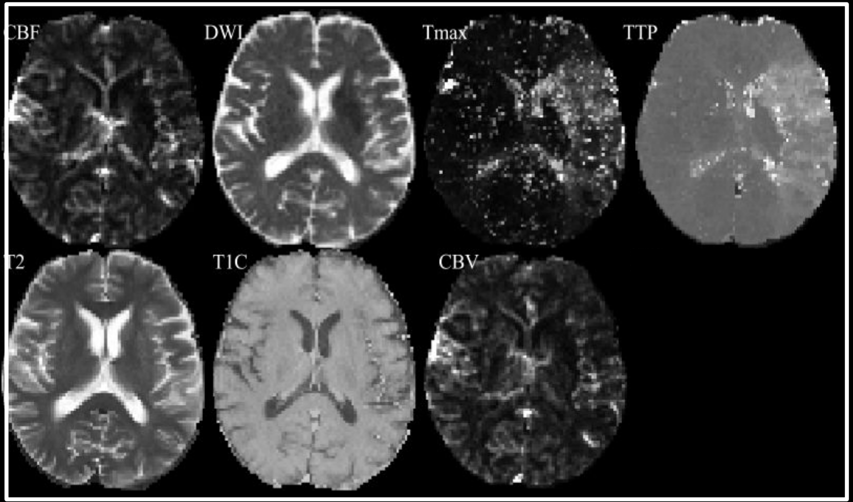

Stroke Perfusion Estimation

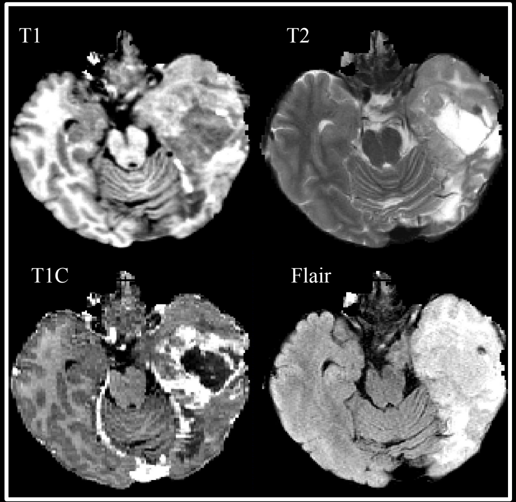

Brain tumor segmentation

Image Credit: Mohammad Havaei

Multi-modal MRI acquisition

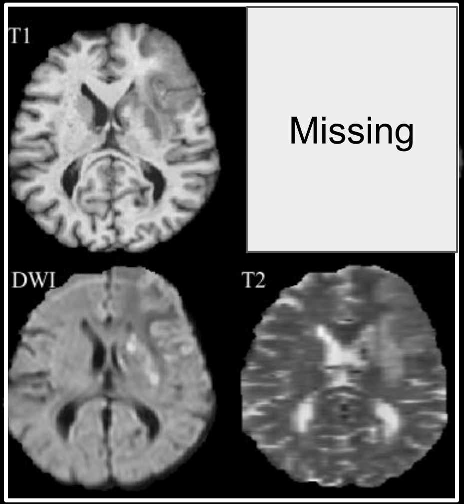

Sub-acute ischemic stroke
lesion segmentation

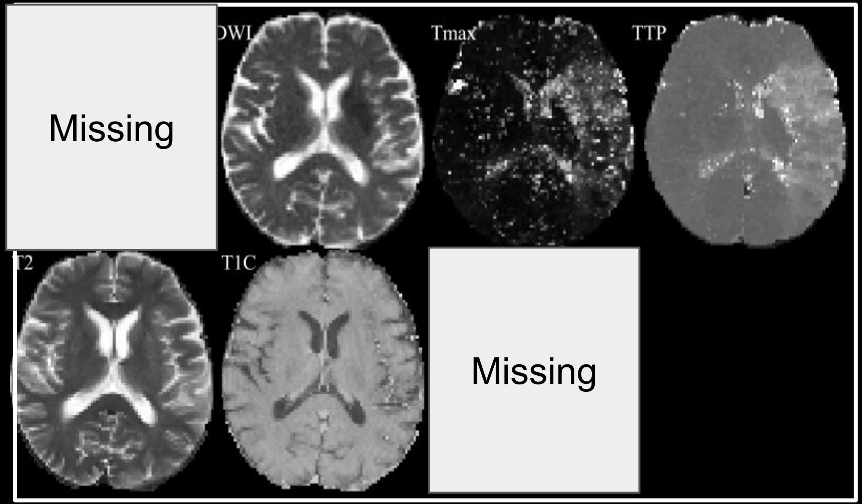

Stroke Perfusion
Estimation

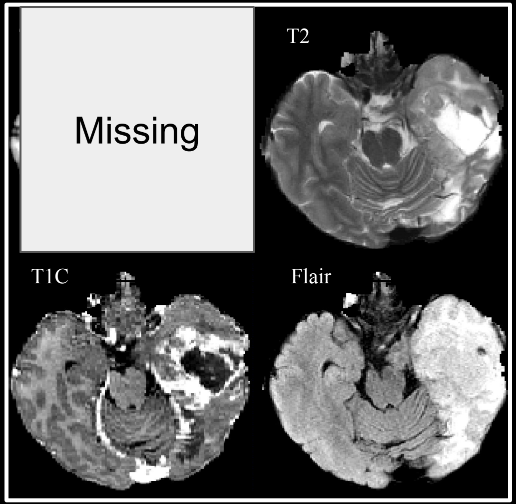

Brain tumor
segmentation

Image Credit: Mohammad Havaei

HeMIS: Hetero-Modal Image Segmentation

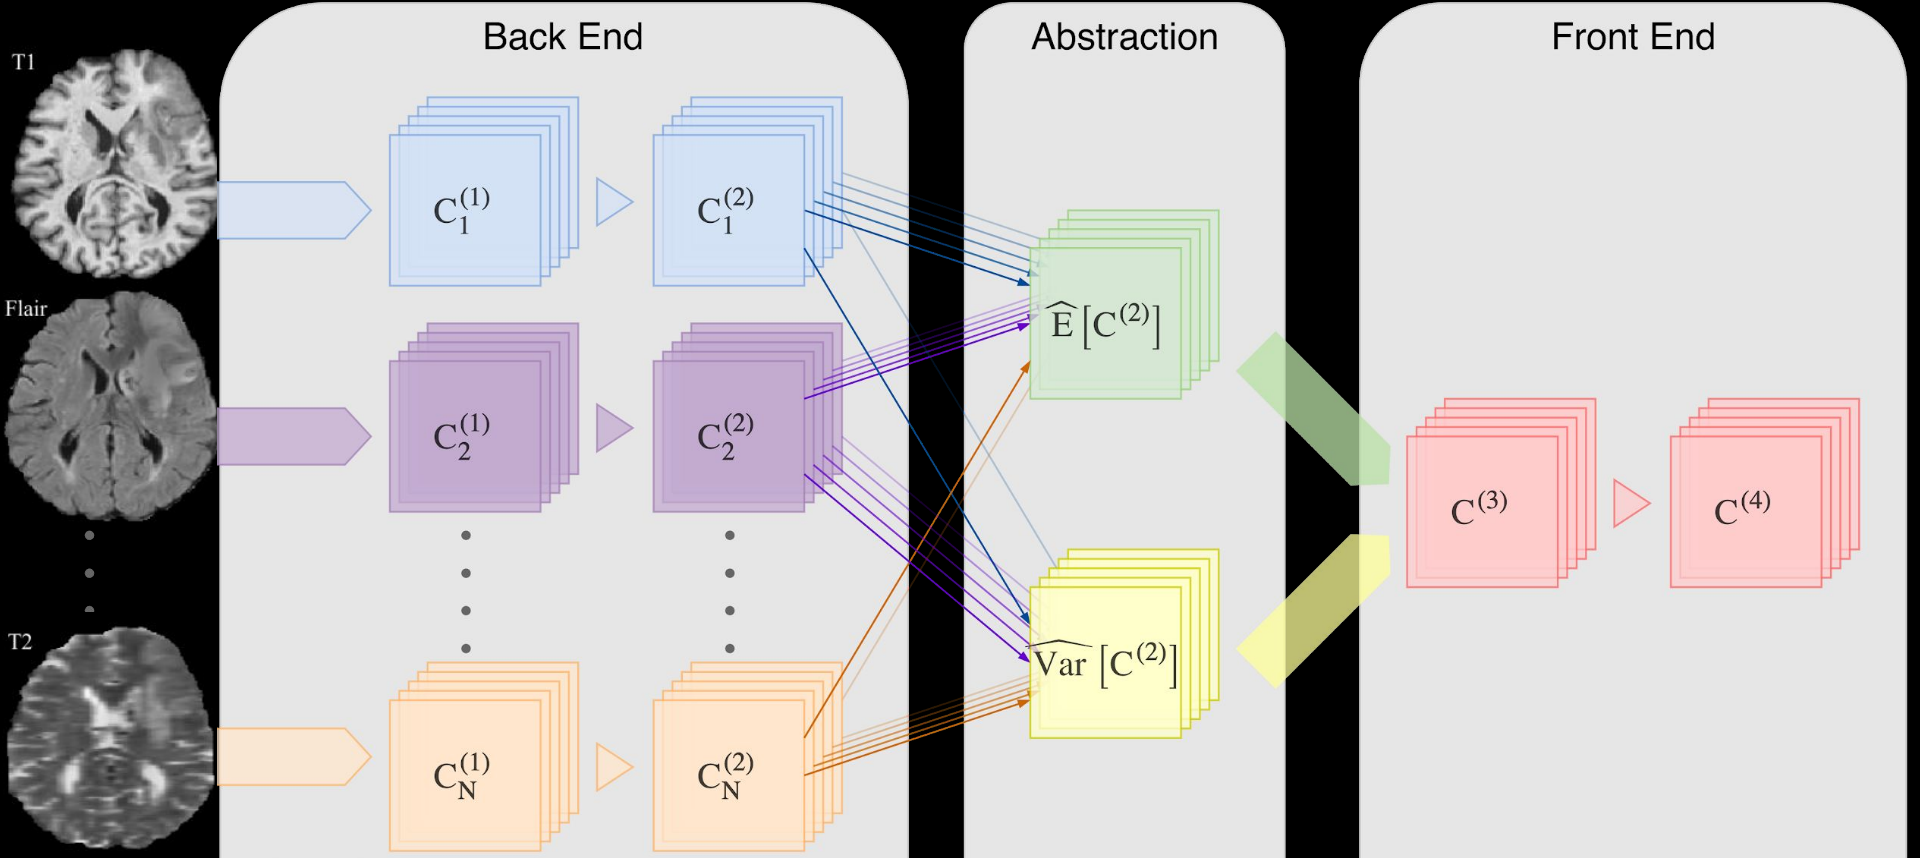

HeMIS: Hetero-Modal Image Segmentation

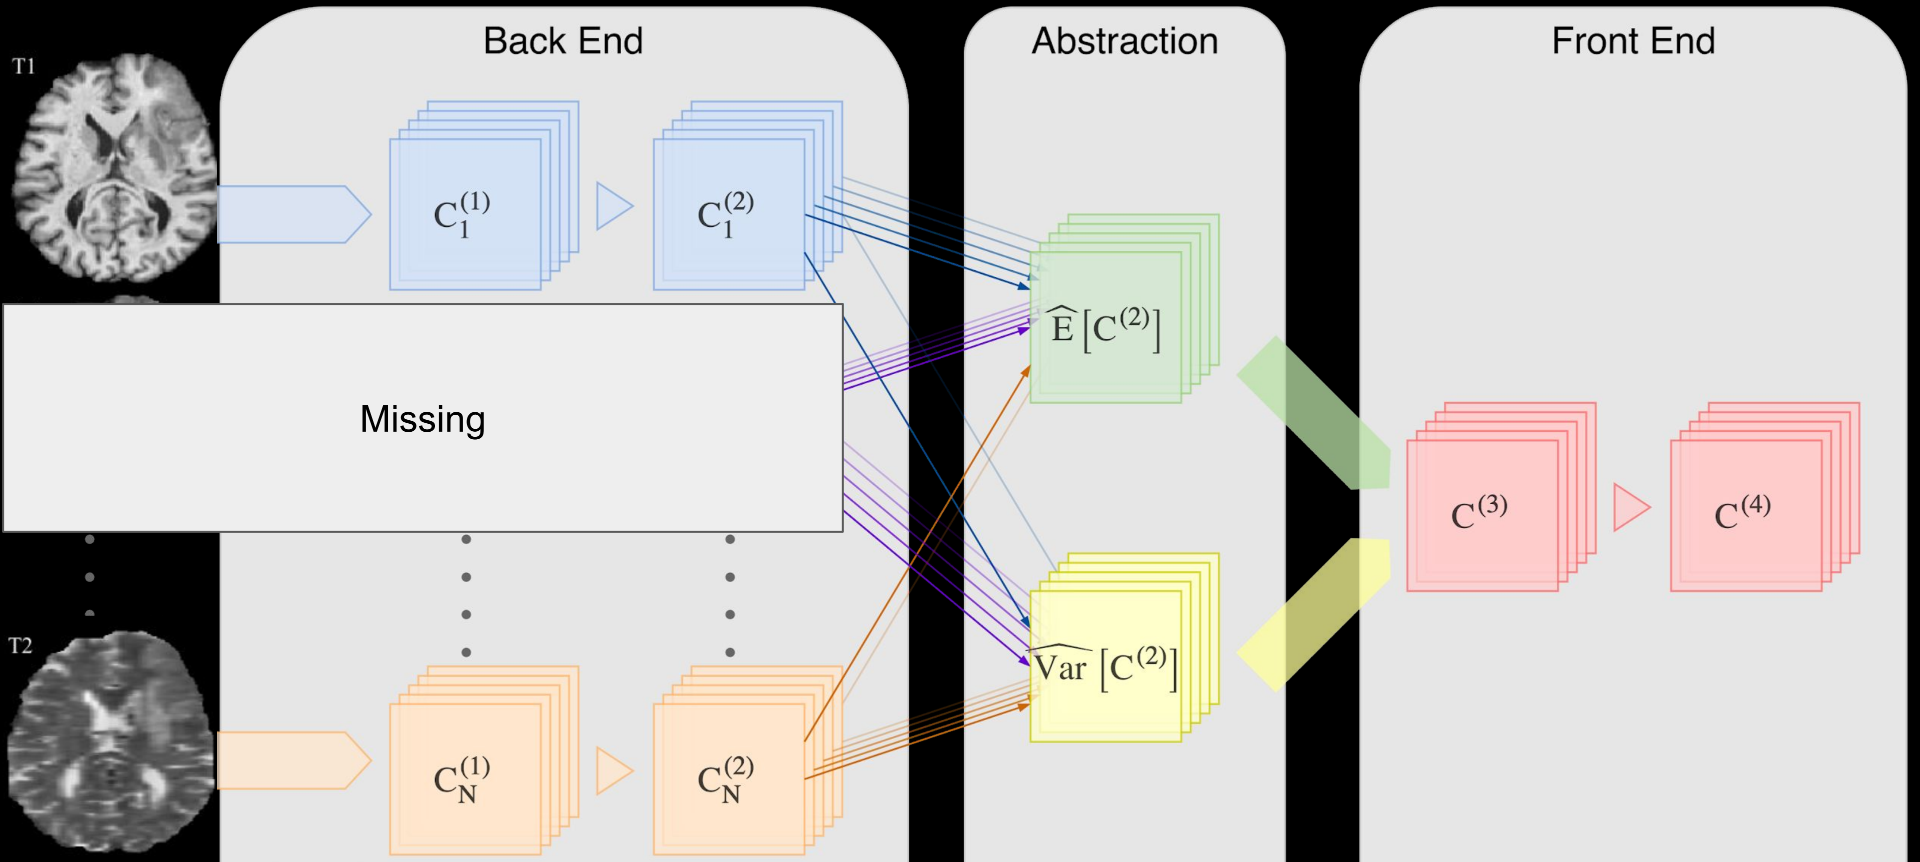

Results (BRATS)

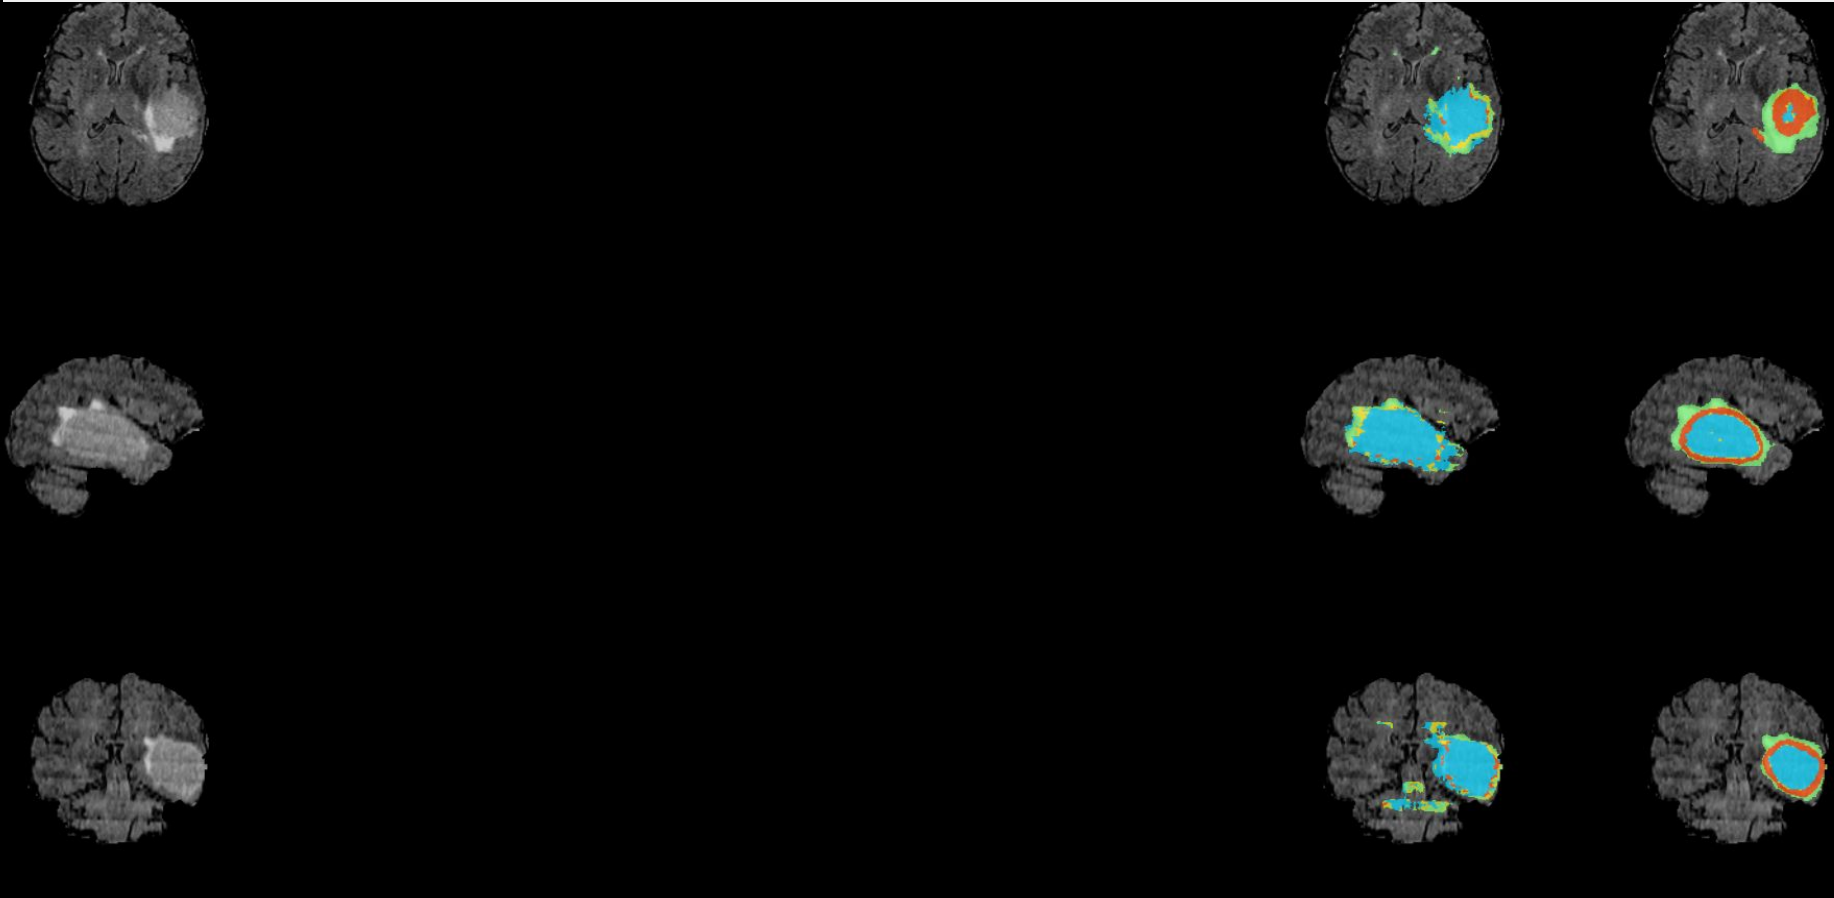

Results (BRATS)

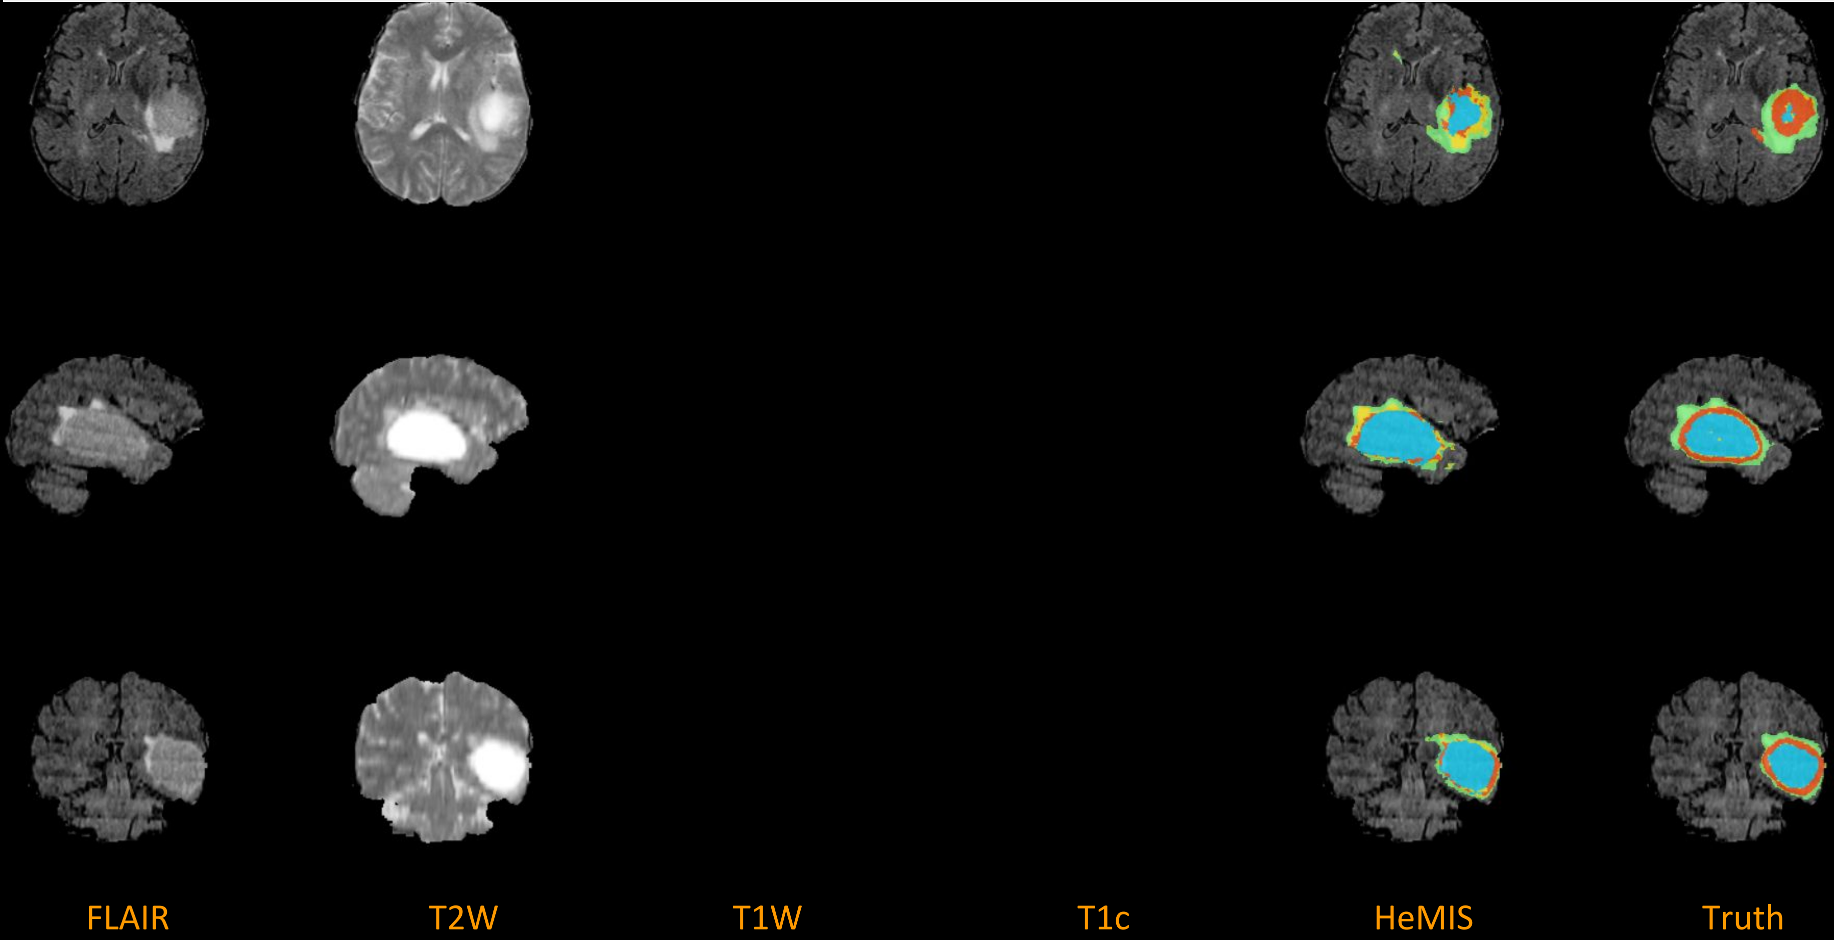

Results (BRATS)

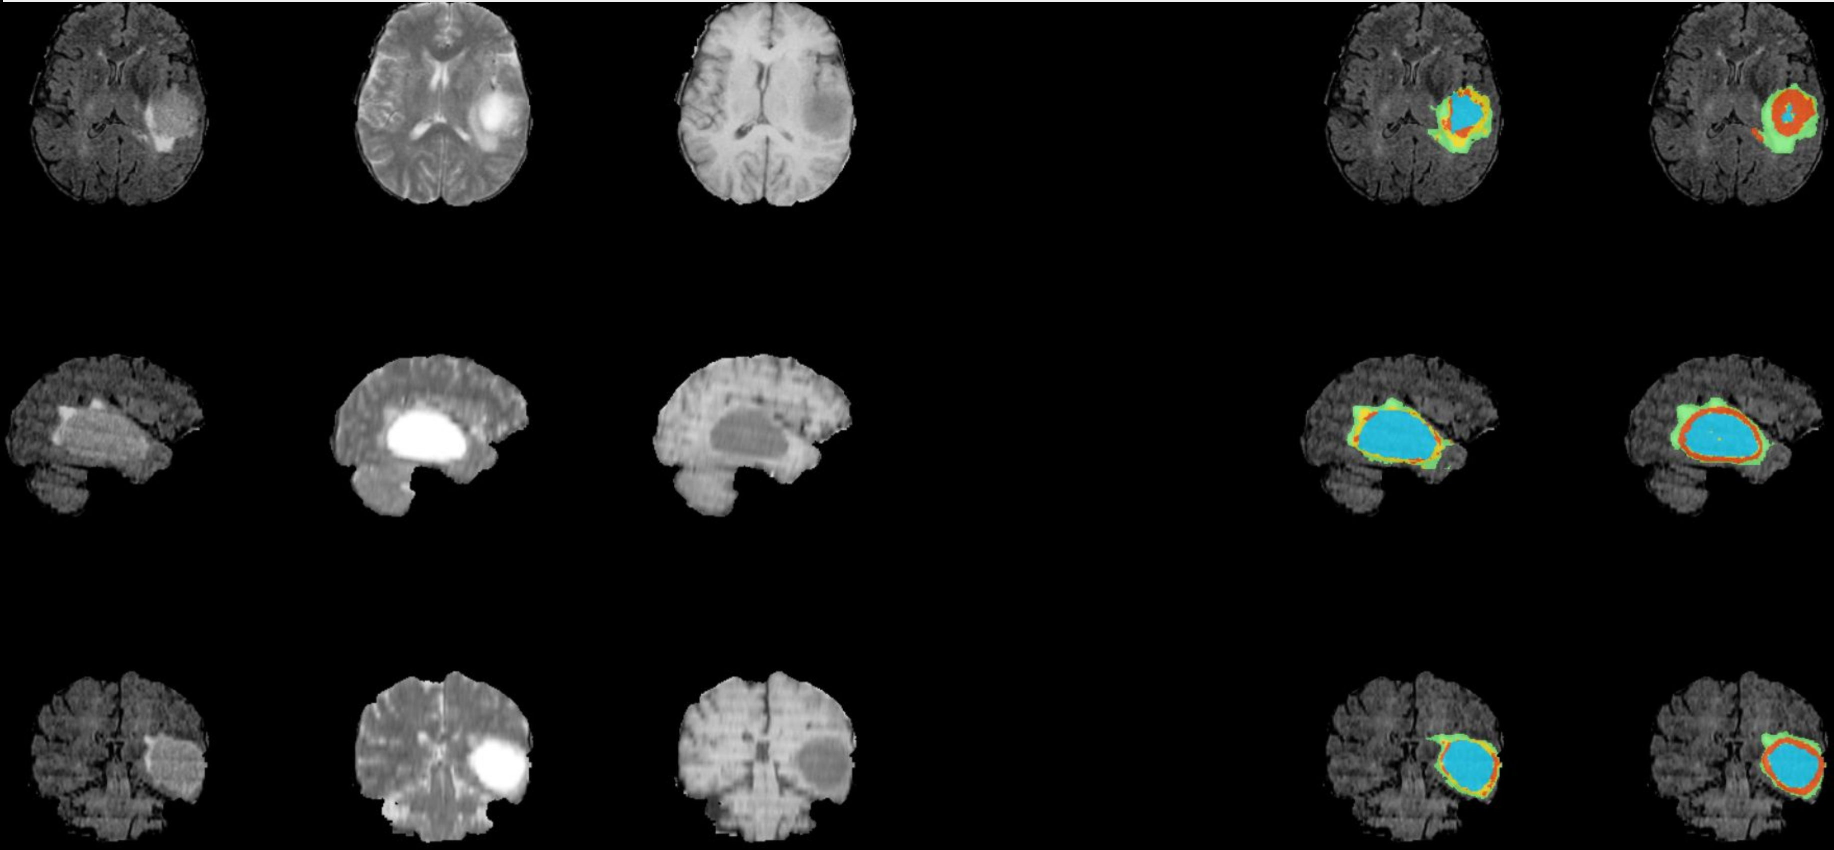

Results (BRATS)

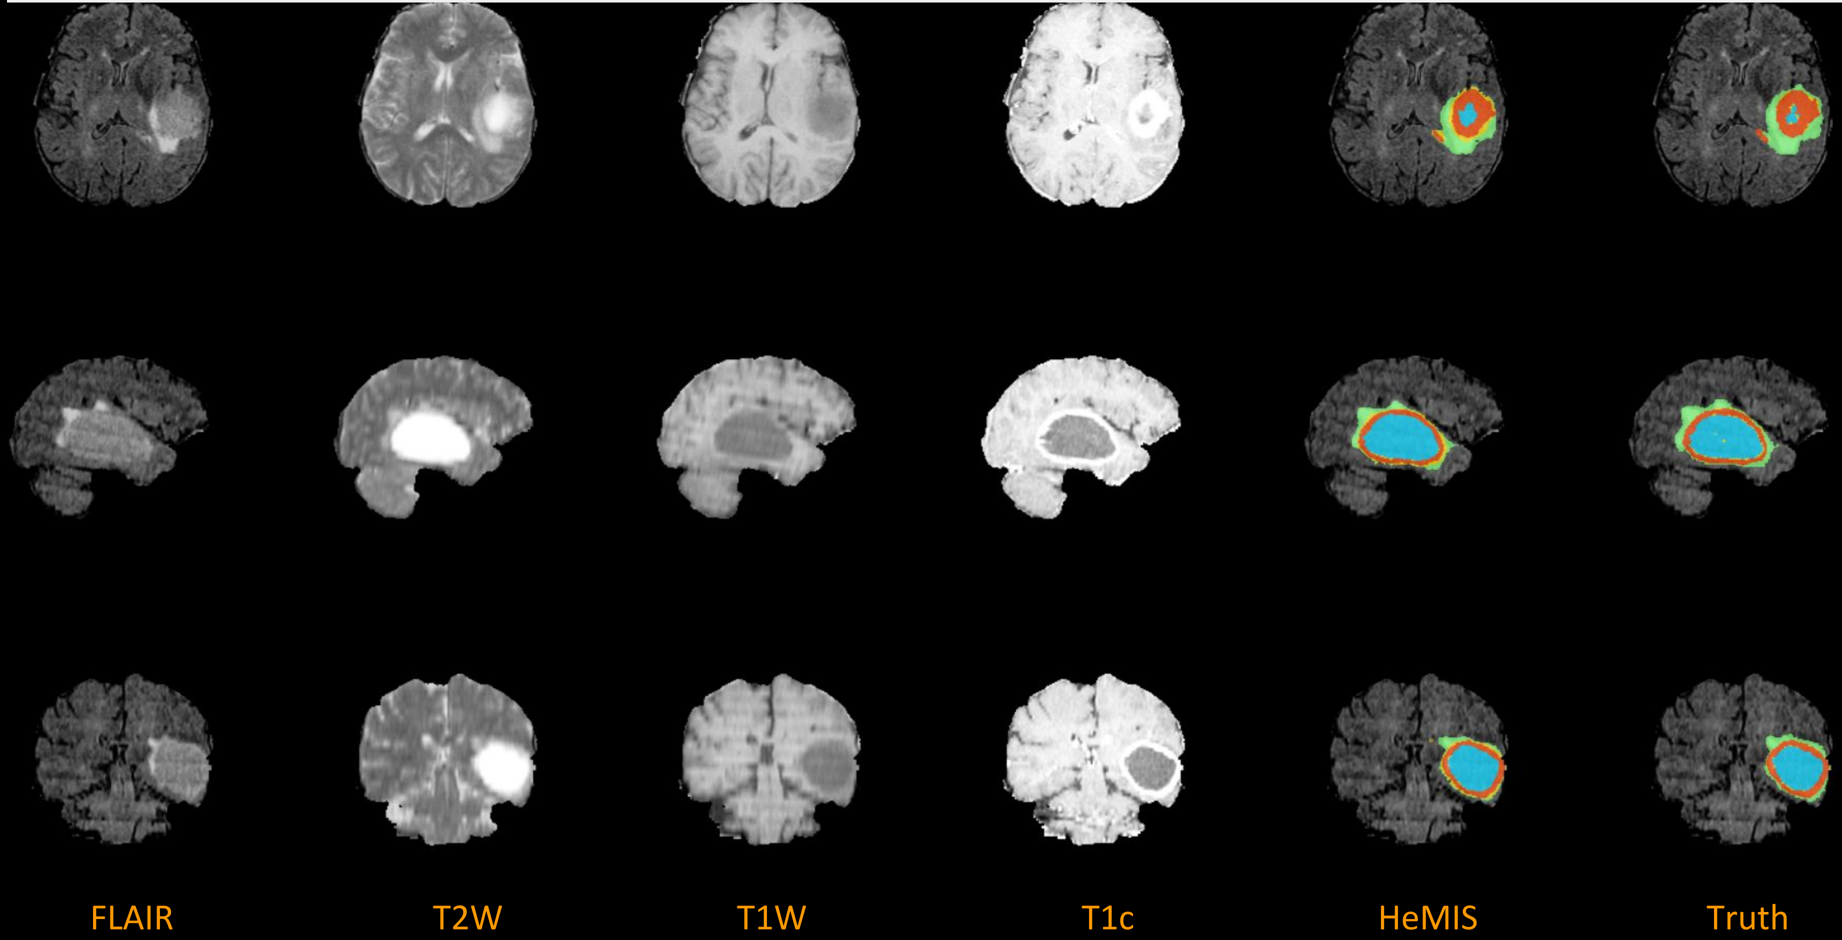

Brain tumor segmentation (BRATS2013 dataset)

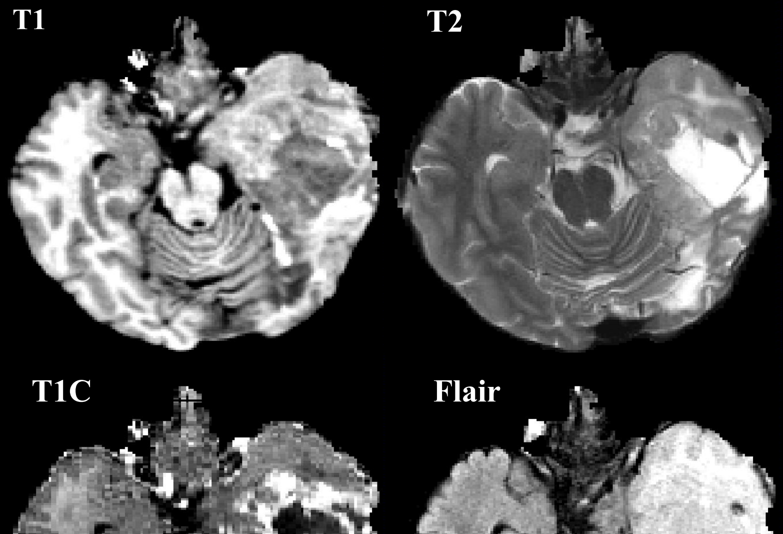

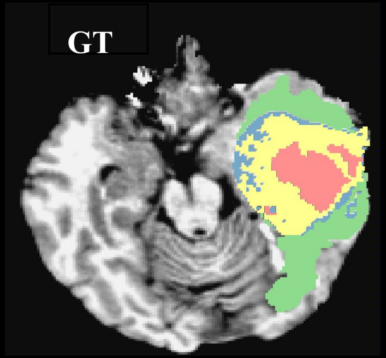

Training data:
220 subjects with high grade and
54 subjects with low grade tumors

Dice Similarity	Leaderboard			Challenge		
Method	Complete	Core	Enhancing	Complete	Core	Enhancing
Tustison	79	65	53	87	**78**	**74**
Zhao	79	59	47	84	70	65
Meier	72	60	53	82	73	69
HeMIS	**83**	**67**	**57**	**88**	75	**74**


HeMIS:
Hetero-Modal Image Segmentation *

Mohammad Havaei[1,2], Nicolas Guizard[1], Nicolas Chapados[1], and Yoshua Bengio[3]

[1] Imagia Inc., Montreal, Qc, Canada
[2] Université de Sherbrooke, Qc, Canada
[3] Université de Montréal, Montréal, Canada
{mohammad.havaei,nicolas.guizard,nicolas.chapados}@imagia.com
seyed.mohammad.havaei@usherbrooke.ca
yoshua.bengio@umontreal.ca
http://www.imagia.com

Abstract. We introduce a deep learning image segmentation framework that is extremely robust to missing imaging modalities. Instead of attempting to impute or synthesize missing data, the proposed approach learns, for each modality, an embedding of the input image into a single latent vector space for which arithmetic operations (such as taking the mean) are well defined. Points in that space, which are averaged over modalities available at inference time, can then be further processed to yield the desired segmentation. As such, any combinatorial subset of available modalities can be provided as input, without having to learn a combinatorial number of imputation models. Evaluated on two neuro-

cs.CV] 18 Jul 2016

[Havaei et al. HeMIS: Hetero-Modal Image Segmentation, MICCAI 2016]


Deep Learning
STANLEY
51-072 20oz
STANLEY

Interpreting Chest X-Rays

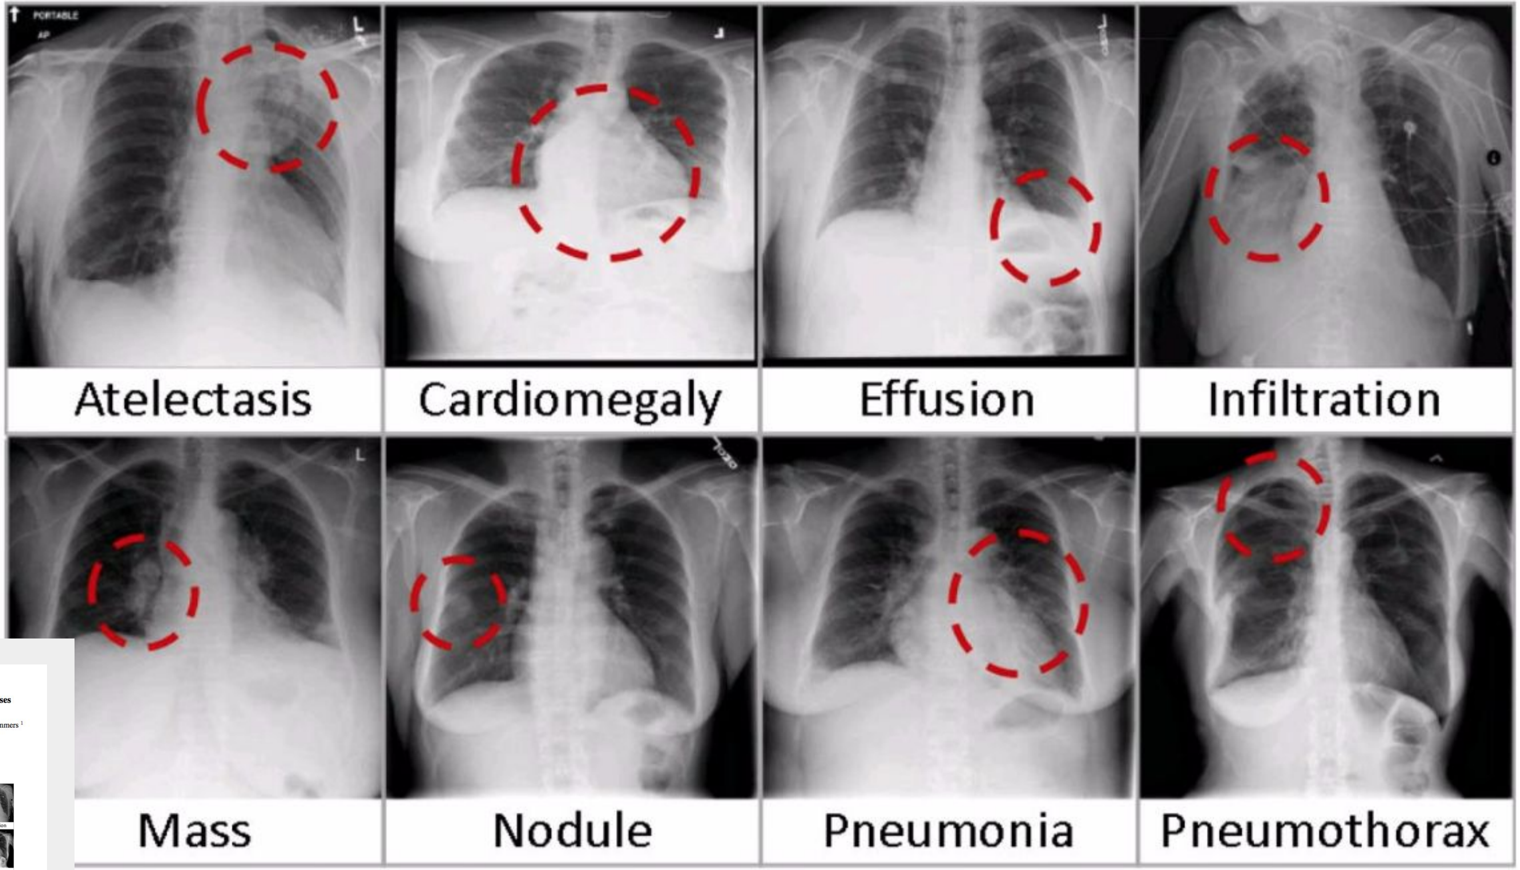

ChestX-ray8: Hospital-scale Chest X-ray Database and Benchmarks on Weakly-Supervised Classification and Localization of Common Thorax Diseases

Xiaosong Wang[1], Yifan Peng [2], Le Lu [1], Zhiyong Lu [2], Mohammadhadi Bagheri [1], Ronald M. Summers [1]
[1]Department of Radiology and Imaging Sciences, Clinical Center,
[2] National Center for Biotechnology Information, National Library of Medicine,
National Institutes of Health, Bethesda, MD 20892
{xiaosong.wang,yifan.peng,le.lu,luzh,mohammad.bagheri,rms}@nih.gov

05.02315v5 [cs.CV] 14 Dec 2017

Abstract

The chest X-ray is one of the most commonly accessible radiological examinations for screening and diagnosis of many lung diseases. A tremendous number of X-ray imaging studies accompanied by radiological reports are accumulated and stored in many modern hospitals' Picture Archiving and Communication Systems (PACS). On the other side, it is still an open question how this type of hospital-size knowledge database containing invaluable imaging informatics (i.e., loosely labeled) can be used to facilitate the data-hungry deep learning paradigms in building truly large-scale high precision computer-aided diagnosis (CAD) systems.

In this paper, we present a new chest X-ray database, namely "ChestX-ray8", which comprises 108,948 frontal-view X-ray images of 32,717 unique patients with the text-mined eight disease image labels (where each image can have multi-labels), from the associated radiological reports

Figure 1. *Eight common thoracic diseases observed in chest X-rays that validate a challenging task of fully-automated diagnosis.*

comparison to previous shallow methodologies built upon hand-crafted image features. Deep neural network representations further make the joint language and vision learning tasks more feasible to solve, in image captioning [49, 24, 33, 48, 23], visual question answering [2, 46, 51, 55]

Wang, ChestX-ray8: Hospital-scale Chest X-ray Database and Benchmarks on Weakly-Supervised Classification and Localization of Common Thorax Diseases, 2017

Convolutional Neural Net Predictions

Setting	Atelectasis	Cardiomegaly	Effusion	Infiltration	Mass	Nodule	Pneumonia	Pneumothorax
AlexNet	0.6458	0.6925	0.6642	0.6041	**0.5644**	0.6487	0.5493	0.7425
GoogLeNet	0.6307	0.7056	0.6876	0.6088	0.5363	0.5579	0.5990	0.7824
VGGNet-16	0.6281	0.7084	0.6502	0.5896	0.5103	0.6556	0.5100	0.7516
ResNet-50	**0.7069**	**0.8141**	**0.7362**	**0.6128**	0.5609	**0.7164**	**0.6333**	**0.7891**

AUCs for each class of a multi-target model

32,717 patients

108,948 X-Rays

Text extracted using NLP from doctors notes

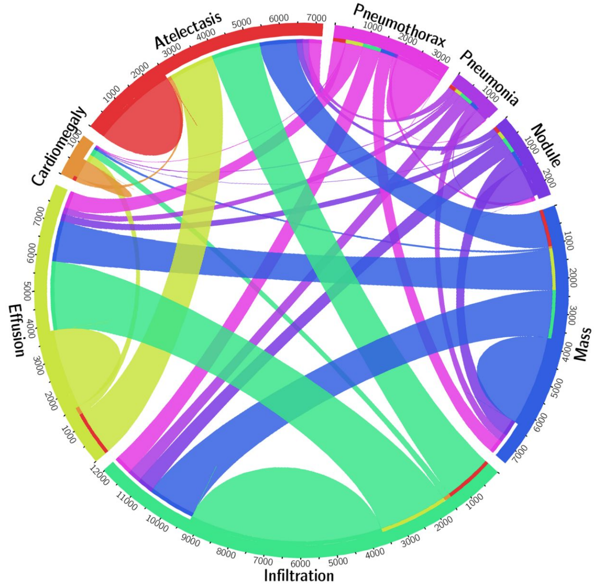

Radiology report	Keyword	Localization Result
findings include: 1. left basilar atelectasis/consolidation. 2. prominent hilum (mediastinal adenopathy). 3. left pic catheter (tip in atriocaval junction). 4. stable, normal appearing cardiomediastinal silhouette. impression: small right pleural effusion otherwise stable abnormal study including left basilar infiltrate/atelectasis, prominent hilum, and position of left pic catheter (tip atriocaval junction).	Effusion; Infiltration; Atelectasis	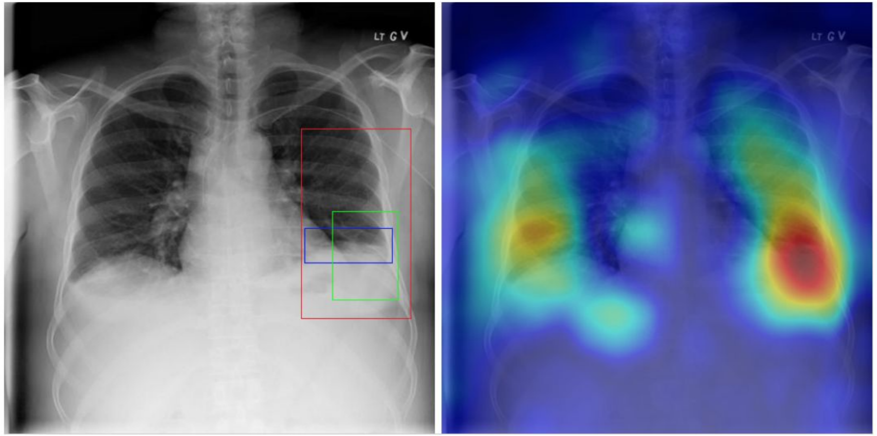

Table 8. *A sample of chest x-ray radiology report, mined disease keywords and localization result from the "Atelectasis" Class. Correct bounding box (in green), false positives (in red) and the ground truth (in blue) are plotted over the original image.*

Radiology report	Keyword	Localization Result
findings: no appreciable change since XX/XX/XX. small right pleural effusion. elevation right hemidiaphragm. diffuse small nodules throughout the lungs, most numerous in the left mid and lower lung. impression: no change with bilateral small lung metastases.	Effusion; Nodule	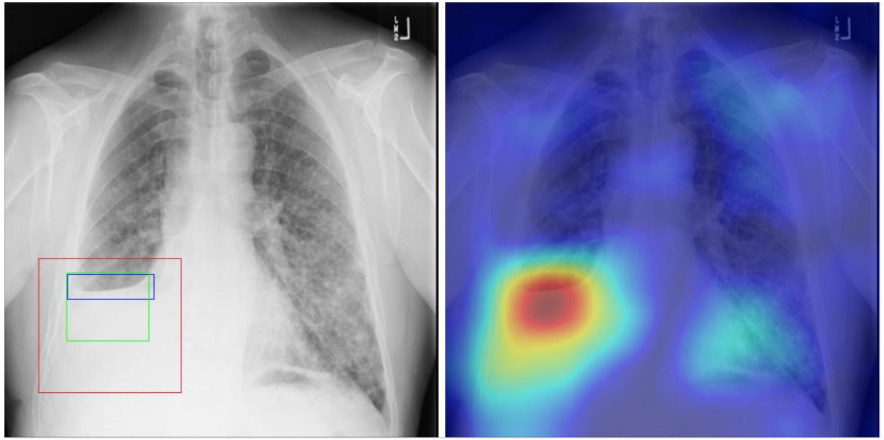

Table 10. *A sample of chest x-ray radiology report, mined disease keywords and localization result from the "Effusion" Class. Correct bounding box (in green), false positives (in red) and the ground truth (in blue) are plotted over the original image.*


Deep Learning
STANLEY
51-072 20oz
STANLEY

Cell counting from images

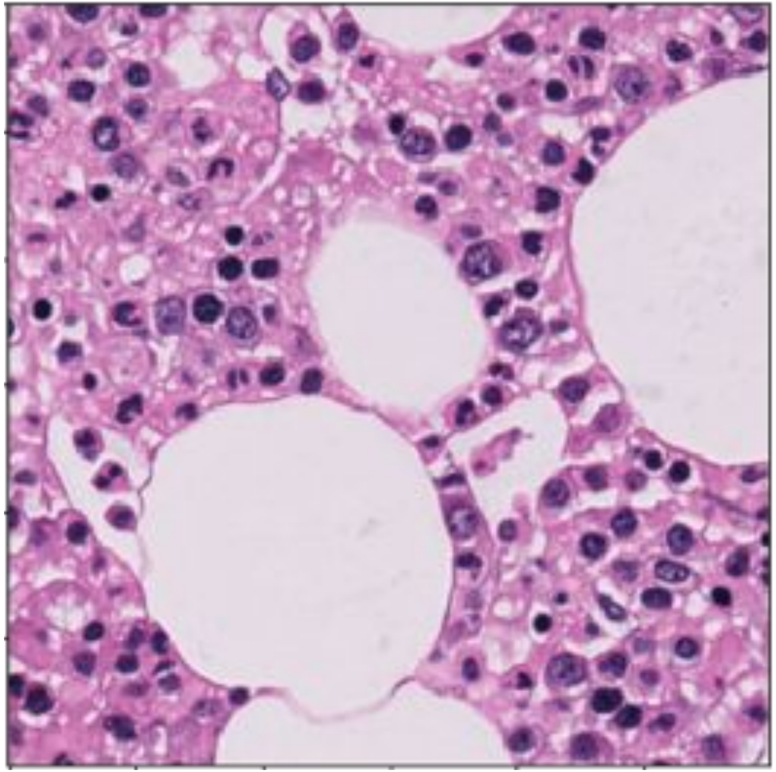

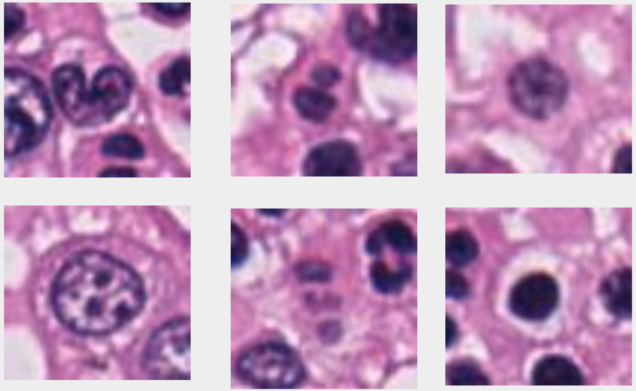

Complicated cell structure

[Cohen et al. 2017]

Cell counting from images

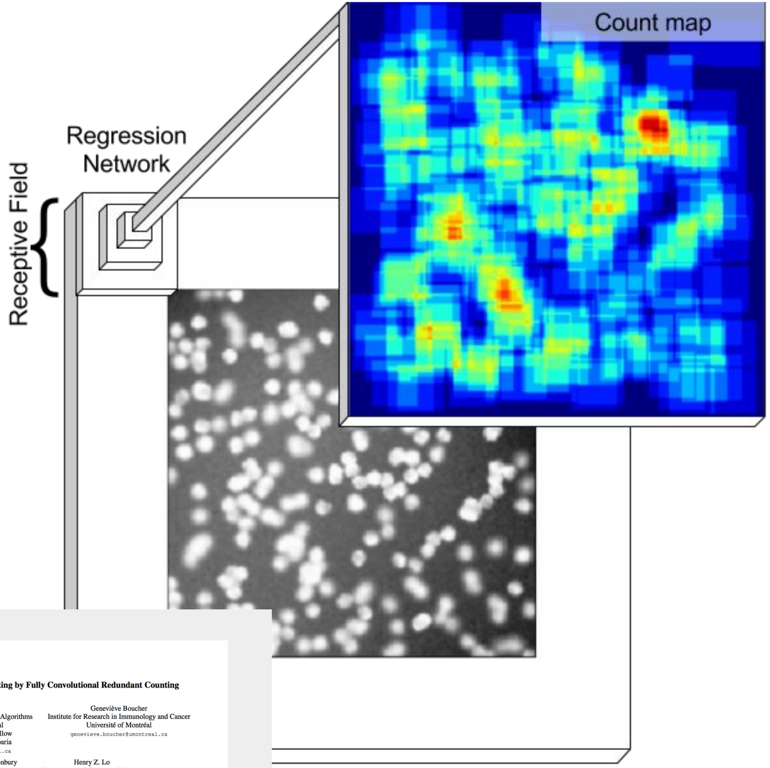

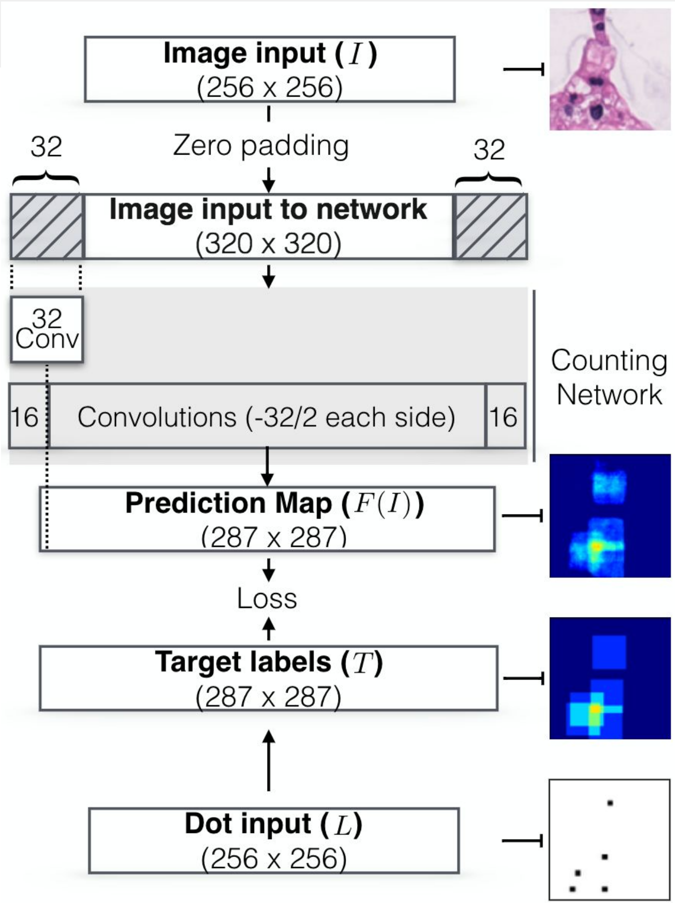

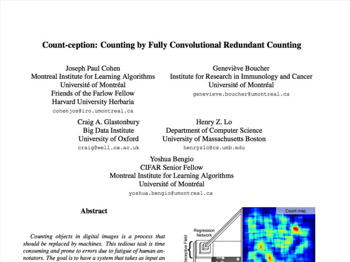

Count-ception: Counting by Fully Convolutional Redundant Counting

Joseph Paul Cohen
Montreal Institute for Learning Algorithms
Université of Montréal
Friends of the Farlow Fellow
Harvard University Herbaria
cohenjos@iro.umontreal.ca

Geneviève Boucher
Institute for Research in Immunology and Cancer
Université of Montréal
genevieve.boucher@umontreal.ca

Craig A. Glastonbury
Big Data Institute
University of Oxford
craig@well.ox.ac.uk

Henry Z. Lo
Department of Computer Science
University of Massachusetts Boston
henryzlo@cs.umb.edu

Yoshua Bengio
CIFAR Senior Fellow
Montreal Institute for Learning Algorithms
Université of Montréal
yoshua.bengio@umontreal.ca

Abstract

Counting objects in digital images is a process that should be replaced by machines. This tedious task is time consuming and prone to errors due to fatigue of human annotators. The goal is to have a system that takes as input an

Cohen et al., "Count-ception: Counting by Fully Convolutional Redundant Counting," 2017

Cell counting from images

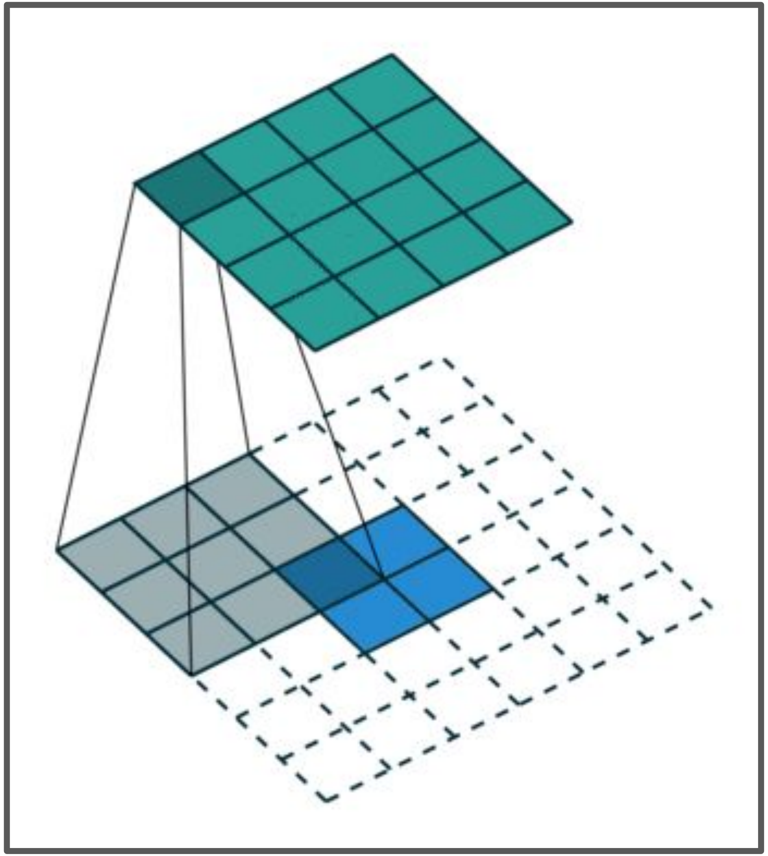

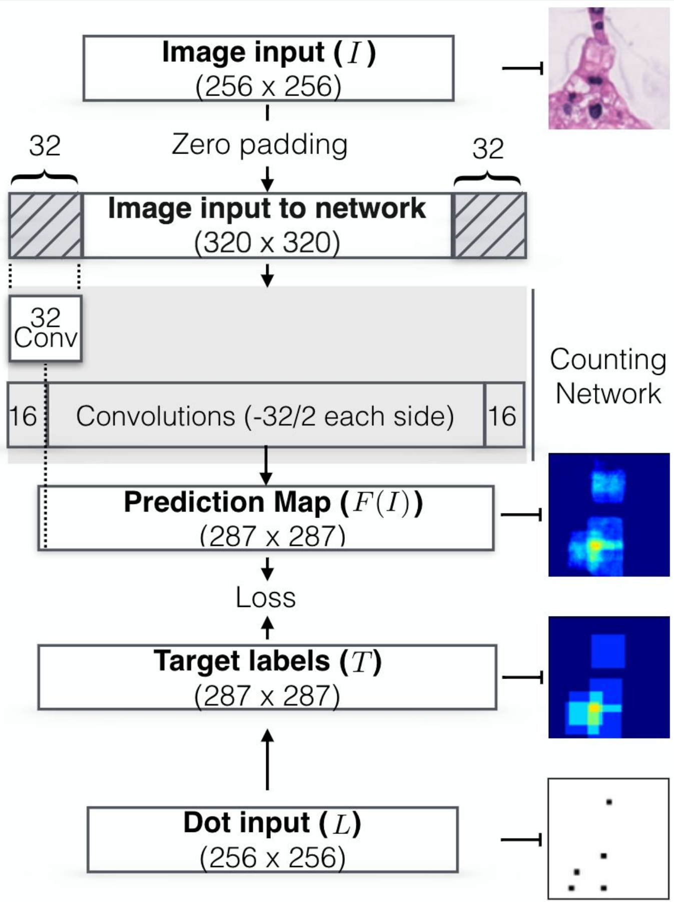

Cohen et al., "Count-ception: Counting by Fully Convolutional Redundant Counting," 2017

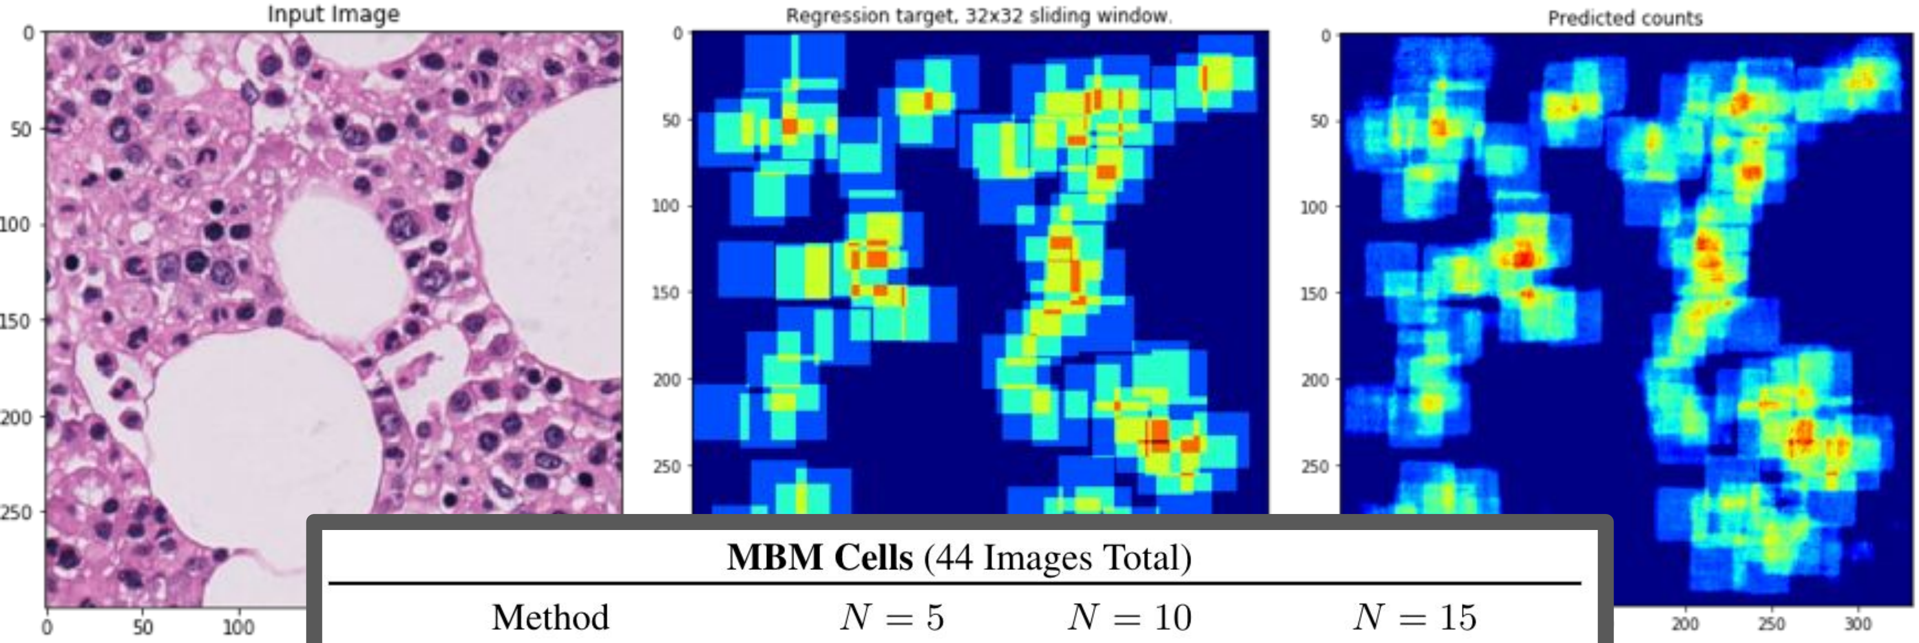

MBM Cells (44 Images Total)

Method	$N = 5$	$N = 10$	$N = 15$
Predict Average Count	29.4 ± 2.3	28.6 ± 1.6	28.2 ± 1.6
Cell Profiler -single		— 19.8 ± 4.2 —	
Cell Profiler -multiple**		— 12.8 ± 3.1 —	
FCRN-A, Xie (2016)	28.9 ± 22.6	22.2 ± 11.6	21.3 ± 9.4
Count-ception (Proposed)	12.6 ± 3.0	10.7 ± 2.5	$\boxed{8.8 \pm 2.3}$

**Cell Profiler results were obtained using a single pipeline (single) and using three different pipelines (multiple) to account for color differences in two of the eleven images.

[Cohen et al. 2017]

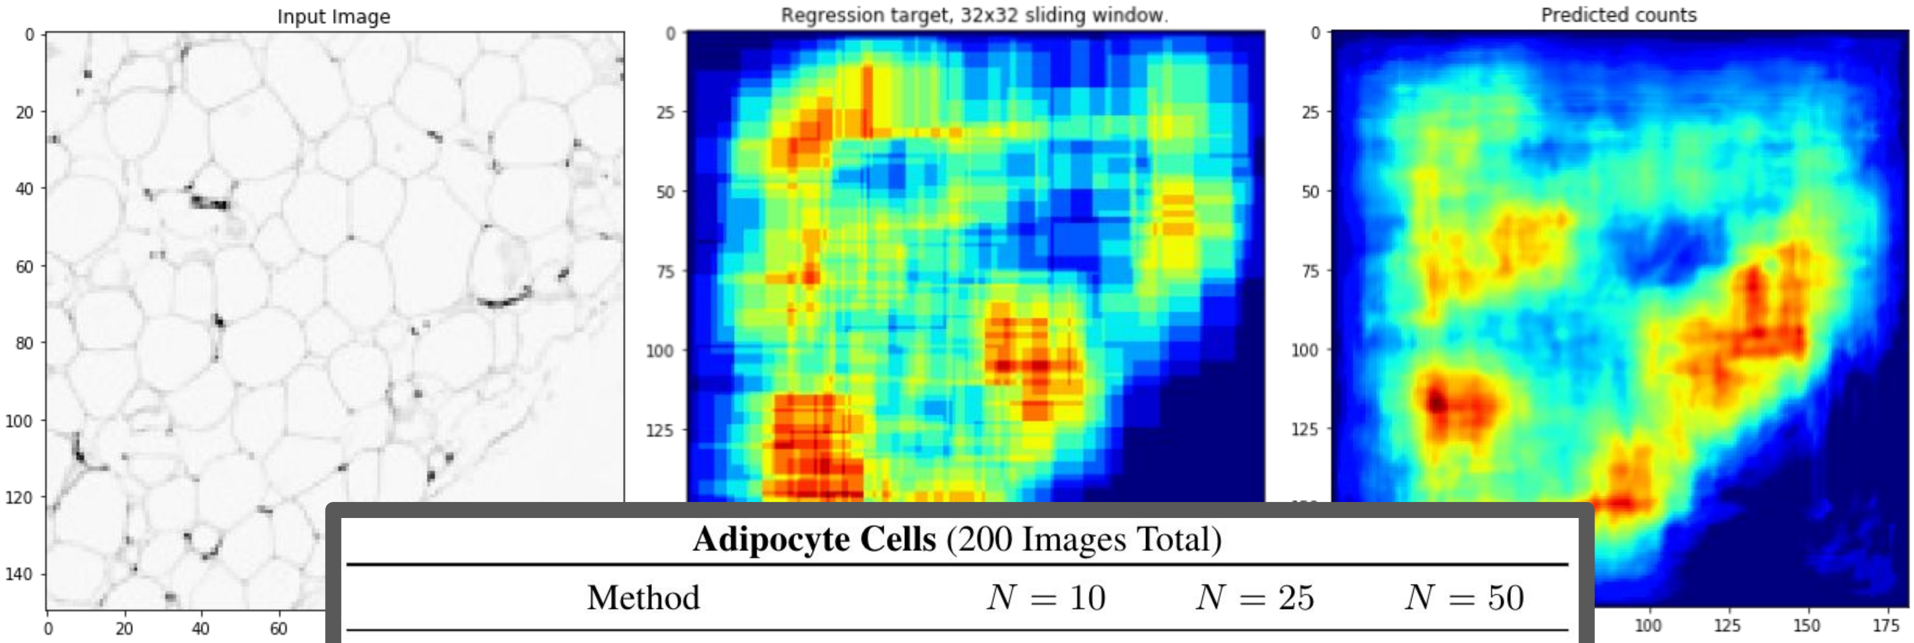

Adipocyte Cells (200 Images Total)			
Method	$N = 10$	$N = 25$	$N = 50$
Predict Average Count	33.8 ± 3.1	33.6 ± 3.0	33.5 ± 2.9
Cell Profiler 150×150		— 40.0 ± 3.9 —	
Cell Profiler 299×299 †		— 25.06 ± 2.6 —	
Adiposoft (2012) 299×299 † ‡		— 45.4 ± 25.47 —	
Adiposoft (2012) 1700×1700 † ‡		— 14.8 ± 13.63 —	
Count-ception (Proposed) 150×150	25.1 ± 2.9	21.9 ± 2.8	$\boxed{19.4 \pm 2.2}$

† This evaluation utilized larger resolution images than then Count-ception evaluation.
‡ This analysis was computed on only 10 images because it must be manually processed.

[Cohen et al. 2017]

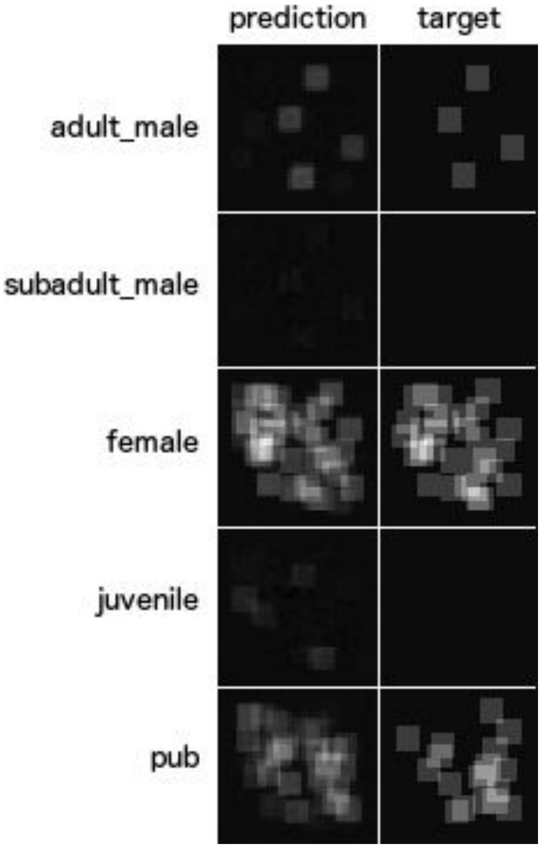

Kaggle sea lion challenge (37th/385 place)
Implemented by Robin Dinse (Universität Koblenz-Landau)


STANLEY
51-072 20oz
Deep Learning
STANLEY

Semi-supervised Segmentation with GANs

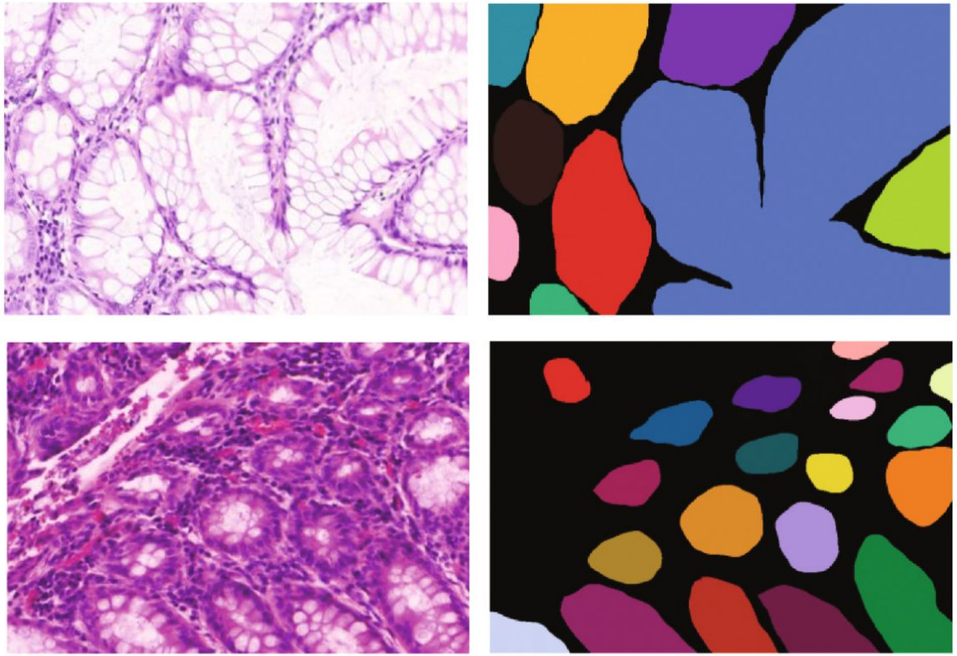

Images with segmentation labels

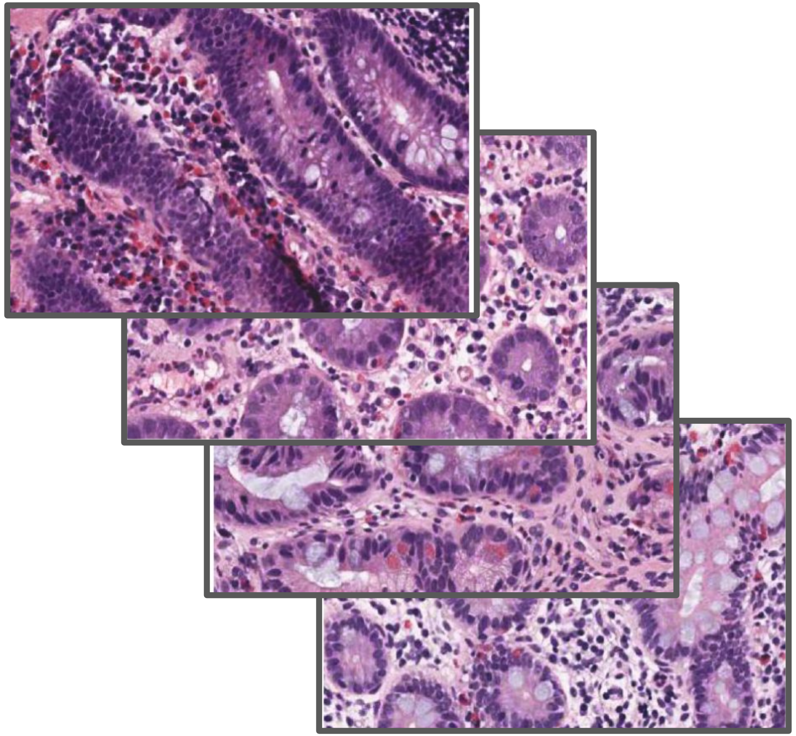

Images without segmentation labels

Semi-supervised Segmentation with GANs

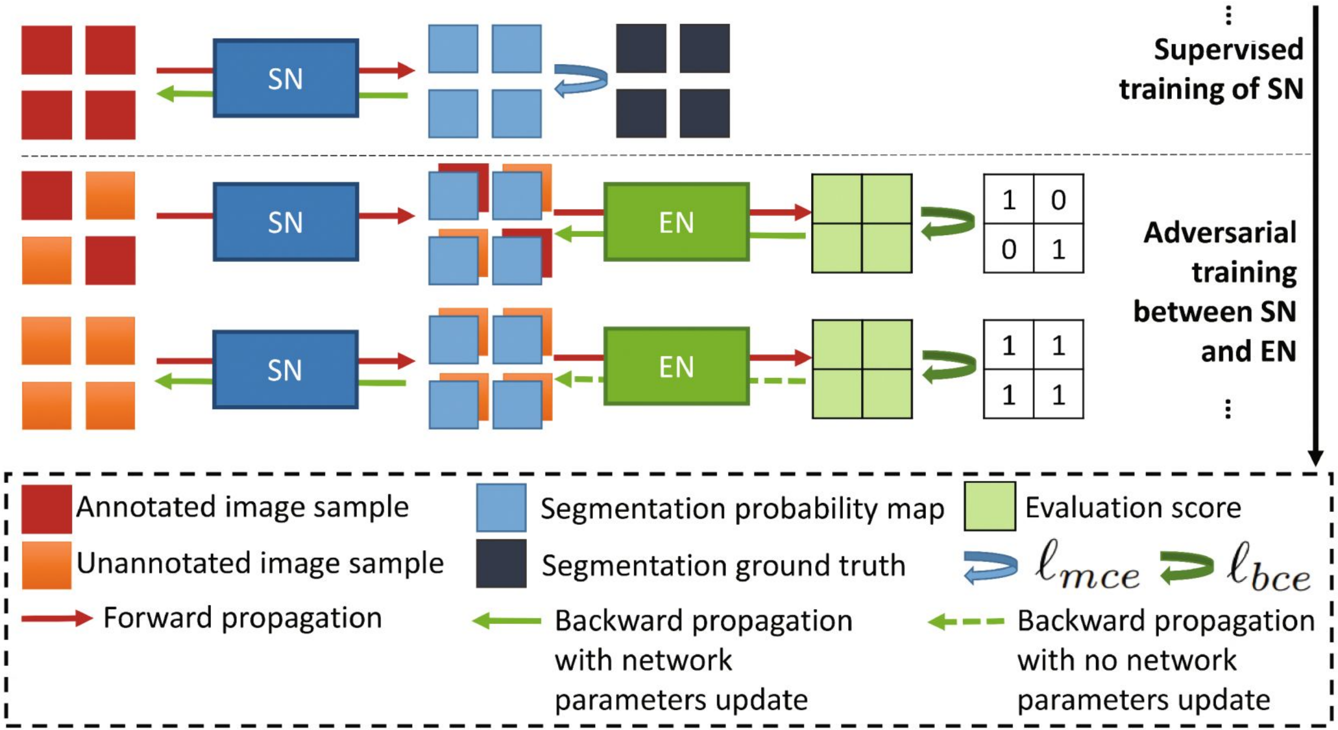

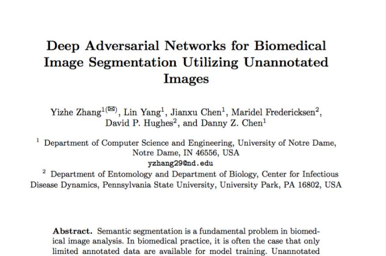

Deep Adversarial Networks for Biomedical Image Segmentation Utilizing Unannotated Images

Yizhe Zhang[1(✉)], Lin Yang[1], Jianxu Chen[1], Maridel Fredericksen[2], David P. Hughes[2], and Danny Z. Chen[1]

[1] Department of Computer Science and Engineering, University of Notre Dame, Notre Dame, IN 46556, USA
yzhang29@nd.edu
[2] Department of Entomology and Department of Biology, Center for Infectious Disease Dynamics, Pennsylvania State University, University Park, PA 16802, USA

Abstract. Semantic segmentation is a fundamental problem in biomedical image analysis. In biomedical practice, it is often the case that only limited annotated data are available for model training. Unannotated

Zhang et al., "Deep Adversarial Networks for Biomedical Image Segmentation Utilizing Unannotated Images," 2017

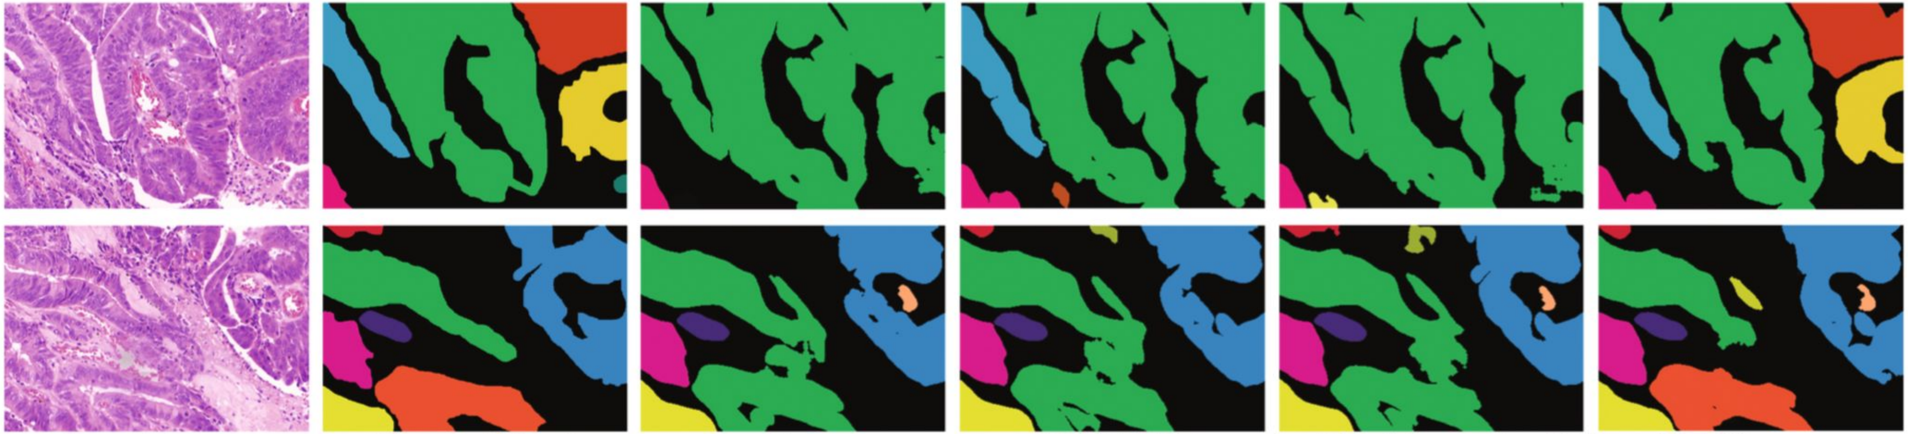

Input Image Ground Truth SN SSAN [6] Ladder Networks [9] DAN

Method	# images used		F_1 Socre		ObjectDice		ObjectHausdorff	
	Anno.	Unanno.	part A	part B	part A	part B	part A	part B
SN (base model)	85	0	0.9071	0.825	0.898	0.826	48.740	126.479
SSAN [6]	85	0	0.9060	0.836	0.886	0.818	53.393	128.385
Ladder networks [9]	85	100	0.9047	0.833	0.893	0.818	45.418	110.984
CUMedVision [2]	85	0	0.912	0.716	0.897	0.718	45.418	160.347
Multichannel1 [16]	85	0	0.858	0.771	0.888	0.815	54.202	129.930
Multichannel2 [15]	85	0	0.893	0.843	**0.908**	0.833	**44.129**	116.821
DAN	85	100	**0.916**	**0.855**	0.903	**0.838**	45.276	**104.982**


STANLEY
51-072 20oz
Deep Learning
STANLEY

Genomics - Gene Understanding

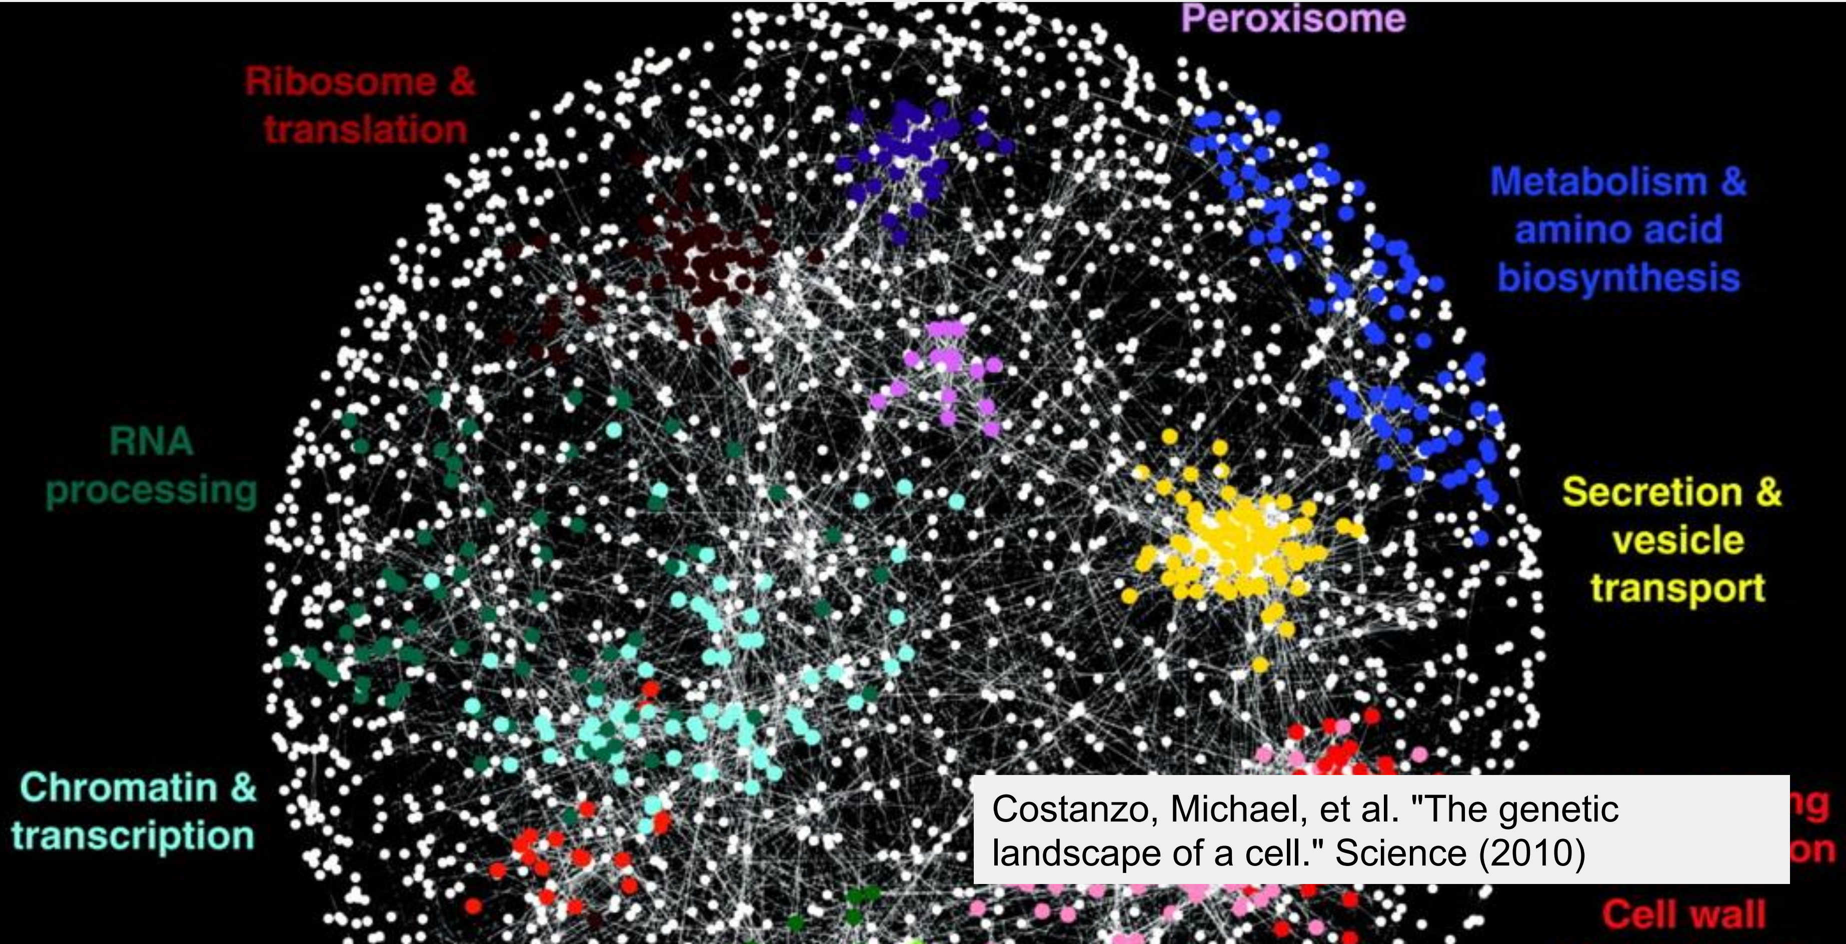

Costanzo, Michael, et al. "The genetic landscape of a cell." Science (2010)

Gene understanding

How can we represent different tissue types?

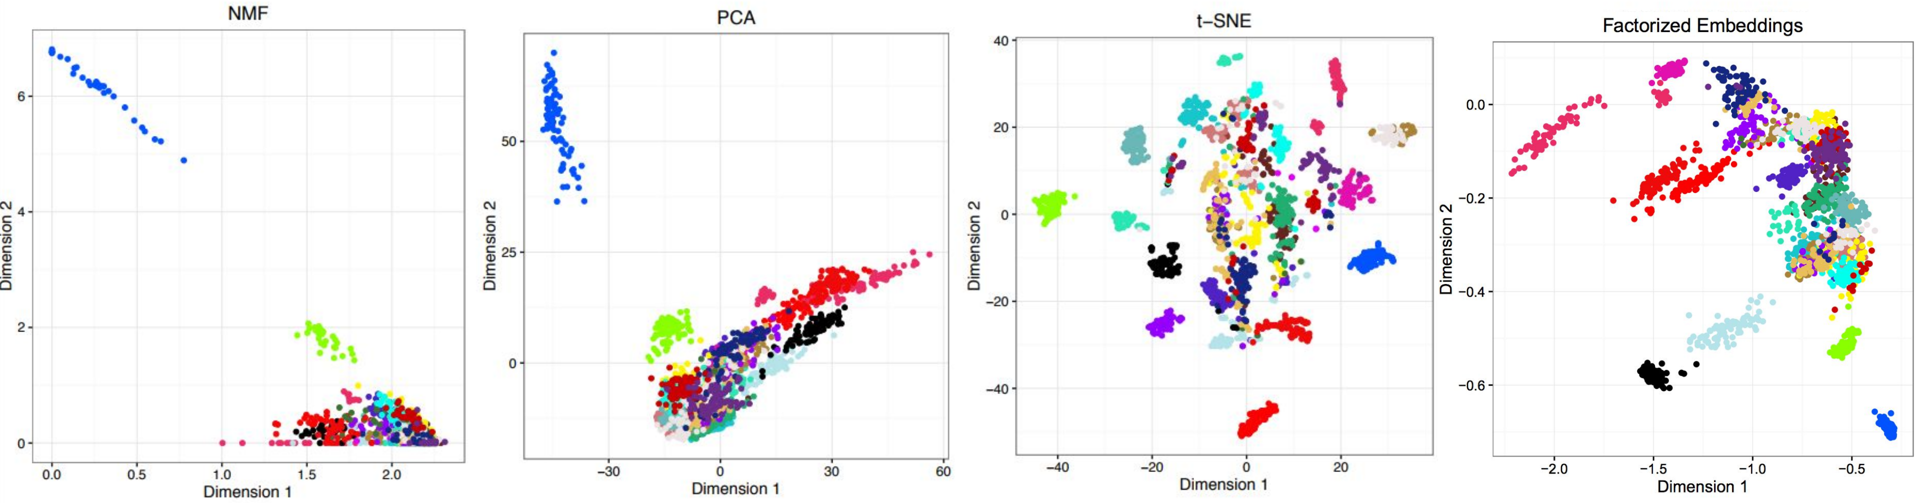

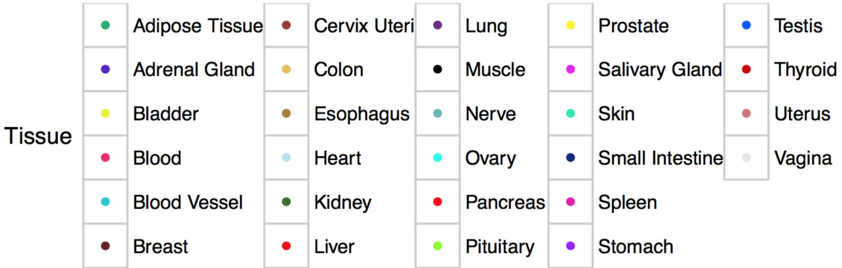

[Trofimov et al. ICML WCB 2017]

Factorized Embeddings

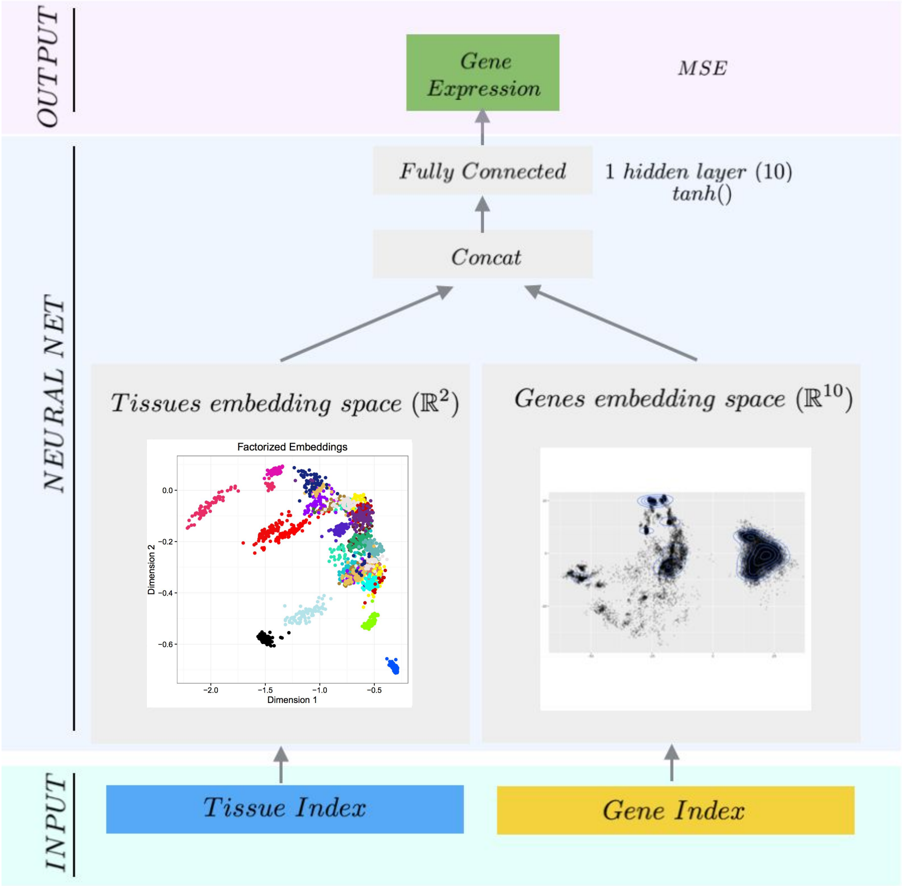

Gene expression is predicted given (Tissue, Gene) pair

Genes condition Tissue Prediction
Tissues condition Gene Prediction

Embedding locations are updated based on function

[Trofimov et al. ICML WCB 2017]

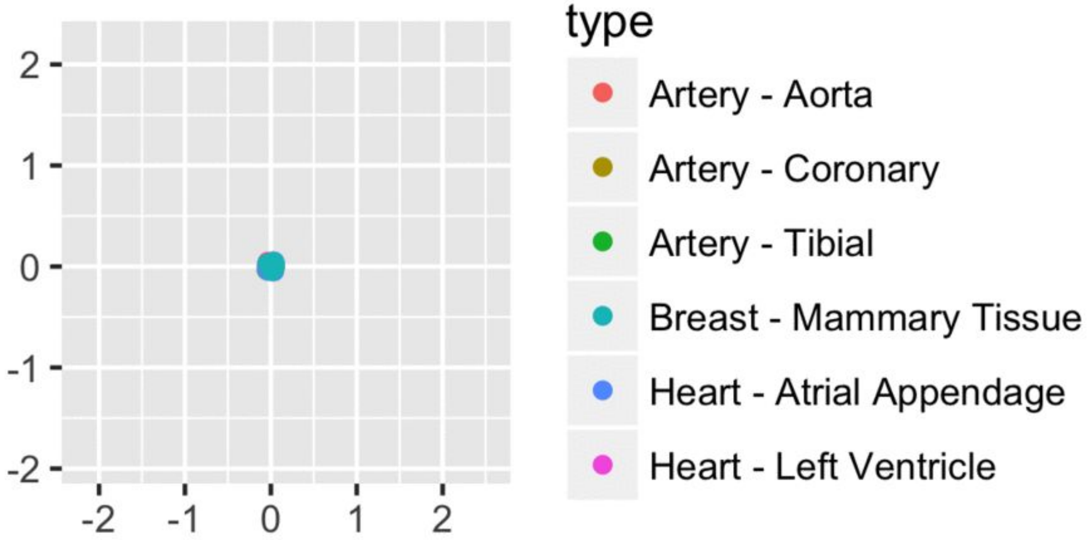
2
1
0
-1
-2
-2
-1
0
1
2
type
Artery - Aorta
Artery - Coronary
Artery - Tibial
Breast - Mammary Tissue
Heart - Atrial Appendage
Heart - Left Ventricle

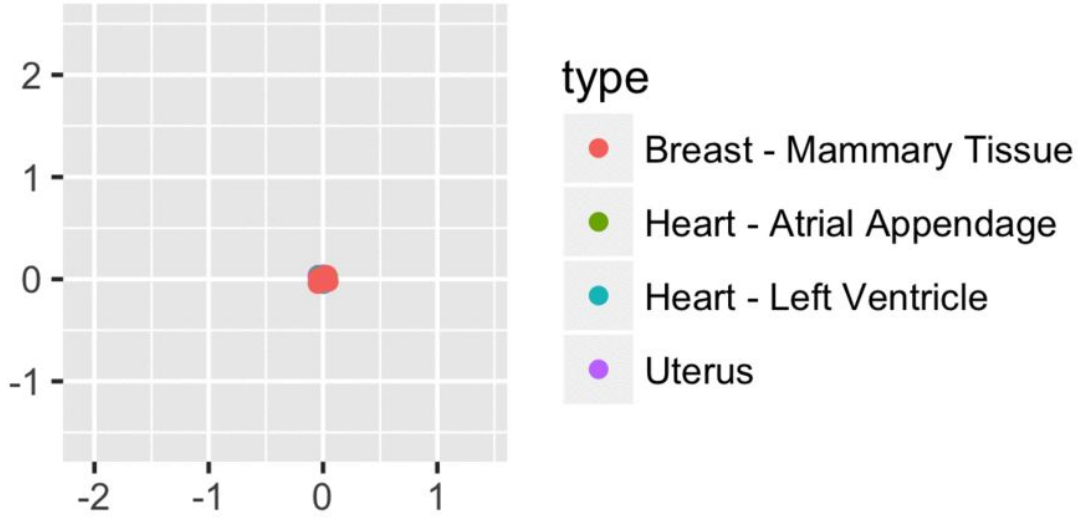
2
1
0
-1
-2
-1
0
1
type
Breast - Mammary Tissue
Heart - Atrial Appendage
Heart - Left Ventricle
Uterus

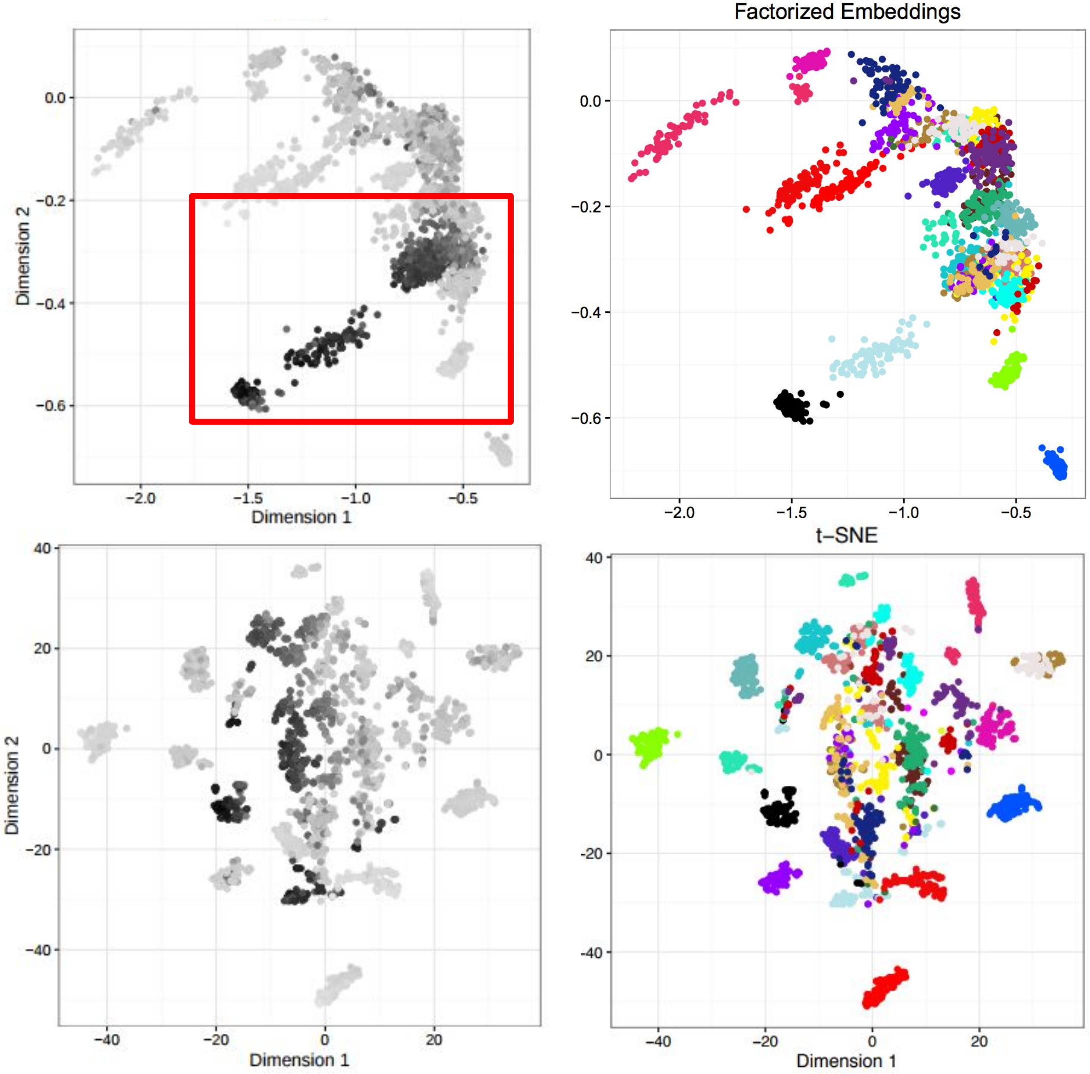

We can visualize an expression level for every tissue embedding

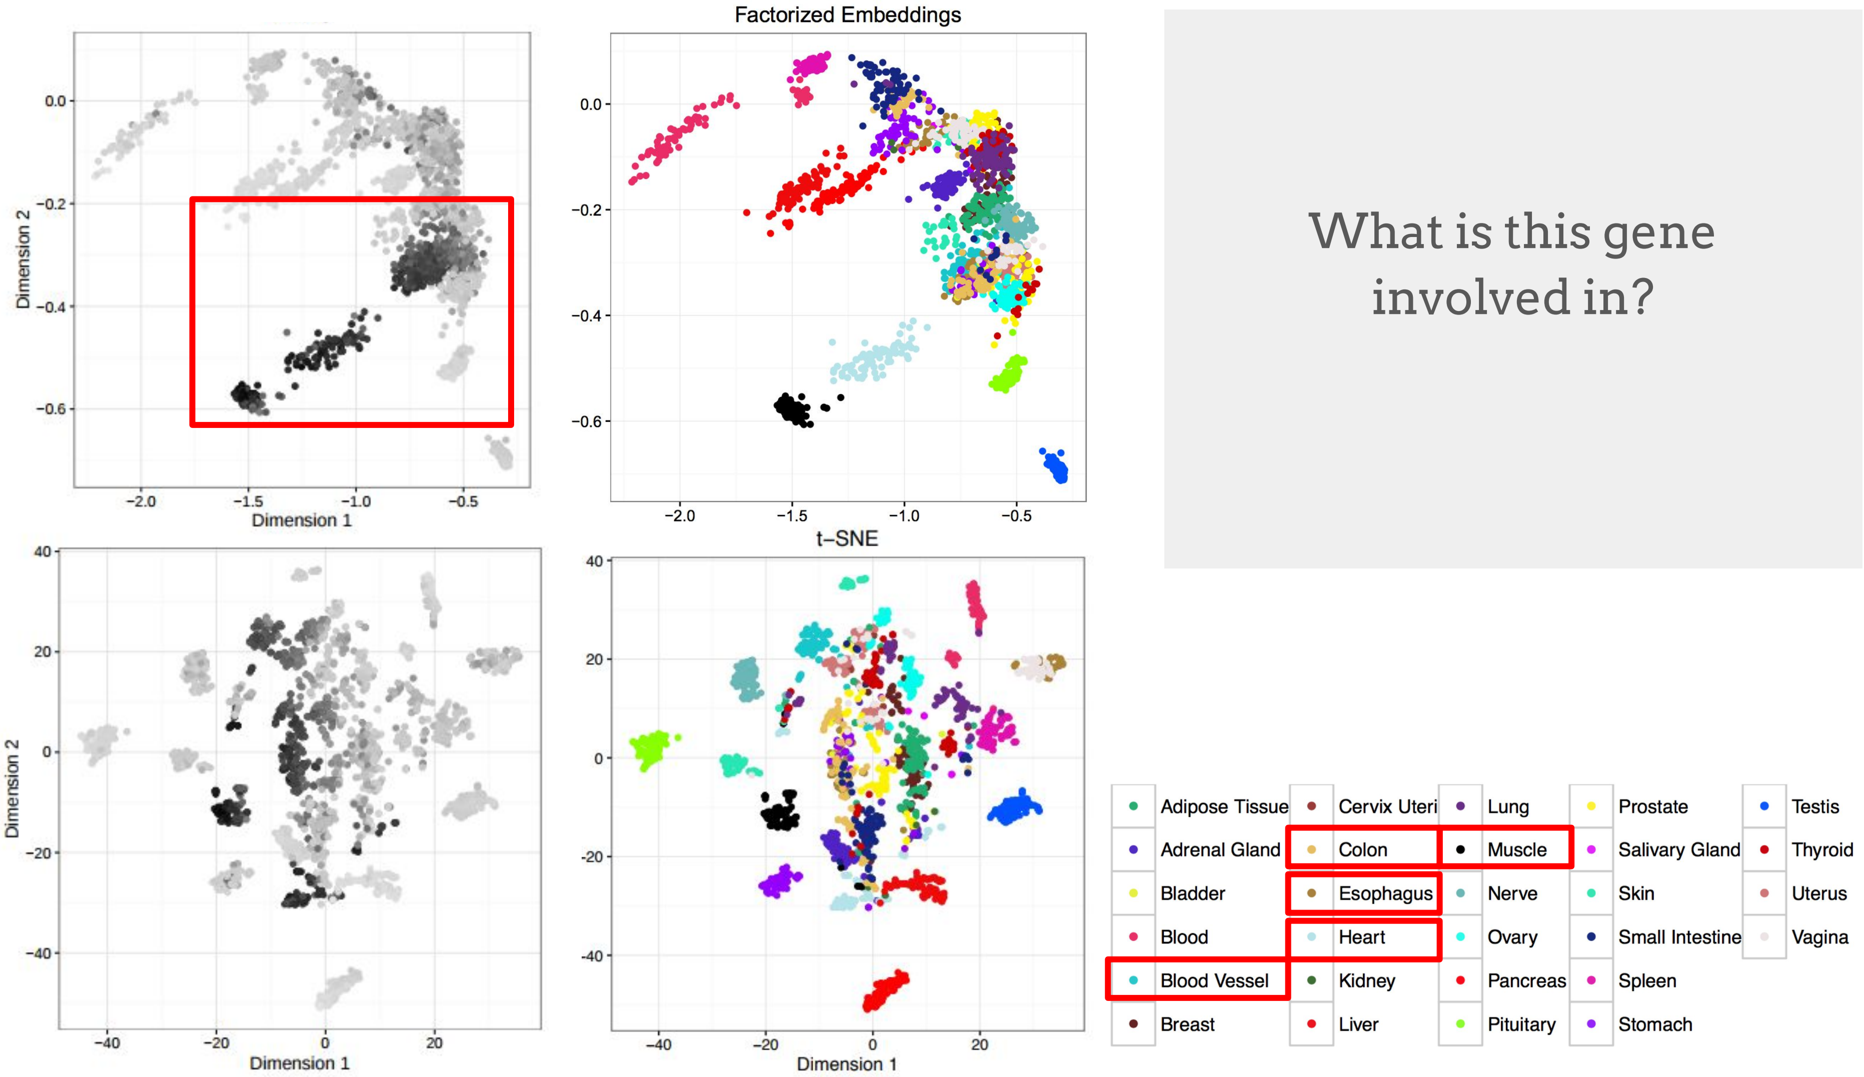

What is this gene involved in?

MYH6

"Muscle contraction"
[uniprot.org]

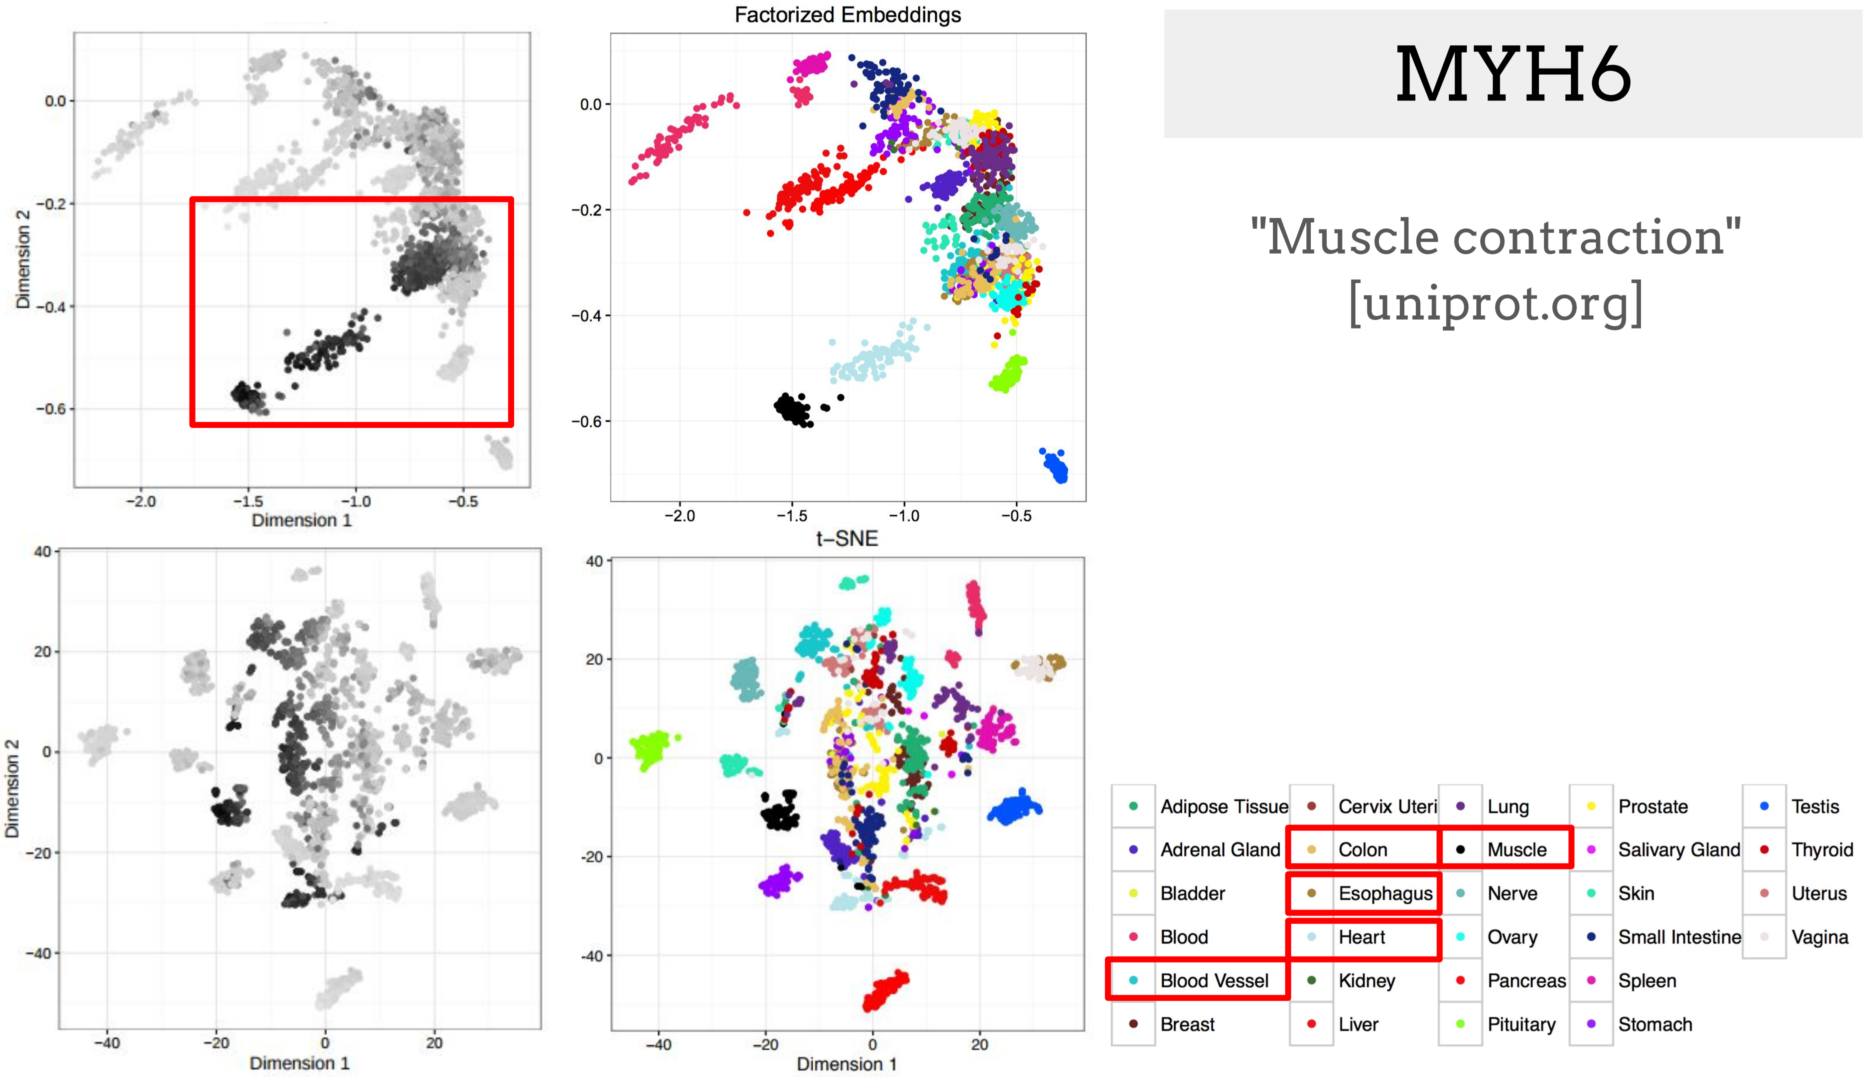

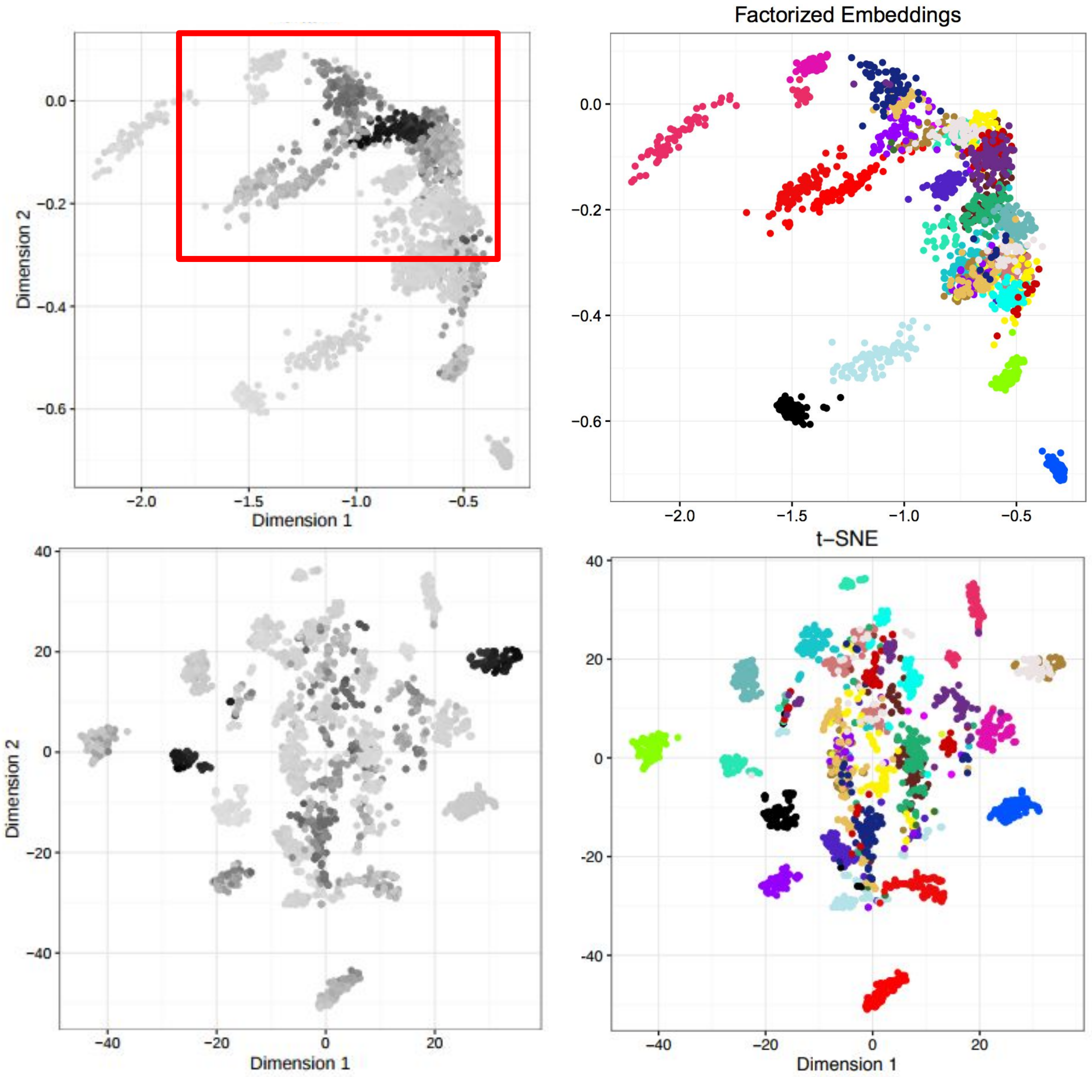

What is this gene involved in?

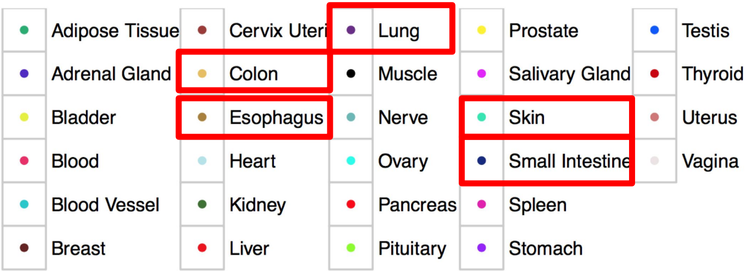

KRT (Keratin)

Keratin is the protein that protects epithelial cells from damage or stress

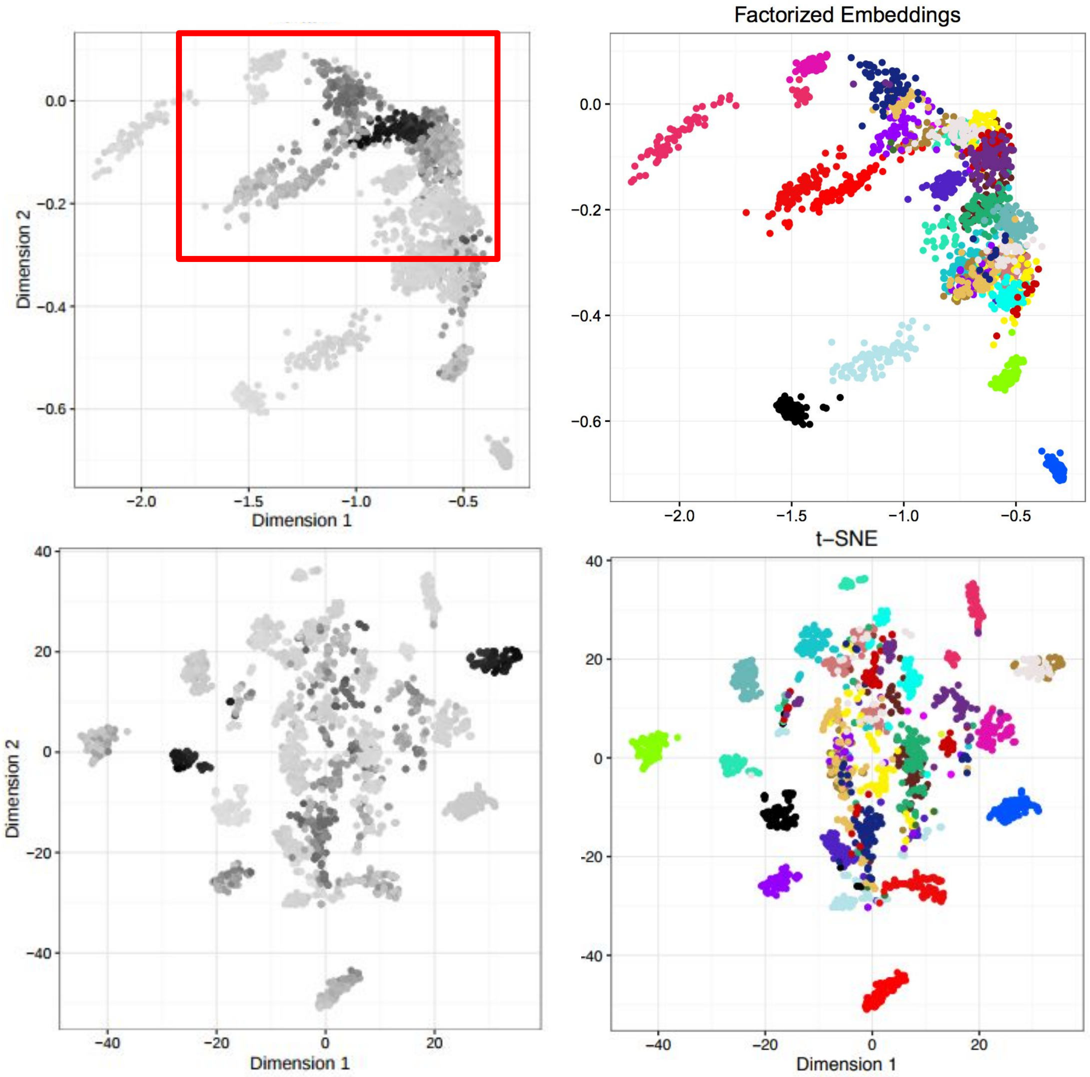

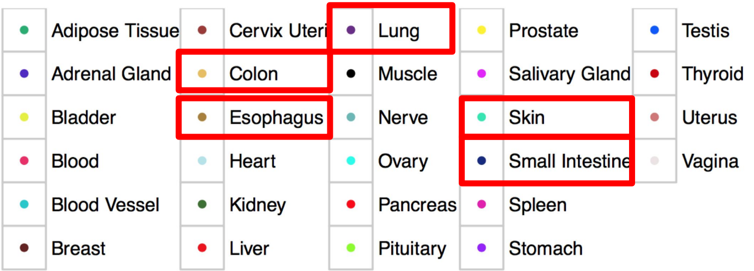

Thanks!

Conferences $ar\chi iv$ ShortScience.org

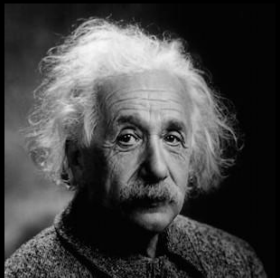

ShortScience.org