JHU vision lab

Artificial Intelligence and Machine Learning in Biomedicine and Health Care

René Vidal

Herschel L. Seder Professor of Biomedical Engineering

Director, Mathematical Institute for Data Science

Johns Hopkins University

JOHNS HOPKINS MATHEMATICAL INSTITUTE for DATA SCIENCE

THE DEPARTMENT OF BIOMEDICAL ENGINEERING

The Whitaker Institute at Johns Hopkins

Center for IMAGING SCIENCE

Brief History of Artificial Intelligence

Beginnings

- Thresholded Logic Unit — 1943
- Perceptron — 1957
- Adaline — 1960

1st Neural Winter

- XOR Problem — 1969

Renaissance

- Multilayer Backprop — 1982, 1986
- CNNs — 1989
- LSTMs — 1997

2nd Neural Winter

- SVMs — 1995

Modern Era

- Deep Nets — 2006
- Alex Net — 2012

1940 | **1950** | **1960** | **1970** | **1980** | **1990** | **2000** | **2010**

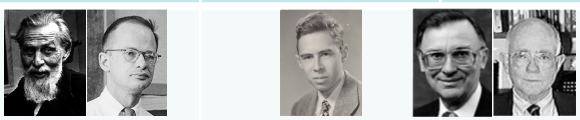

S. McCulloch - W. Pitts

R. Rosenblatt

B. Widrow - M. Hoff

M. Minsky - S. Papert

P. Werbos

D. Rumelhart - G. Hinton - R. Williams

Y. Lecun

J. Schmidhuber

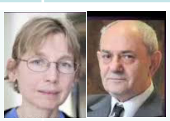

C. Cortes - V. Vapnik

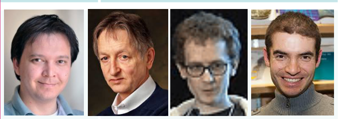

R. Salakhutdinov - J. Hinton - A. Krizhevsky - I. Sutskever

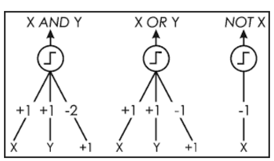

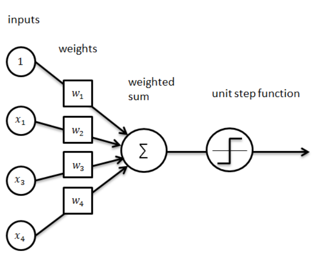

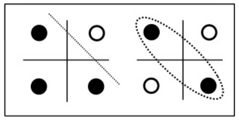

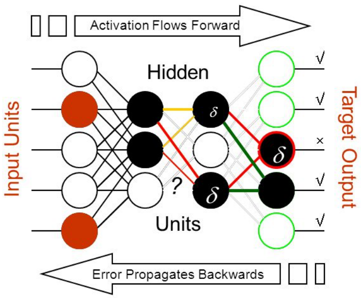

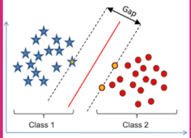

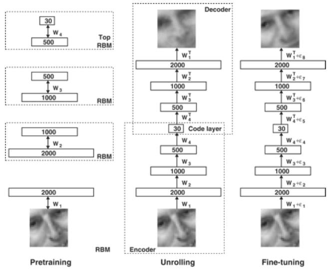

JOHNS HOPKINS
MATHEMATICAL INSTITUTE
for DATA SCIENCE

Center for IMAGING SCIENCE

Impact of AI/DL in Computer Vision

- 2012-2014 classification results in ImageNet

CNN
non-CNN

2012 Teams	%error
Supervision (Toronto)	15.3
ISI (Tokyo)	26.1
VGG (Oxford)	26.9
XRCE/INRIA	27.0
UvA (Amsterdam)	29.6
INRIA/LEAR	33.4

2013 Teams	%error
Clarifai (NYU spinoff)	11.7
NUS (singapore)	12.9
Zeiler-Fergus (NYU)	13.5
A. Howard	13.5
OverFeat (NYU)	14.1
UvA (Amsterdam)	14.2
Adobe	15.2
VGG (Oxford)	15.2
VGG (Oxford)	23.0

2014 Teams	%error
GoogLeNet	6.6
VGG (Oxford)	7.3
MSRA	8.0
A. Howard	8.1
DeeperVision	9.5
NUS-BST	9.7
TTIC-ECP	10.2
XYZ	11.2
UvA	12.1

- 2015 results: ResNet under 3.5% error using 150 layers!

Slide from Yann LeCun's CVPR'15 plenary and ICCV'15 tutorial intro by Joan Bruna

Impact of AI/DL in Speech Recognition

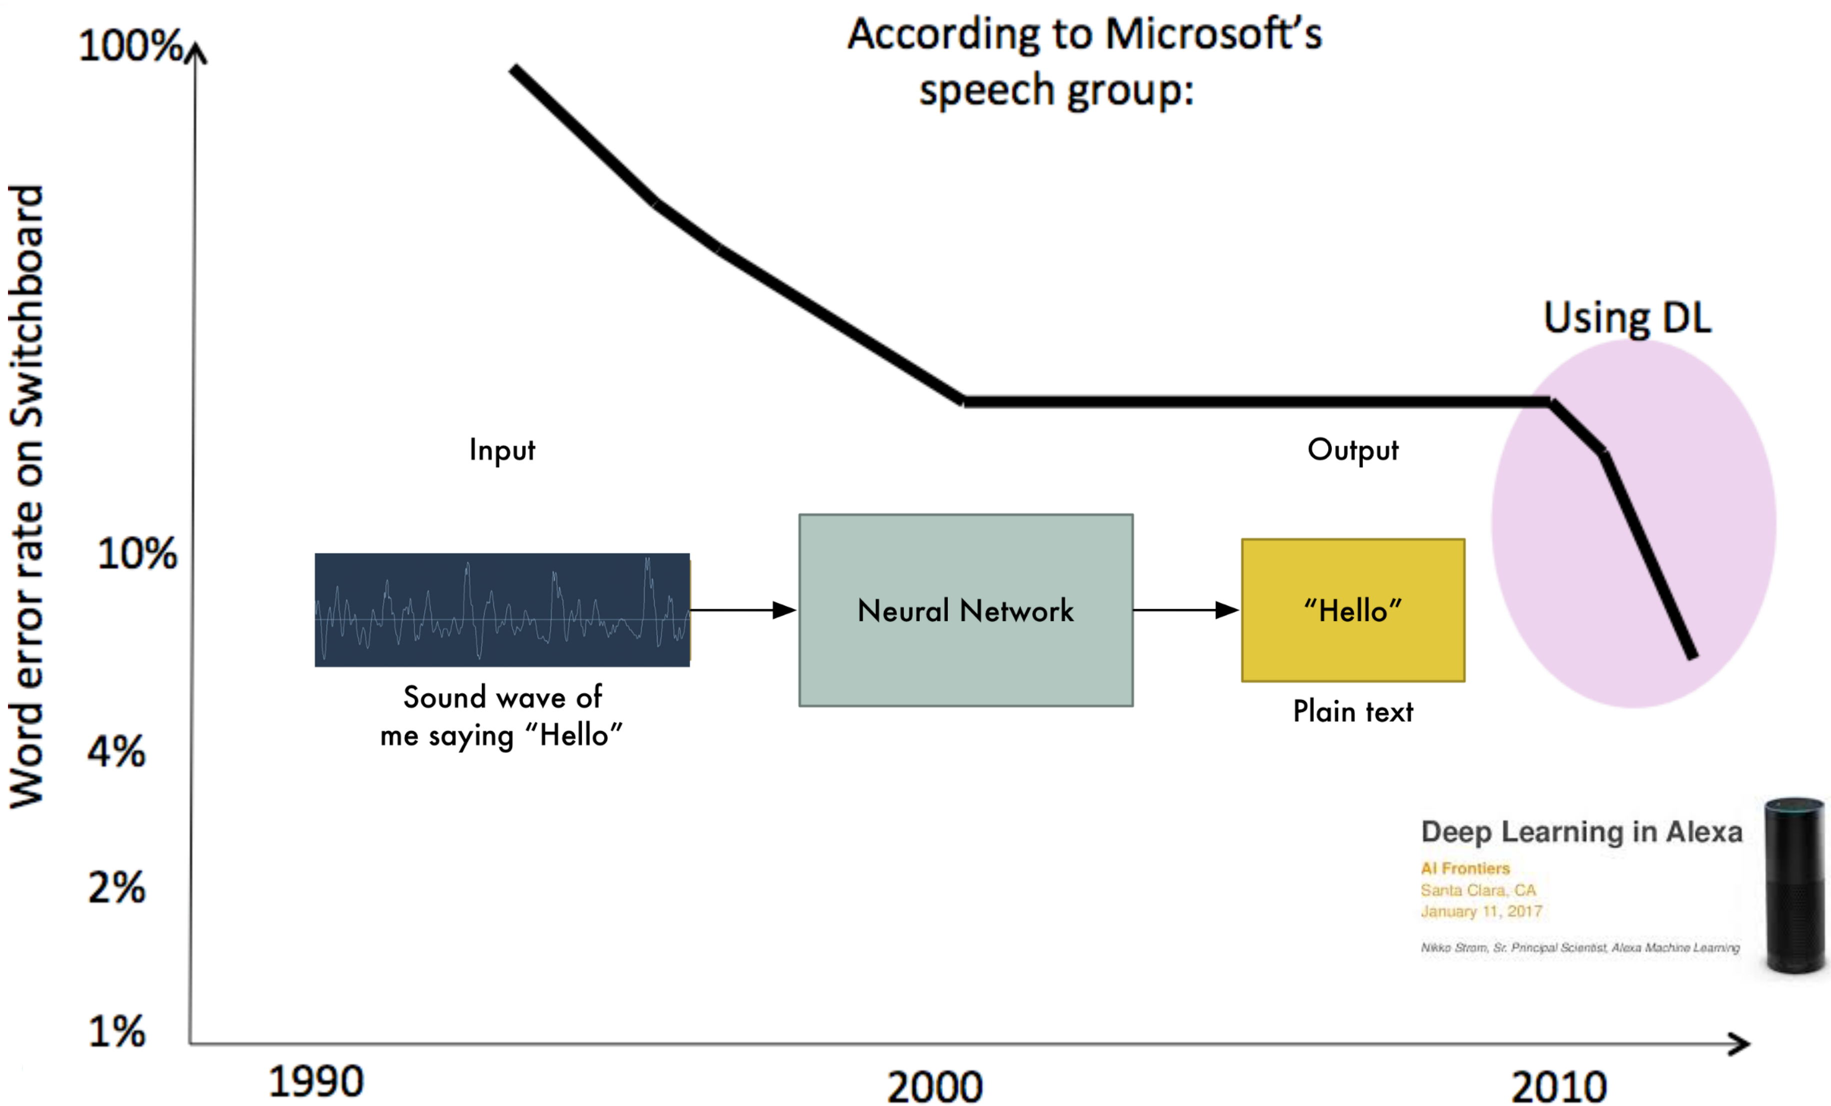

Impact of AI/DL in Game Playing

- **AlphaGo**: the first computer program to ever beat a professional player at the game of Go [1]

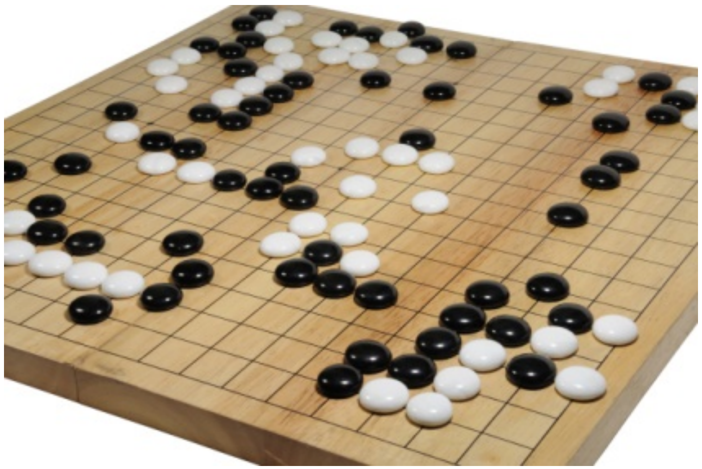

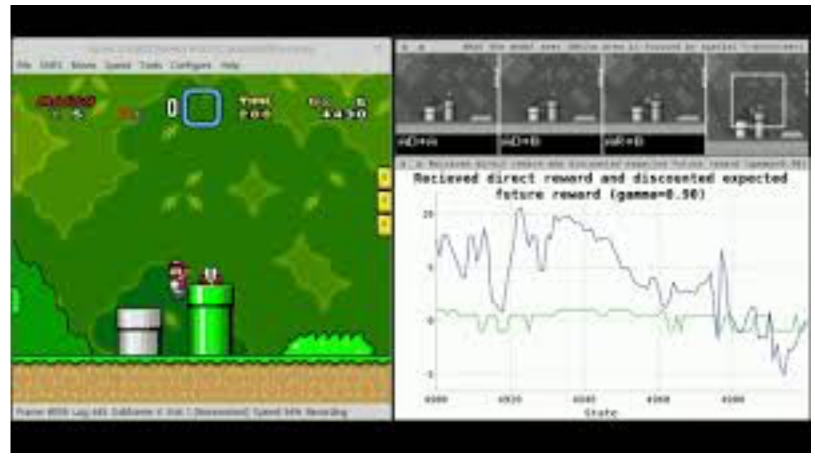

- Similar deep reinforcement learning strategies developed to play **Atari Breakout**, **Super Mario**

Silver et al. Mastering the game of Go with deep neural networks and tree search, Nature 2016
Artificial intelligence learns Mario level in just 34 attempts, https://www.engadget.com/2015/06/17/super-mario-world-self-learning-ai/, https://github.com/aleju/mario-ai

Potential Impact of AI/DL in Biomedicine

https://www.forbes.com/sites/robertpearl/2018/03/13/artificial-intelligence-in-healthcare/#470c8be1d750

Potential Impact of AI/DL in Biomedicine

Google DeepMind Health

IBM Watson Path

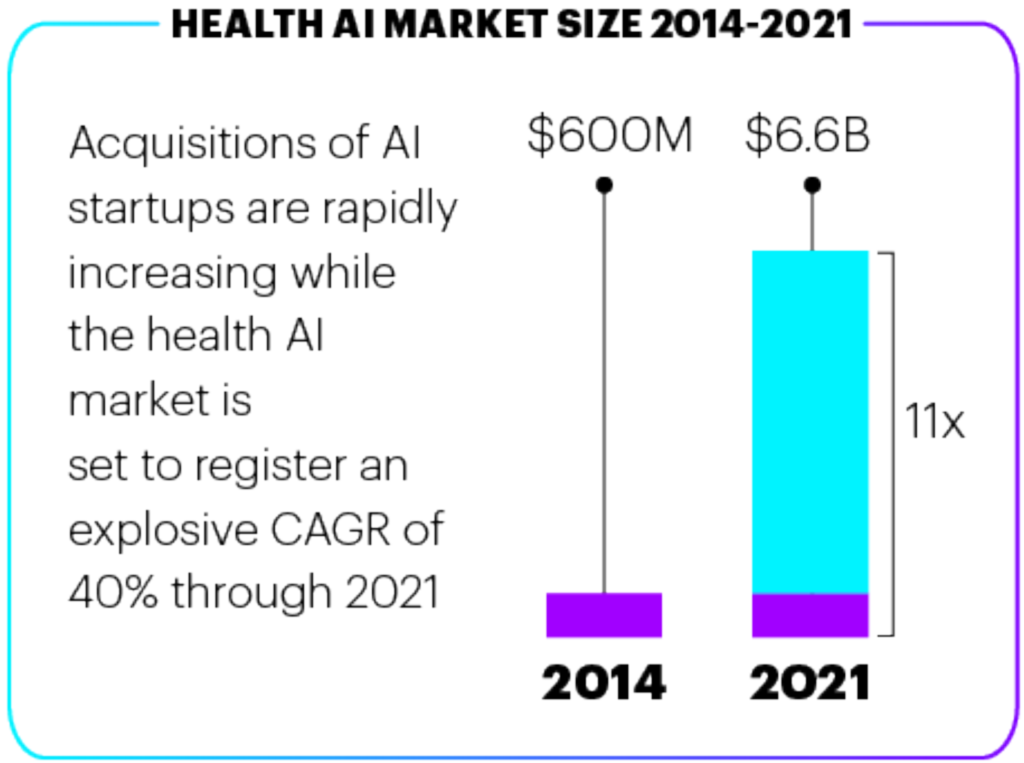

Source: Accenture analysis

https://www.accenture.com/us-en/insight-artificial-intelligence-healthcare
http://medicalfuturist.com/top-artificial-intelligence-companies-in-healthcare/

What is Biomedical Data Science About?

Data

EEG/ECG

Calcium imaging

D-MRI

Surgical data

Electronic health records

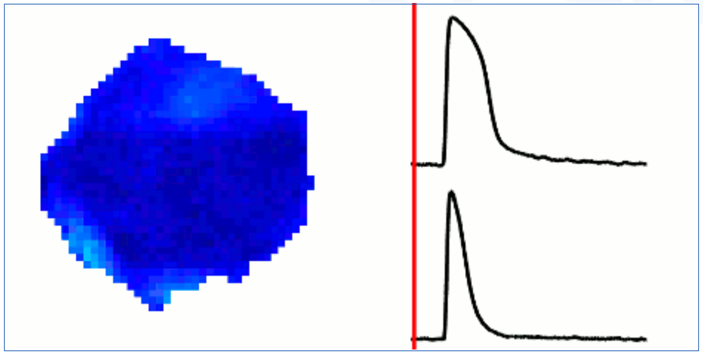

Machine Learning

Interpretation

Cell phenotype

Neural shape, spiking activity

Brain networks

Surgeon skill

Diagnosis and prognosis

Fundamental Challenges: Scale/Modalities

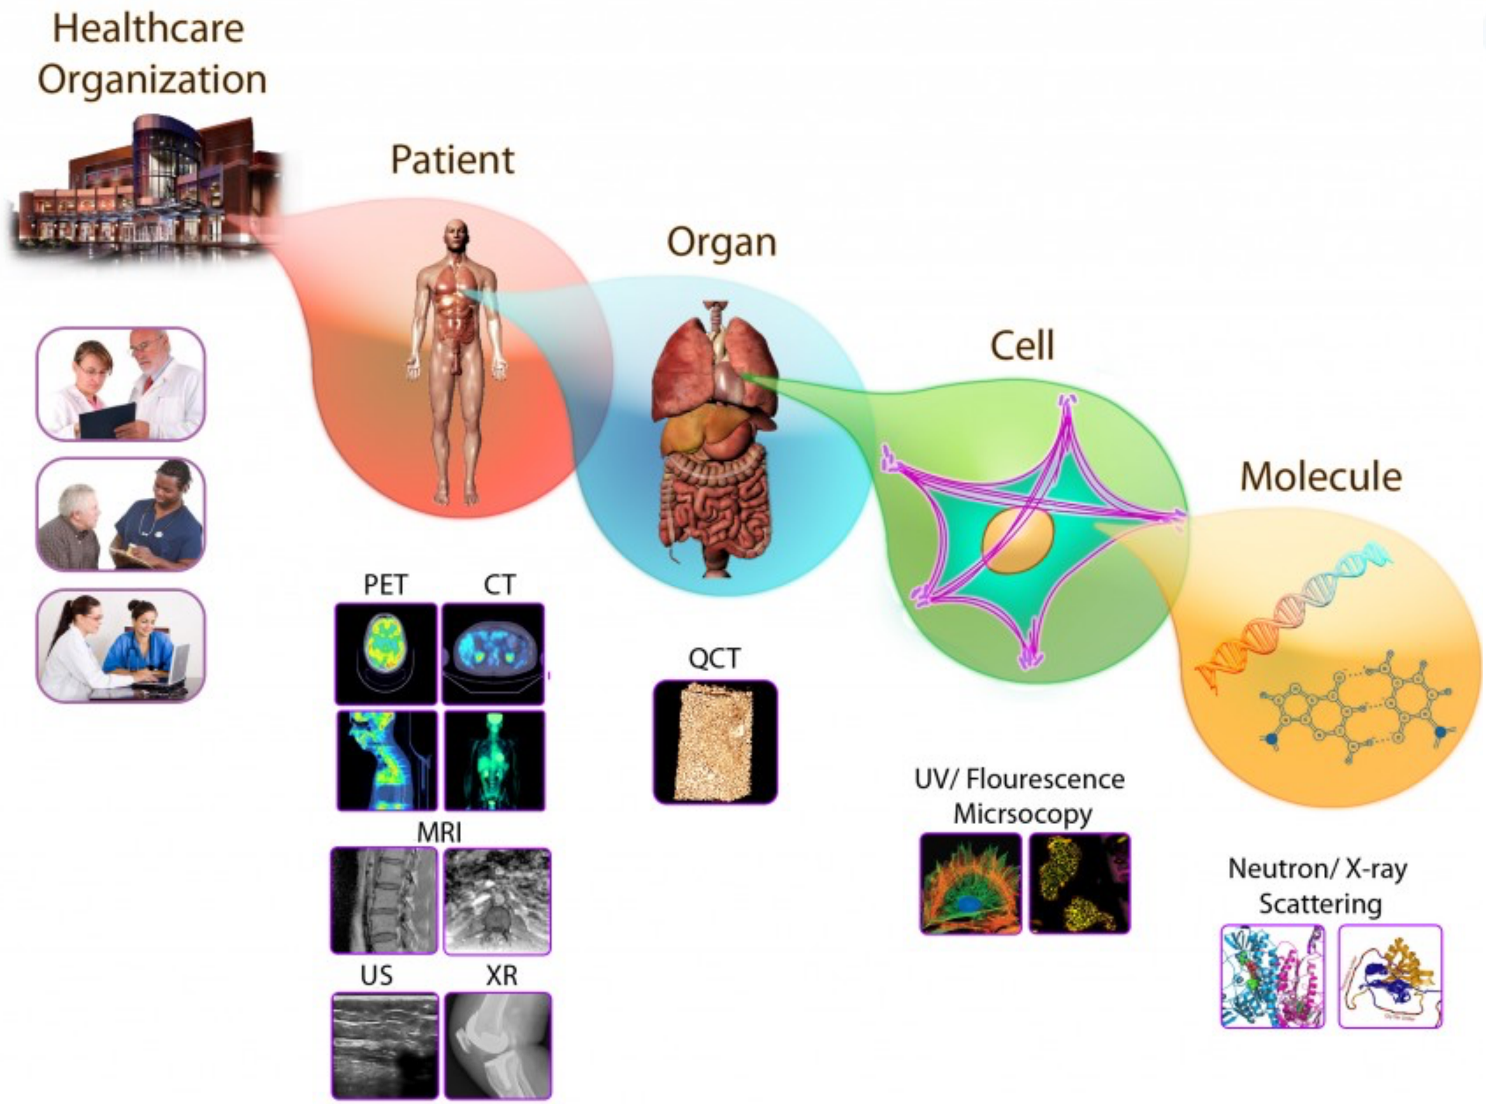

http://ibme.utk.edu/research/

Fundamental Challenges: Big/Complex Data

- 400 million procedures/year involve at least 1 medical image
- Medical image archives are increasing by 20-40 percent each year
- 1 billion medical images stored in the US (2012)
- 1/3 of global storage is medical image information

http://www.corp.att.com/healthcare/docs/medical_imaging_cloud.pdf

Fundamental Challenges: Small Annotations

- State of the art methods in ML are data hungry
  - Cleaning & annotation of biomedical data is very costly

- State of the art methods in ML are not fully interpretable

- Strong need for
  - Sharing clean, highly annotated data
  - Developing methods that require minimal supervision
  - Developing methods that are interpretable to physicians

Our Research in Biomedical Data Science

- **Cellular Level**
  - Dictionary learning for blood cell detection, classification, counting
  - Structured matrix factorization for segmentation of neural activity in calcium imaging
  - Metamorphosis for classification of embryonic cardio-myocytes

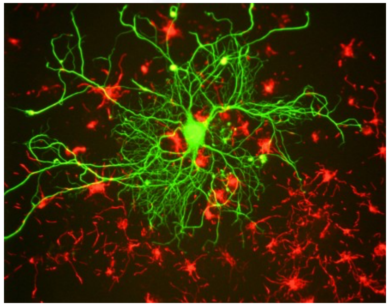

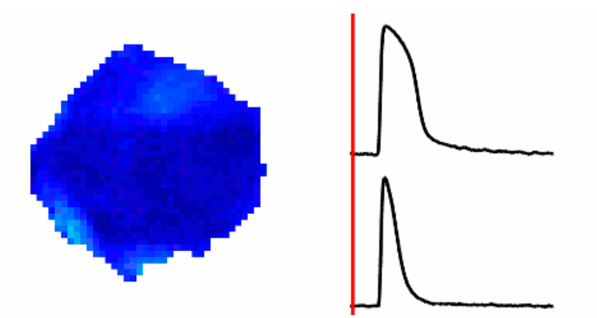

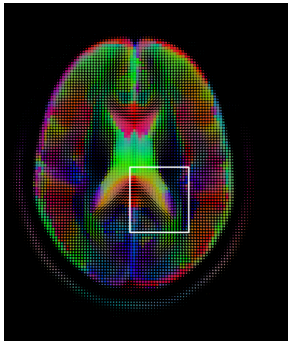

- **Organ Level**
  - Compressed sensing and Riemannian geometry for processing diffusion MRI

- **Patient/Surgeon Level**
  - Assessing surgical skill
  - Assessing children's motions

Ghoreyshi ISBI07; Singaraju CVPR08, CVPR09, TPAMI11; Goh ECCV06, ISBI06, ECCV08, CVPR09; MICCAI09, Neuroimage 12; Cetingul PPMIA09, ISBI09, TBME11, ISBI11, ISBI12, TBME14; Schwab IPMI13; Gorospe MICCAI13, TBME13; Haeffele ICML14; Tao IPCAI12, MICCAI13; Bejar-Zapella MICCAI12, Media13, MICCAI13; Lea: WACV15, ECCV16

JHU vision lab

Machine Learning in Hematology

René Vidal
Herschel L. Seder Professor of Biomedical Engineering
Director, Mathematical Institute for Data Science
Johns Hopkins University

Thermometer

Hematology Analyzer

Machine Learning in Hematology

- Complete Blood Count (CBC) using hematology analyzer

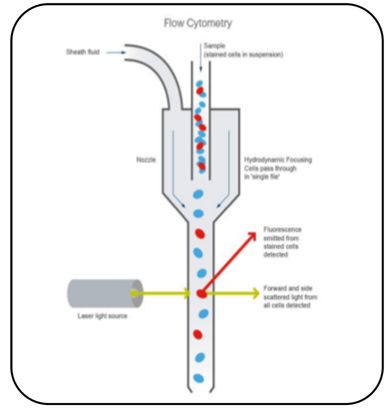

		Results
WBC		5.65×10^3 cells / µL
RBC	L	3.88×10^6 cells / µL
HGB		13.8 g / dL
HCT		41.3 %
MCV	H	106.5 fL
MCH	H	35.7 pg
MCHC		33.5 g / dL
CHCM		33.3 g / dL
CH		35.2 pg
RDW	H	15.5 %
HDW		2.74 g / dL
PLT	L	87×10^3 cells / µL
MPV		9.4 fL

- CBC using lens-free imaging

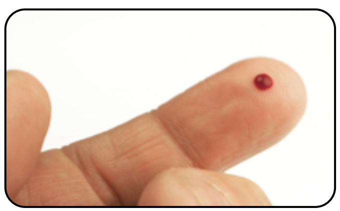

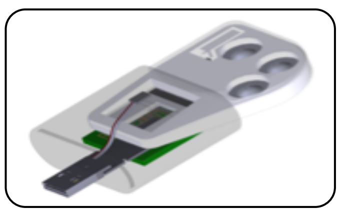

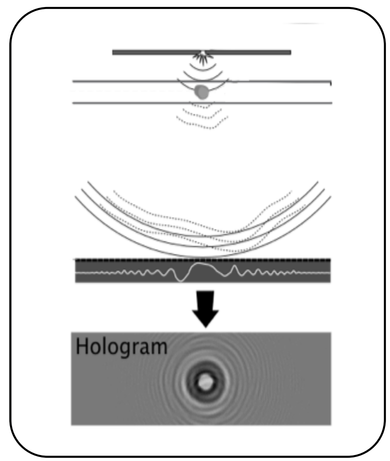

		Results
WBC		5.65×10^3 cells / µL
RBC	L	3.88×10^6 cells / µL
HGB		13.8 g / dL
HCT		41.3 %
MCV	H	106.5 fL
MCH	H	35.7 pg
MCHC		33.5 g / dL
CHCM		33.3 g / dL
CH		35.2 pg
RDW	H	15.5 %
HDW		2.74 g / dL
PLT	L	87×10^3 cells / µL
MPV		9.4 fL

Computer Vision for Blood Counting & Classification

Granulocyte

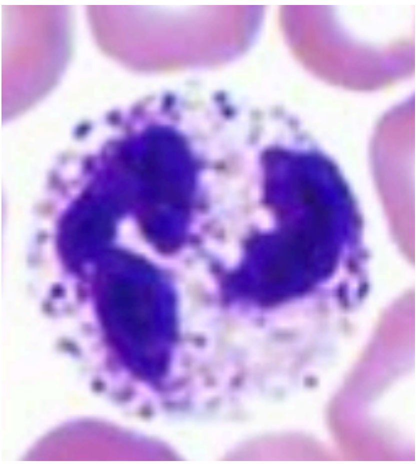

Lymphocyte

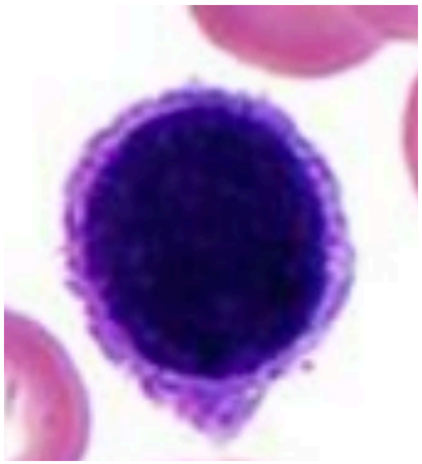

Monocyte

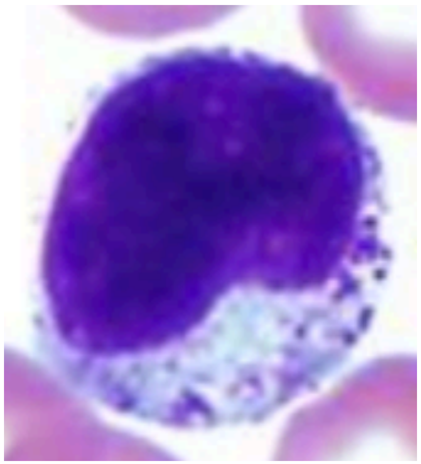

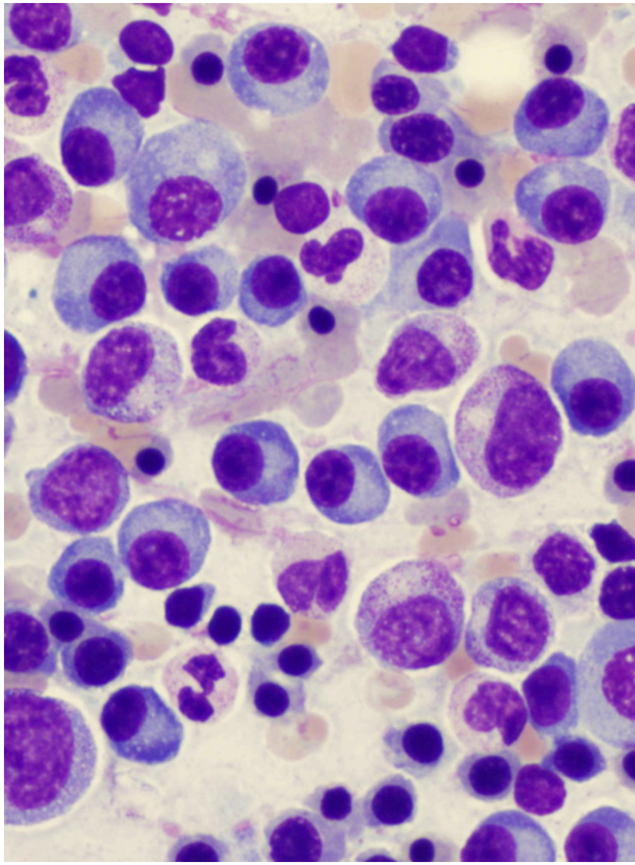

Standard Reconstruction of Lens-Free Images

Granulocyte Lymphocyte Monocyte

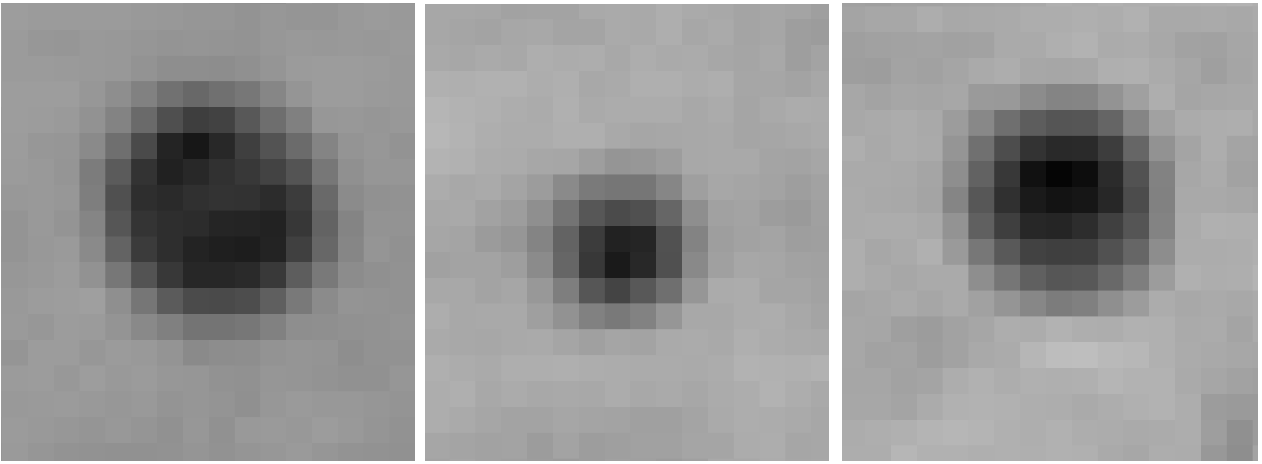

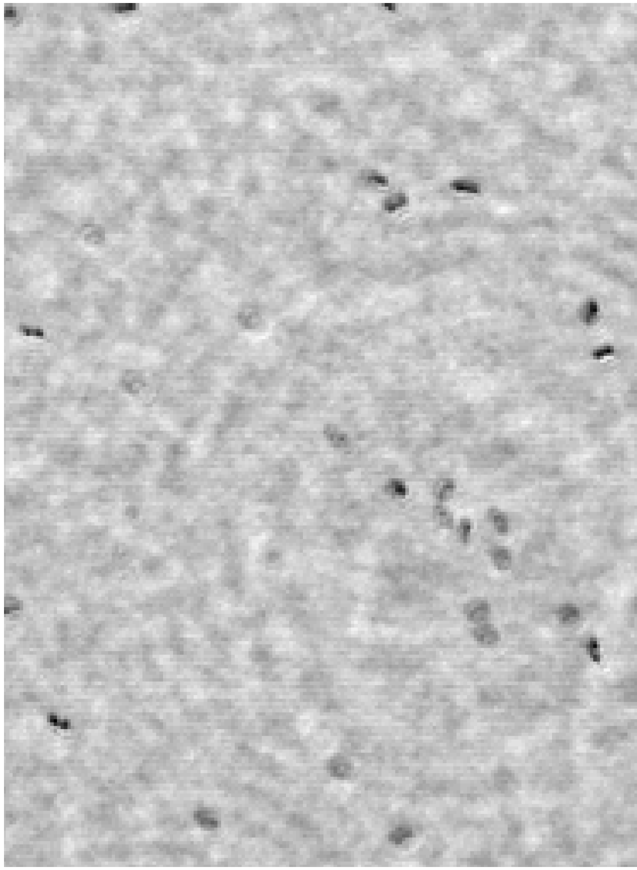

Sparse Phase Recovery Reconstruction of Lens-Free Images

$$\min_{X,W,\mu} \frac{1}{2}\|H \odot W - \mu\mathbf{1} - T(z) * X\|_F^2 + \lambda\|X\|_1 \quad \text{s.t.} \quad |W| = \mathbf{1}$$

Granulocyte | Lymphocyte | Monocyte

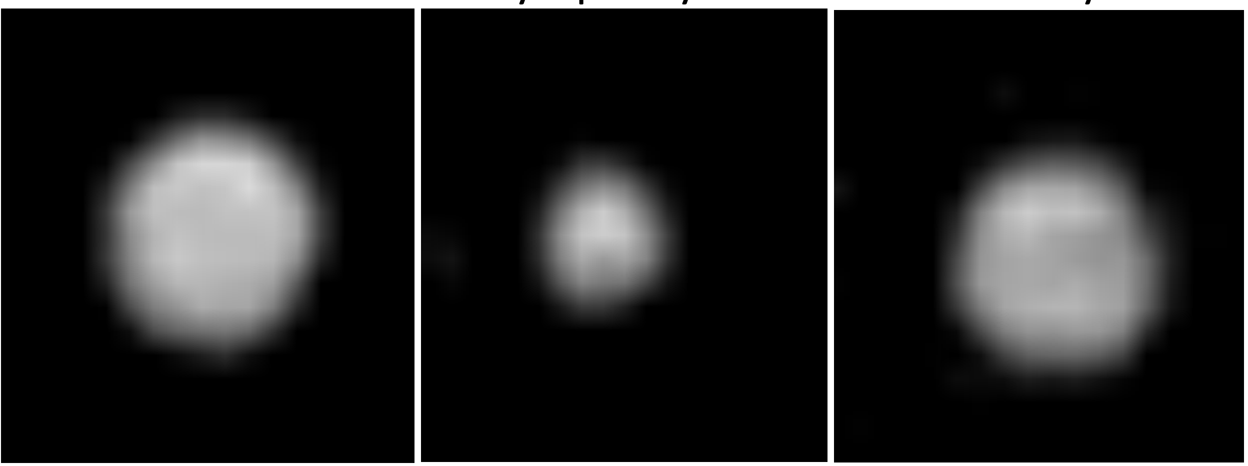

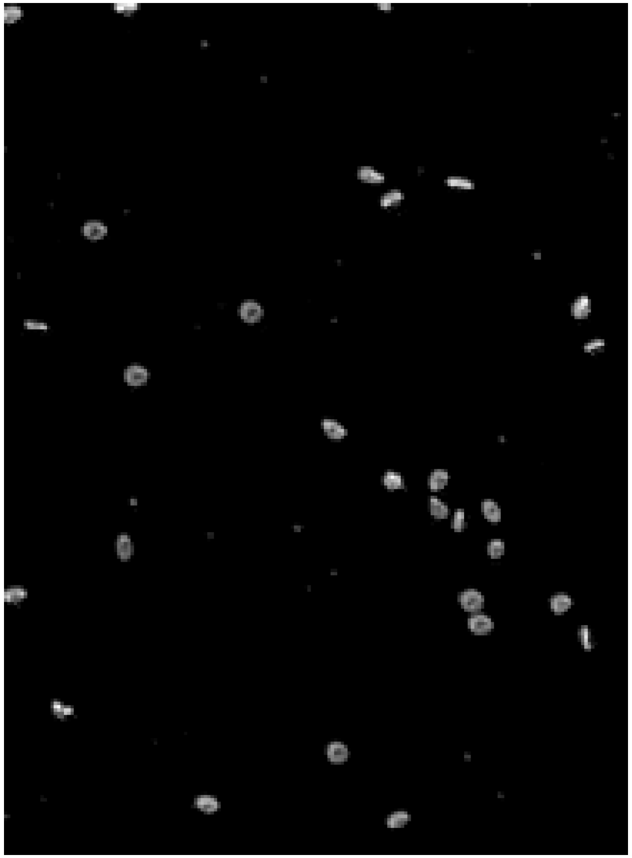

- ✓ Remove background
- ✓ Recover phase

B. Haeffele, R. Stahl, G. Vanmeerbeeck, and R. Vidal, "Efficient Reconstruction of Holographic Lens-Free Images by Sparse Phase Recovery ." MICCAI, 2017.

Lens-Free Images of White Blood Cells

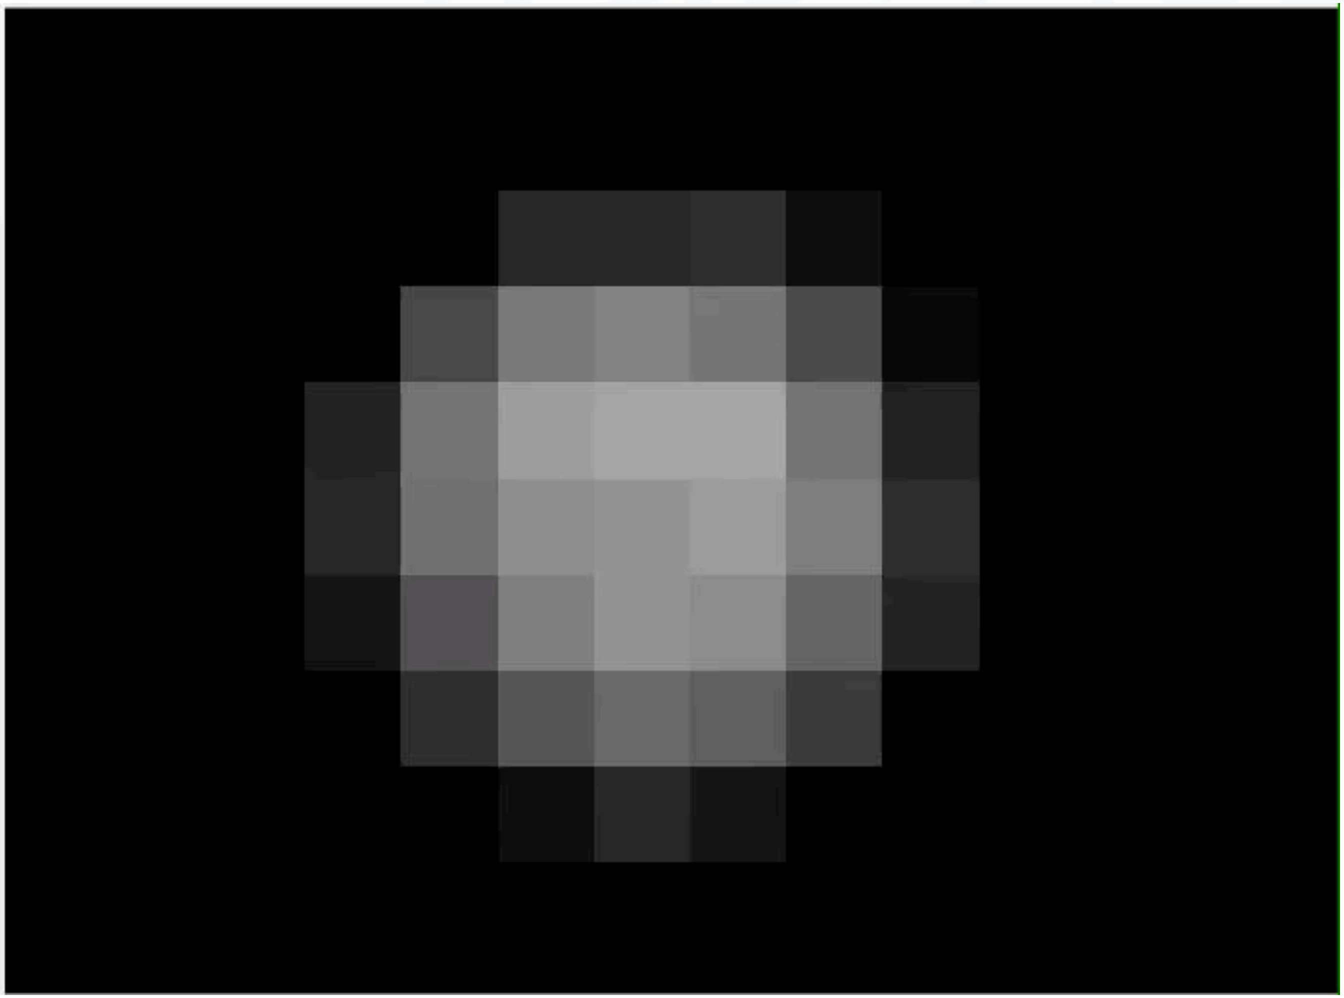

✗ Low resolution
✓ Compact size
✓ Low-cost
✓ Large field of view

Possible Approach to Detection & Classification

Threshold

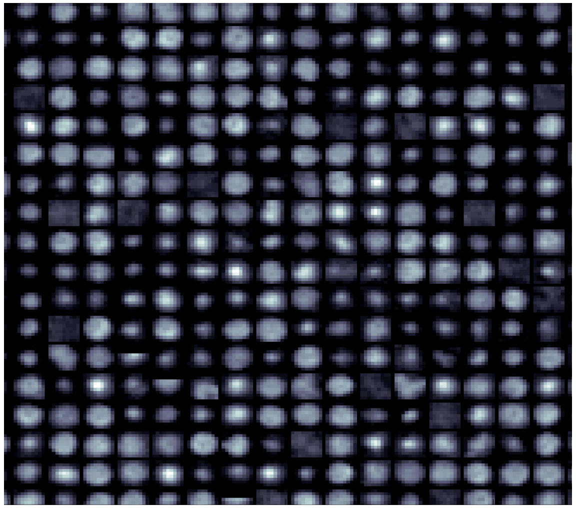

CNN

Class Labels

Our Approach to Detection & Classification

Method to detect, count and classify cell populations in LFI

1. Generative probabilistic model for multi-object detection & classification

2. Weakly supervised learning without per-object bounding boxes

3. Efficient inference method to detect, count, and classify populations of hundreds to thousands of cells per image

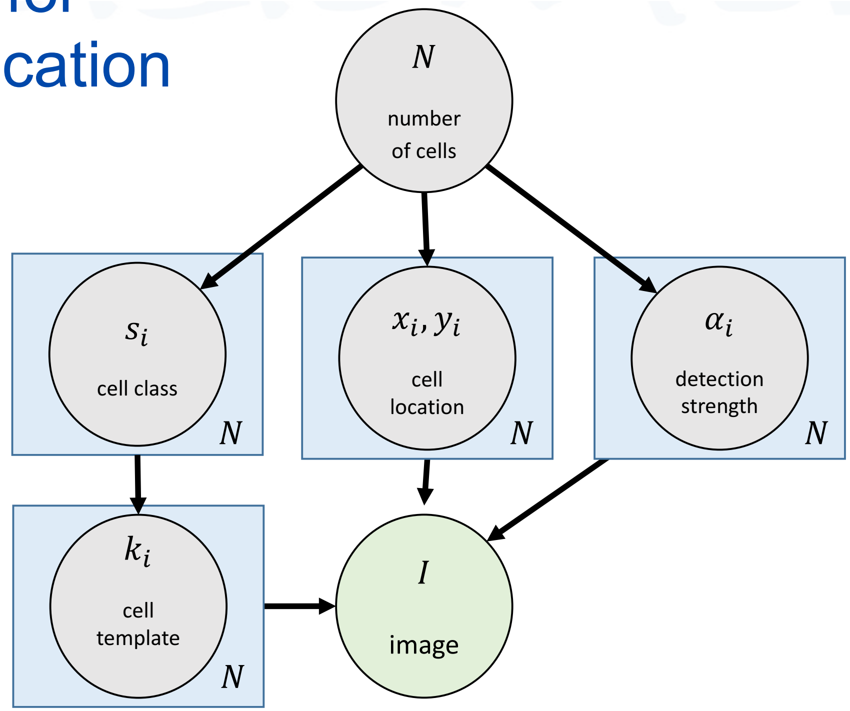

Florence Yellin, Benjamin D Haeffele, Sophie Roth, René Vidal, "Multi-Cell Detection and Classification Using a Generative Convolutional Model" CVPR, 2018.

Learning Image Parameters from Purified Blood

- Experimentally isolate cells from a single subclass to obtain purified blood and use detected cells as training examples

- **Domain adaptation challenge**: purified data may not be representative mixed-cell populations

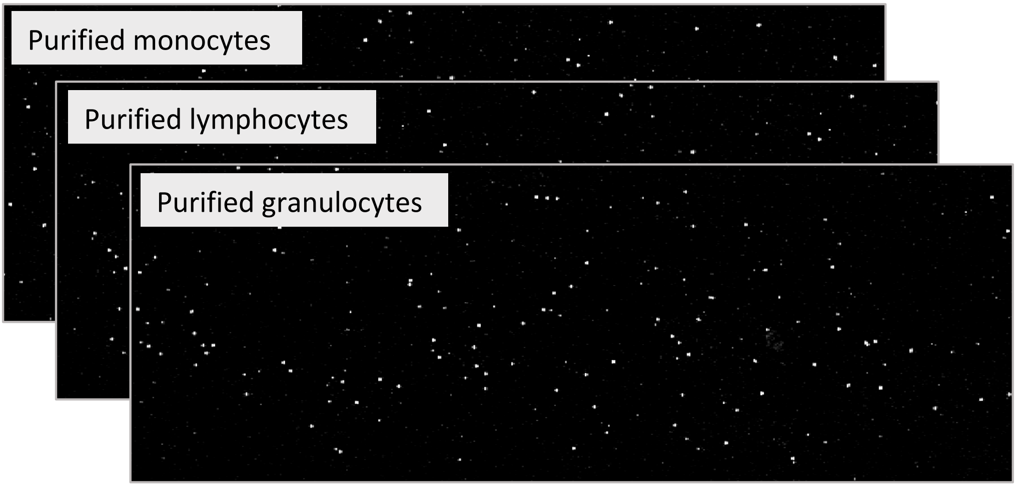

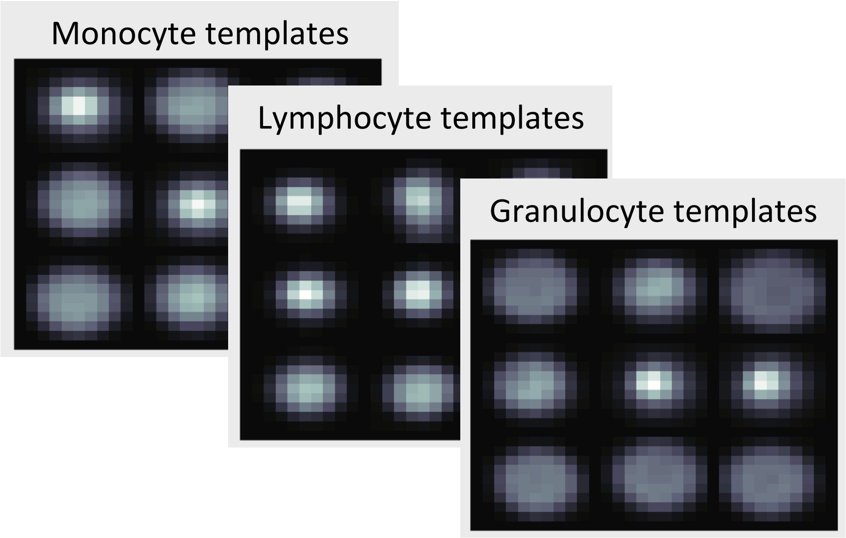

Class Proportion Results

Mean Absolute Error	CNN	CSC	Ours
Granulocytes - normal	27.8	31.1	**4.5**
Lymphocytes - normal	12.8	9.5	**4.6**
Monocytes - normal	15.9	21.9	**4.7**
Granulocytes - all	28.6	30.1	**6.8**
Lymphocytes - all	11.6	8.3	**5.6**
Monocytes - all	18.9	22.3	**5.5**

JHU vision lab

Machine Learning in Neuroscience

René Vidal
Herschel L. Seder Professor of Biomedical Engineering
Director, Mathematical Institute for Data Science
Johns Hopkins University

JOHNS HOPKINS MATHEMATICAL INSTITUTE for DATA SCIENCE

THE DEPARTMENT OF BIOMEDICAL ENGINEERING
The Whitaker Institute at Johns Hopkins

Center for IMAGING SCIENCE

Computer Vision for Neuroscience

Single Electrode

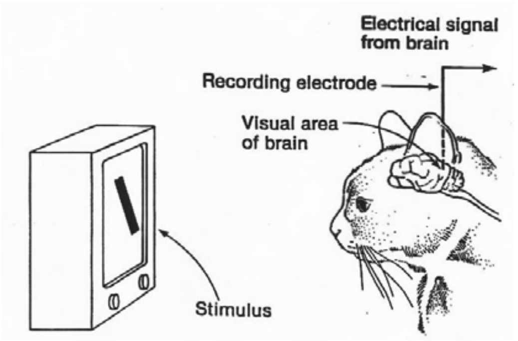

- Classical gold standard
- One neuron at a time

Functional MRI

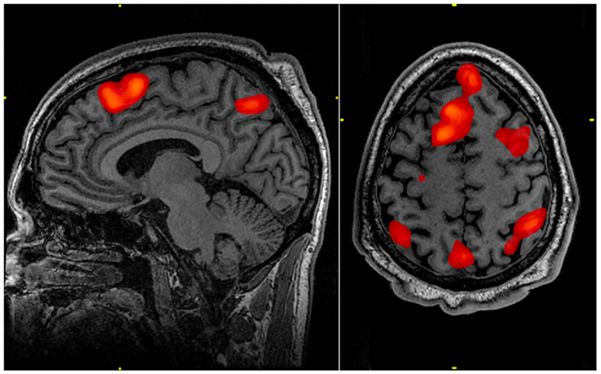

- Whole brain activity
- Averaged over 1000s of neurons

Electrode Arrays

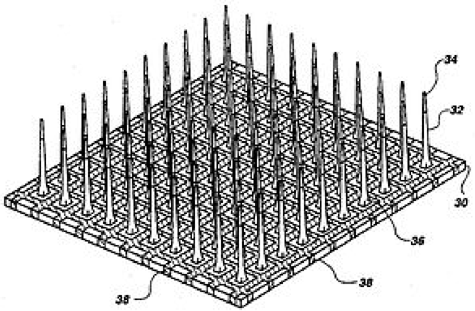

- Multiple neurons simultaneously
- Limited spatial information

Computer Vision for Neuroscience

- Fluorescent microscopy technique
  - Optical recording of brain activity
  - Neurons “flash” when active electrically

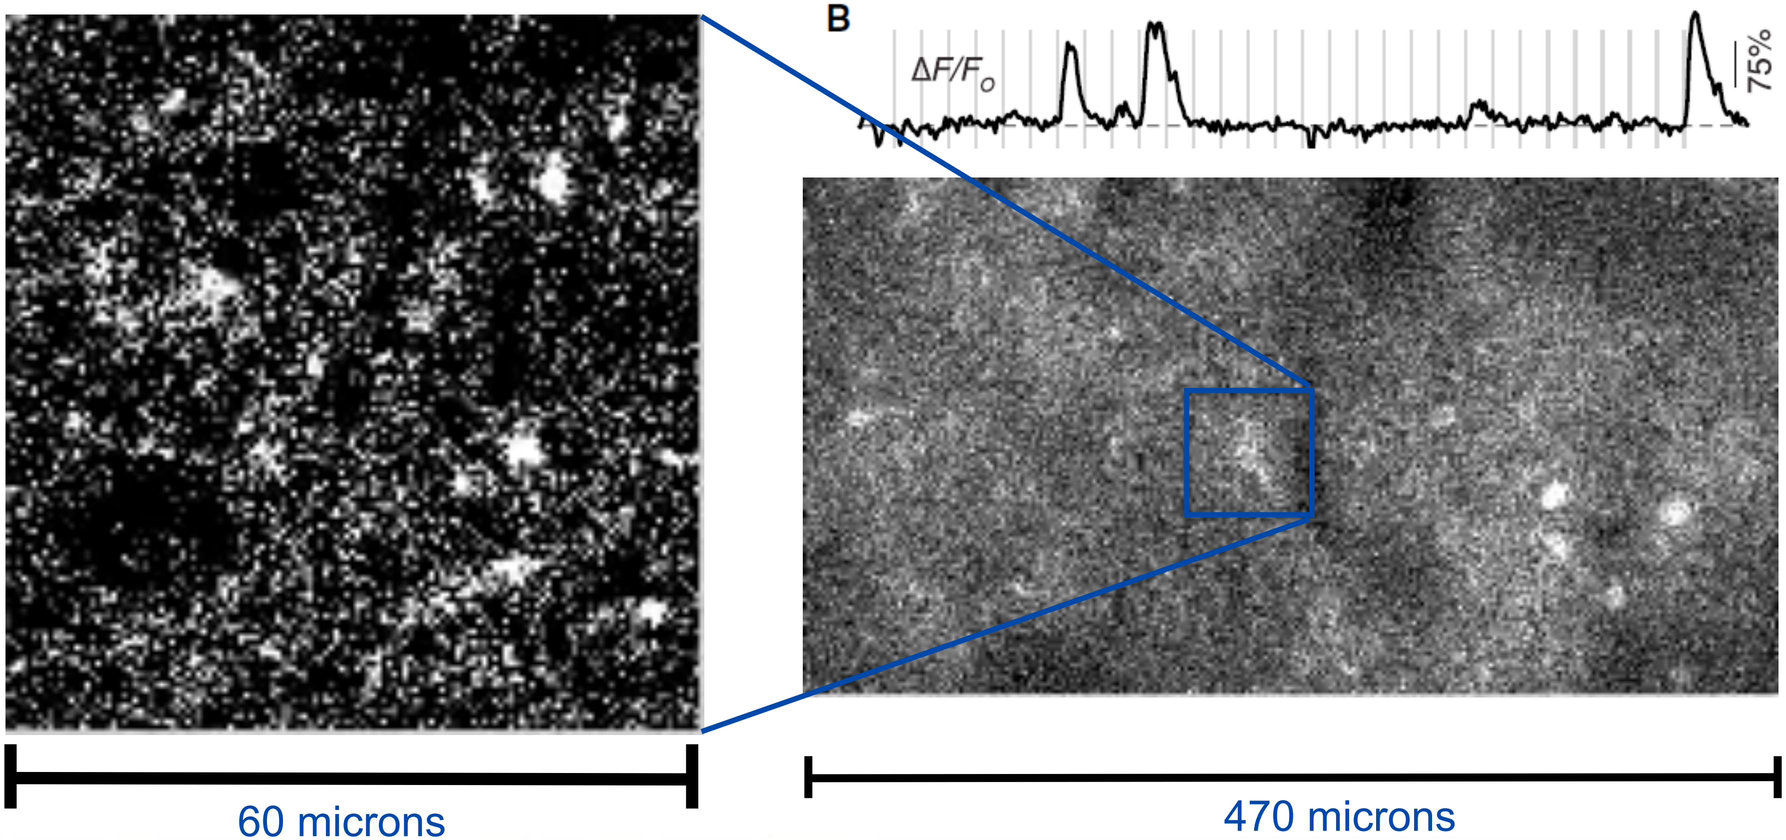

Stosiek, Garaschuk, Holthoff, Konnerth, “In vivo two-photon calcium imaging of neuronal networks,” PNAS, 100(12):7319–7324, 2003.

Computer Vision for Neuroscience

- Video can be approximated as the sum of rank-1 matrices

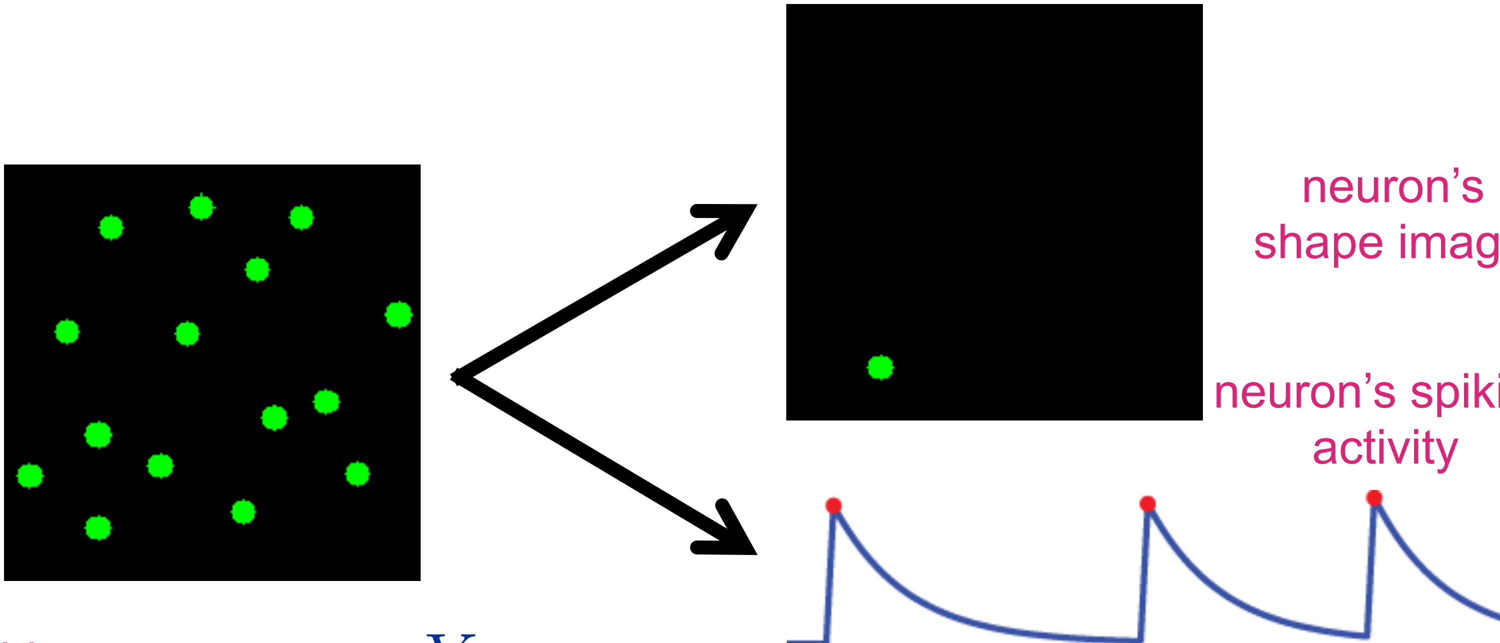

neuron's shape image $= U_i$

neuron's spiking activity $= V_i$

video sequence $= Y$

neuron's temporal signal $= V_i(t) \star D(t)$

number of neurons

$$Y \approx \sum_{i=1}^{r} U_i V_i^\top D$$

Haeffele, Young, Vidal. Structured Low-Rank Matrix Factorization: Optimality, Algorithm, and Applications to Image Processing, ICML '14

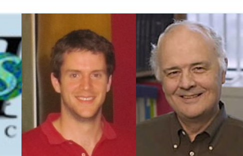

Computer Vision for Neuroscience

- Desired properties of shape and spike matrices
  - U_i should be sparse (low L1)
  - U_i should be compact (low TV)
  - V_i should be sparse (low L1)

$$Y \approx \sum_{i=1}^{r} U_i V_i^\top D$$

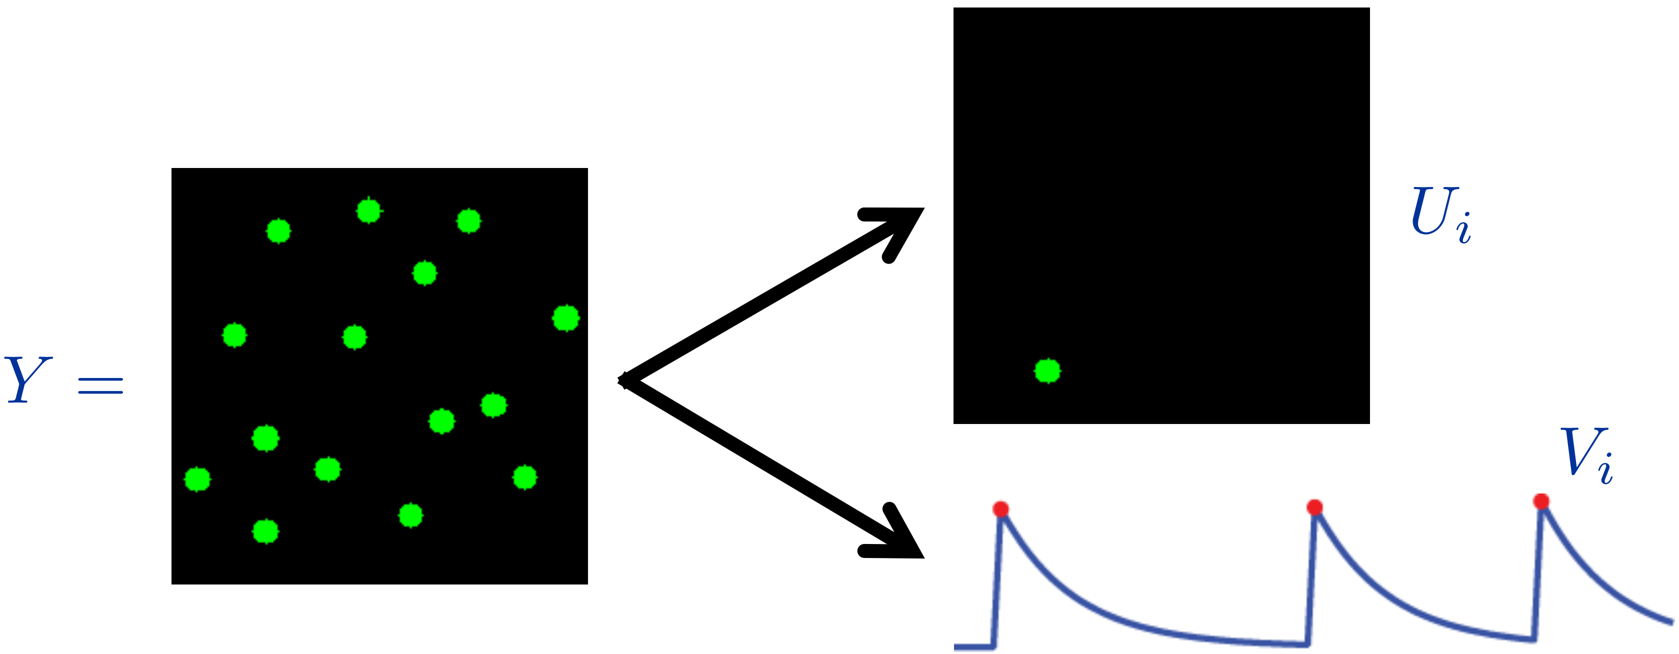

Haeffele, Young, Vidal. Structured Low-Rank Matrix Factorization: Optimality, Algorithm, and Applications to Image Processing, ICML ’14

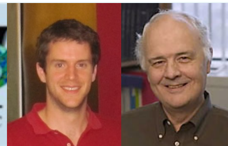

Computer Vision for Neuroscience

$$\min_{U,V} \|Y - UV^\top D\|_F^2 + \lambda \sum_{i=1}^{r} \|U_i\|_u \|V_i\|_v$$

few neurons

$$\|\cdot\|_u = \|\cdot\|_2 + \|\cdot\|_1 + \|\cdot\|_{TV}$$

compact

$$\|\cdot\|_v = \|\cdot\|_2 + \|\cdot\|_1$$

sparse

60 microns

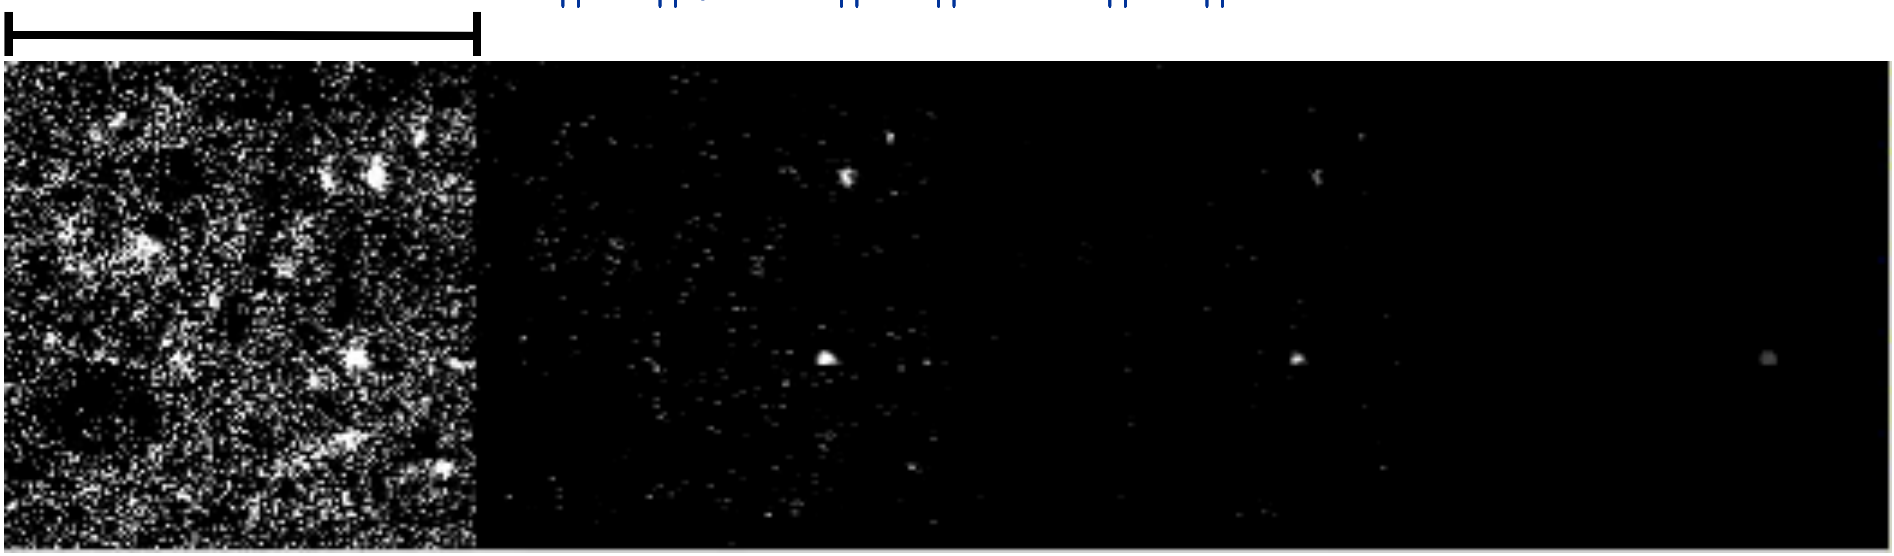

Raw Data | Sparse | + Low Rank | + Total Variation

Haeffele, Young, Vidal. Structured Low-Rank Matrix Factorization: Optimality, Algorithm, and Applications to Image Processing, ICML '14

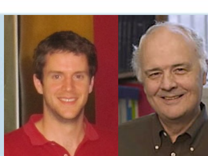

JHU vision lab

Machine Learning in Regenerative Medicine

René Vidal

Herschel L. Seder Professor of Biomedical Engineering

Director, Mathematical Institute for Data Science

Johns Hopkins University

JOHNS HOPKINS MATHEMATICAL INSTITUTE *for* DATA SCIENCE

THE DEPARTMENT OF BIOMEDICAL ENGINEERING

The Whitaker Institute at Johns Hopkins

Center for IMAGING SCIENCE

Machine Learning in Regenerative Medicine

- Cardiovascular disease is the world's leading cause of death.
  - 17.3 million deaths per year.
  - 787,000 in the US (2011).
- Myocardial infarction is one of the leading causes of sudden cardiac arrest.
- Stem cells present a potential avenue to treating myocardial infarction.

https://www.heart.org/

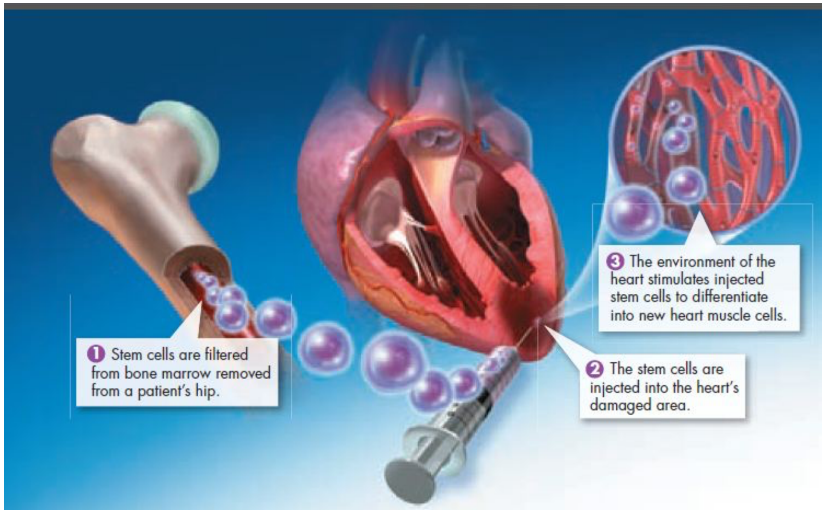

http://www.georgeinstitute.org/units/cardiovascular

Laflamme, Chen, Naumova, Muskheli, Fugate, Dupras, Reinecke, Xu, Hassanipour, Police, O'Sullivan, Collins, Chen, Minami, Gill, Ueno, Yuan, Gold, Murray. Cardiomyocytes derived from human embryonic stem cells in pro-survival factors enhance function of infarcted rat hearts. Nature biotechnology 25(9):1015–24, 2007

Machine Learning in Regenerative Medicine

- Predict cell phenotype from the shape of its action potential
- Maturation process affects the shape of the action potential

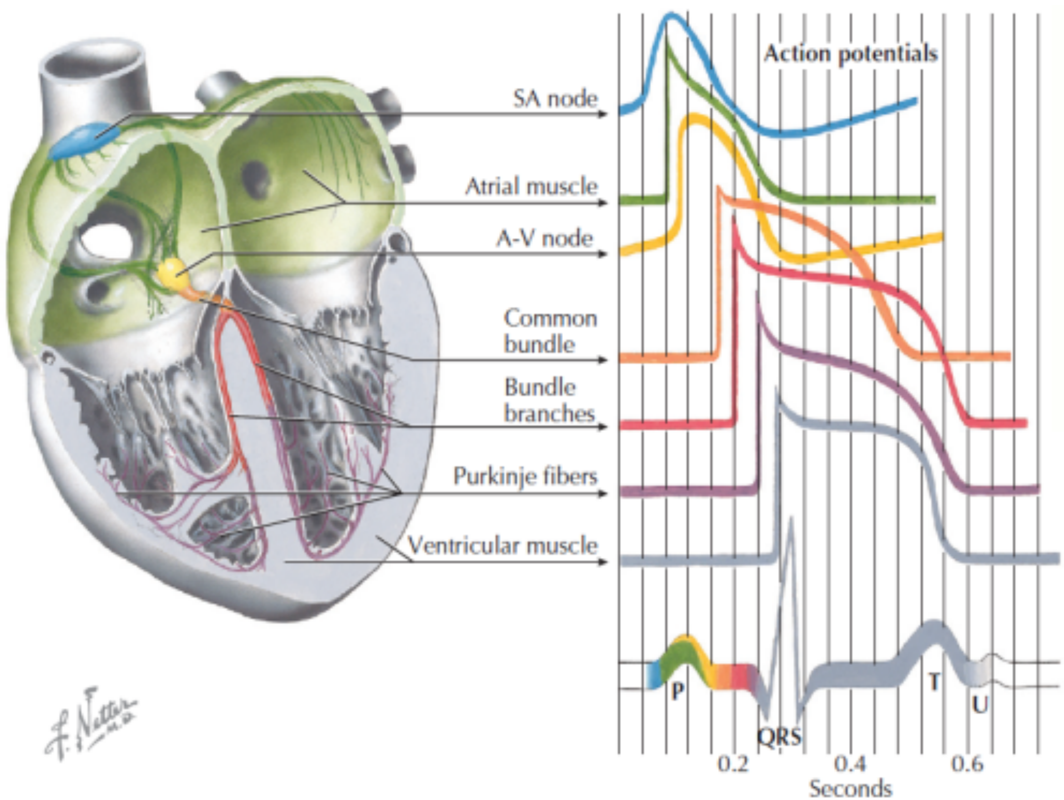

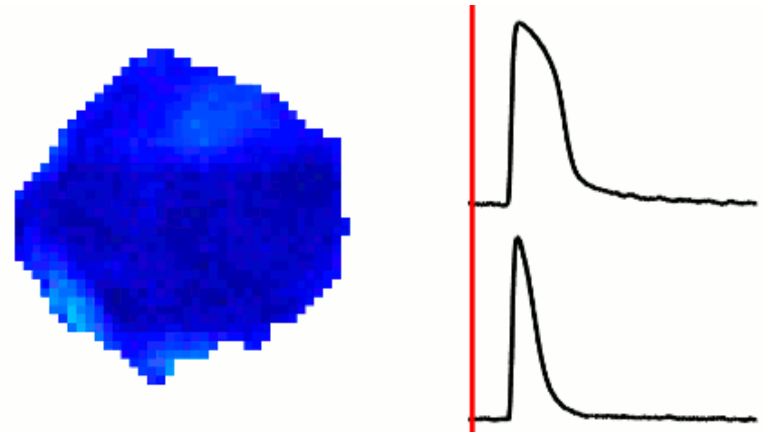

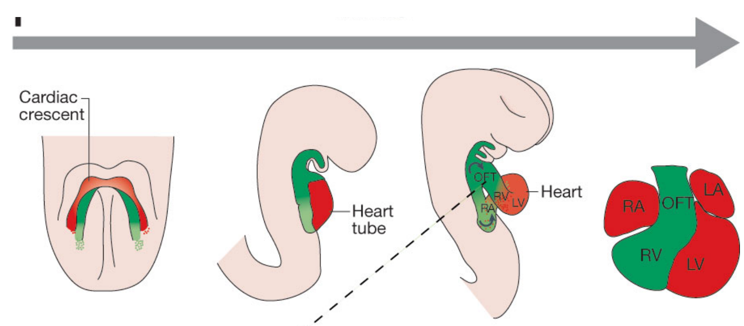

JOHNS HOPKINS
MATHEMATICAL INSTITUTE
for DATA SCIENCE

Gorospe, Zhu, Milrod, Zambidis, Tung, Vidal. Automated grouping of action potentials of human embryonic stem cell-derived cardiomyocytes. TBME13.
Gorospe, Younes, Tung, Vidal. A metamorphosis distance for embryonic cardiac action potential interpolation and classification. MICCAI13.
Gorospe, Zhu, He, Tung, Younes, Vidal. Efficient metamorphosis computation for classifying embryonic cardiac action potentials. MICCAI15.

Machine Learning in Regenerative Medicine

- **Metamorphosis** interpolation produces intermediate shapes that better resemble those of a cardiac action potential.

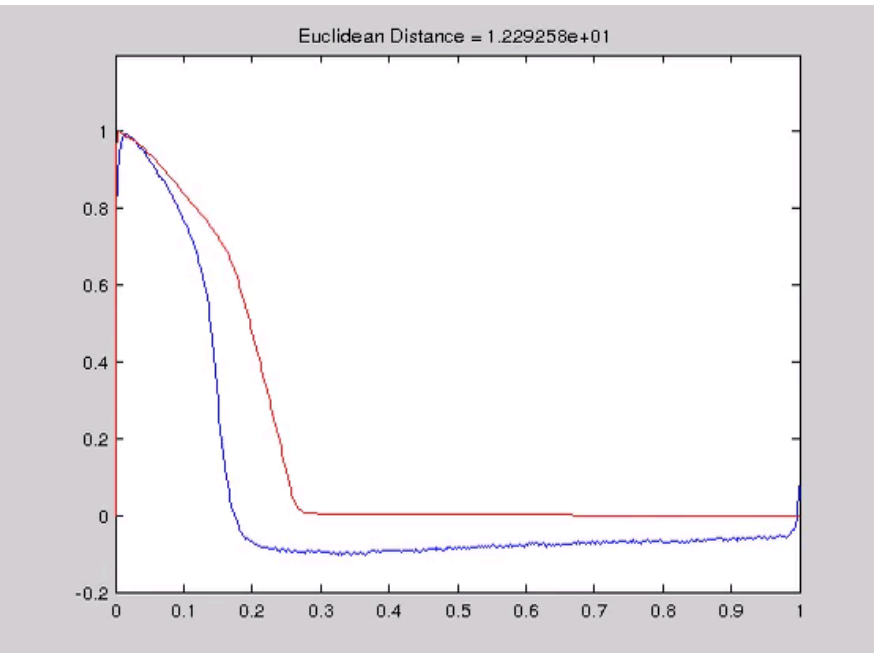

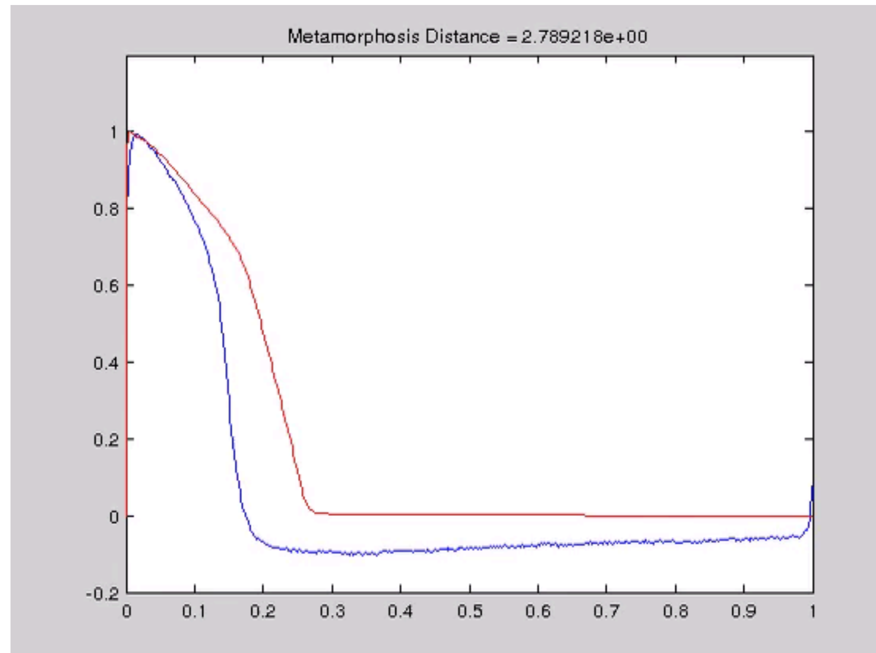

$$d^2_{\mathcal{M}}(I_0, I_1) = \inf_{\substack{v,I \\ I(t,0)=I_0(t) \\ I(t,1)=I_1(t)}} \int_0^1 \frac{1}{2}\left(\|v(t,\tau)\|^2_V + \frac{1}{\sigma^2}\left\| \frac{\partial I}{\partial \tau}(t,\tau) + \frac{\partial I}{\partial t}(t,\tau)\ v(t,\tau) \right\|^2_{L^2} \right) d\tau$$

JOHNS HOPKINS
MATHEMATICAL INSTITUTE
for DATA SCIENCE

Gorospe, Zhu, Milrod, Zambidis, Tung, Vidal. Automated grouping of action potentials of human embryonic stem cell-derived cardiomyocytes. TBME14.
Gorospe, Younes, Tung, Vidal. A metamorphosis distance for embryonic cardiac action potential interpolation and classification. MICCAI13.
Gorospe, Zhu, He, Tung, Younes, Vidal. Efficient metamorphosis computation for classifying embryonic cardiac action potentials. MICCAI15.

Clustering Embryonic CM Action Potentials

- Clustering results for 9 cell clusters of 6940 embryonic APs

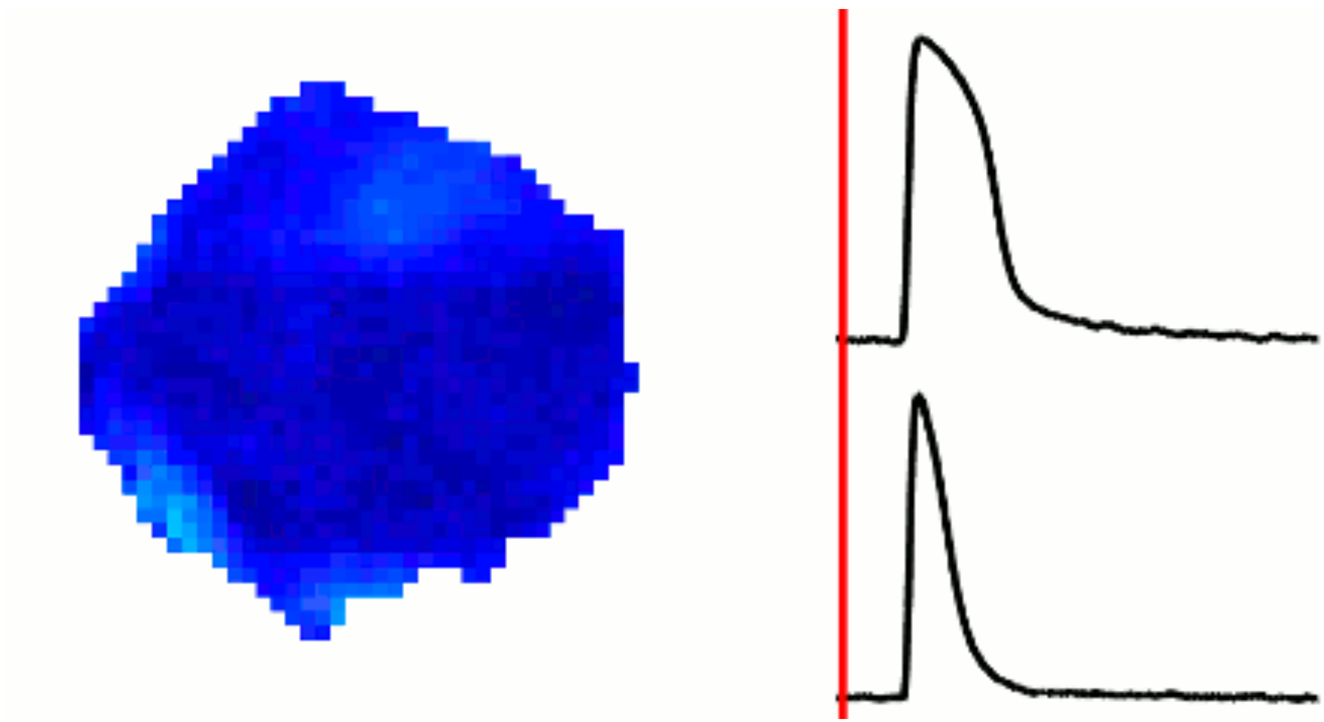

JOHNS HOPKINS
MATHEMATICAL INSTITUTE
for DATA SCIENCE

Gorospe, Zhu, Milrod, Zambidis, Tung, Vidal. Automated grouping of action potentials of human embryonic stem cell-derived cardiomyocytes. TBME14.
Gorospe, Younes, Tung, Vidal. A metamorphosis distance for embryonic cardiac action potential interpolation and classification. MICCAI13.
Gorospe, Zhu, He, Tung, Younes, Vidal. Efficient metamorphosis computation for classifying embryonic cardiac action potentials. MICCAI15.

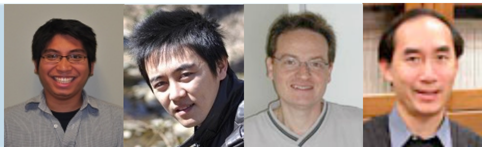

JHU vision lab

Machine Learning in Brain Image Analysis

René Vidal
Herschel L. Seder Professor of Biomedical Engineering
Director, Mathematical Institute for Data Science
Johns Hopkins University

Johns Hopkins Mathematical Institute for Data Science

The Department of Biomedical Engineering
The Whitaker Institute at Johns Hopkins

Center for Imaging Science

Machine Learning for Brain Image Analysis

- The human brain contains 100 billion neurons connected in a complex network of fiber bundles.
- Studies show that in neurological disorders such as schizophrenia, autism and Alzheimer's disease, these connections are damaged.
- By mapping brain connections, we can compare neuro anatomies of healthy and diseased brains to develop new tools for diagnosis of neurological diseases.

Machine Learning for Brain Image Analysis

Finding biomarkers of neurological disease in diffusion MRI data

2. Processing

3. Segmentation

Weighted Signals

Filtered ODF field

Fiber Bundles

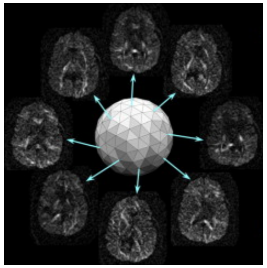

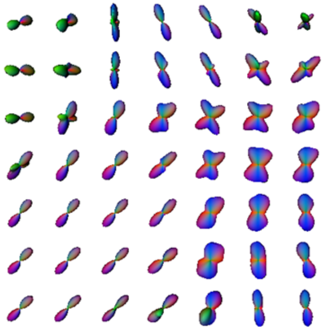

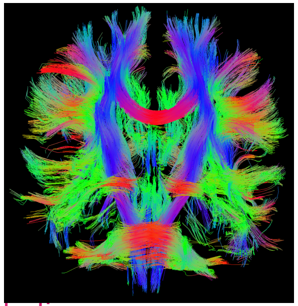

1. Estimation

5.

4. Registration

ODF Field

Validation

Atlas Construction

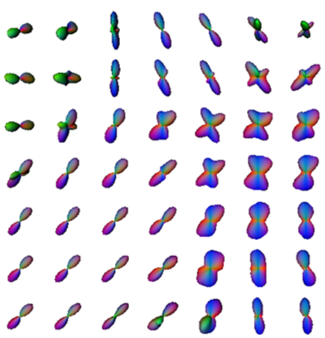

Study gives more proof that intelligence is largely inherited

UCLA researchers find that genes determine brain's processing speed

By Mark Wheeler | 3/17/2009

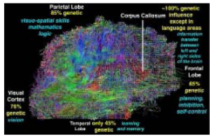

They say a picture tells a thousand stories, but can it also tell how smart you are? Actually, say UCLA researchers, it can.

In a study published in the Journal of Neuroscience Feb. 18, UCLA neurolo professor Paul Thompson and colleagues used a new type of brain-imaging scanner to show that intelligence is strongly influenced by the quality of t brain's axons, or wiring that sends signals throughout the brain. The faste the signaling, the faster the brain processes information. And since the integrity of the brain's wiring is influenced by genes, the genes we inherit play a far greater role in intelligence than was previously thought.

Genes appear to influence intelligence by determining how well nerve axo are encased in myelin — the fatty sheath of "insulation" that coats our axons and allows for fast signaling bursts in our brains. The thicker the myelin, the faster the nerve impulses.

JOHNS HOPKINS
MATHEMATICAL INSTITUTE
for DATA SCIENCE

Goh et al. "A nonparametric Riemannian framework for processing high angular resolution diffusion images and its applications to ODF-based morphometry" ECCV06, ECCV08, CVPR09, MICCAI09, Neuroimage11.
Cetingul et al. "Group action induced averaging for HARDI processing" ISBI12, ISBI12, TBME14
Schwab et al. "Rotation invariant features for HARDI" MICCAI12, IPMI13, CDMRI15
Wolfers et al. "Nonnegative ODF estimation via optimal constraint

Machine Learning for Brain Image Analysis

- Reconstruction of dMRI images

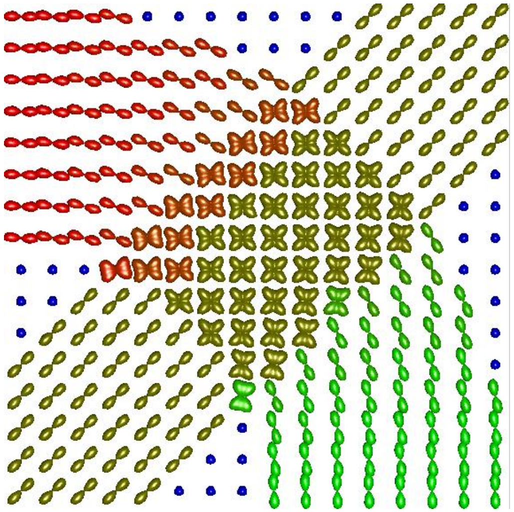

- **Problem**: dMRI currently too slow for clinical use (requires 100s of MRIs)
- **Question**: How can we accelerate acquisition but still estimate accurate fiber tracks?
- **Solution**: Compressed Sensing

If sparse representation of data → only need few measurements → Accelerate acquisition

JOHNS HOPKINS
MATHEMATICAL INSTITUTE
for DATA SCIENCE

E Schwab, R Vidal, N Charon. (k, q)-Compressed Sensing for dMRI with Joint Spatial-Angular Sparsity Prior, arXiv 2017
E Schwab, R Vidal, N Charon. Efficient Global Spatial-Angular Sparse Coding for Diffusion MRI with Separable Dictionaries, arXiv 2016
E Schwab, R Vidal, N Charon, Spatial-Angular Sparse Coding for HARDI. MICCAI 2017

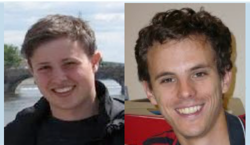

Machine Learning for Brain Image Analysis

- High throughput neuroinformatics: bits of neuroscience at 1mm scale
  - 3000 brains
  - 1000x1000x500x100 dimensions
  - 1000-2000 relevant variables

- BrainGPS generates a machine learnable feature vector, a BRAINPRINT.

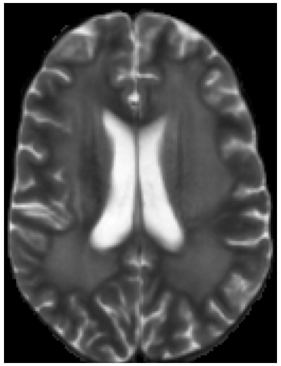

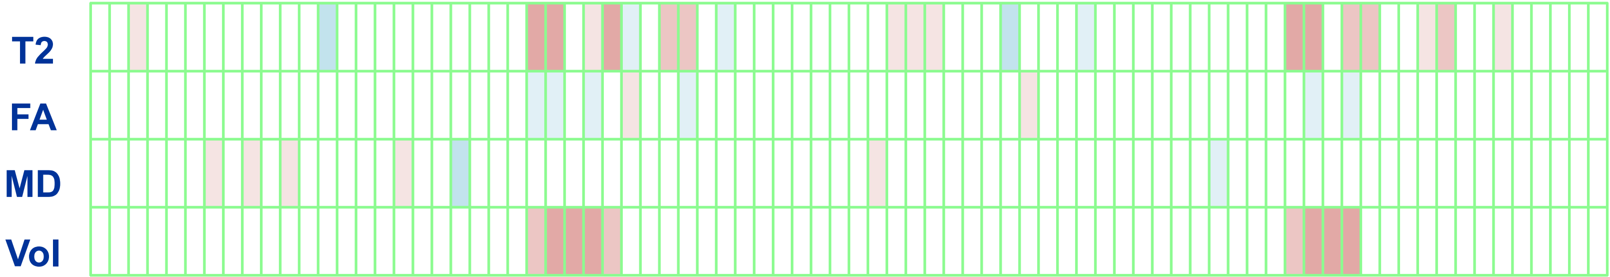

Brain Science Institute, Johns Hopkins University
Michael Miller, Susumu Mori, Andreia Faria, Kenichi Oishi

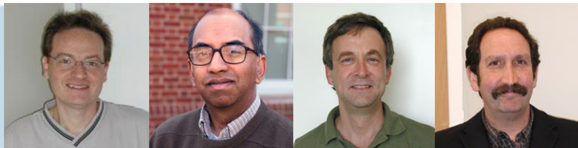

JHU vision lab

Machine Learning in Surgery

René Vidal
Herschel L. Seder Professor of Biomedical Engineering
Director, Mathematical Institute for Data Science
Johns Hopkins University

JOHNS HOPKINS MATHEMATICAL INSTITUTE for DATA SCIENCE

THE DEPARTMENT OF BIOMEDICAL ENGINEERING
The Whitaker Institute at Johns Hopkins

Center for IMAGING SCIENCE

Machine Learning in Surgery

- RMIS has the potential to revolutionize our understanding of modeling, teaching and evaluating human manipulation skills.

- The goal of the project is to develop quantitative methods for modeling surgical tasks and evaluation of surgical skill.

Machine Learning in Surgery

Recognizing surgical gestures and assessing the skill level of a surgeon in kinematic and video data of a surgical procedure

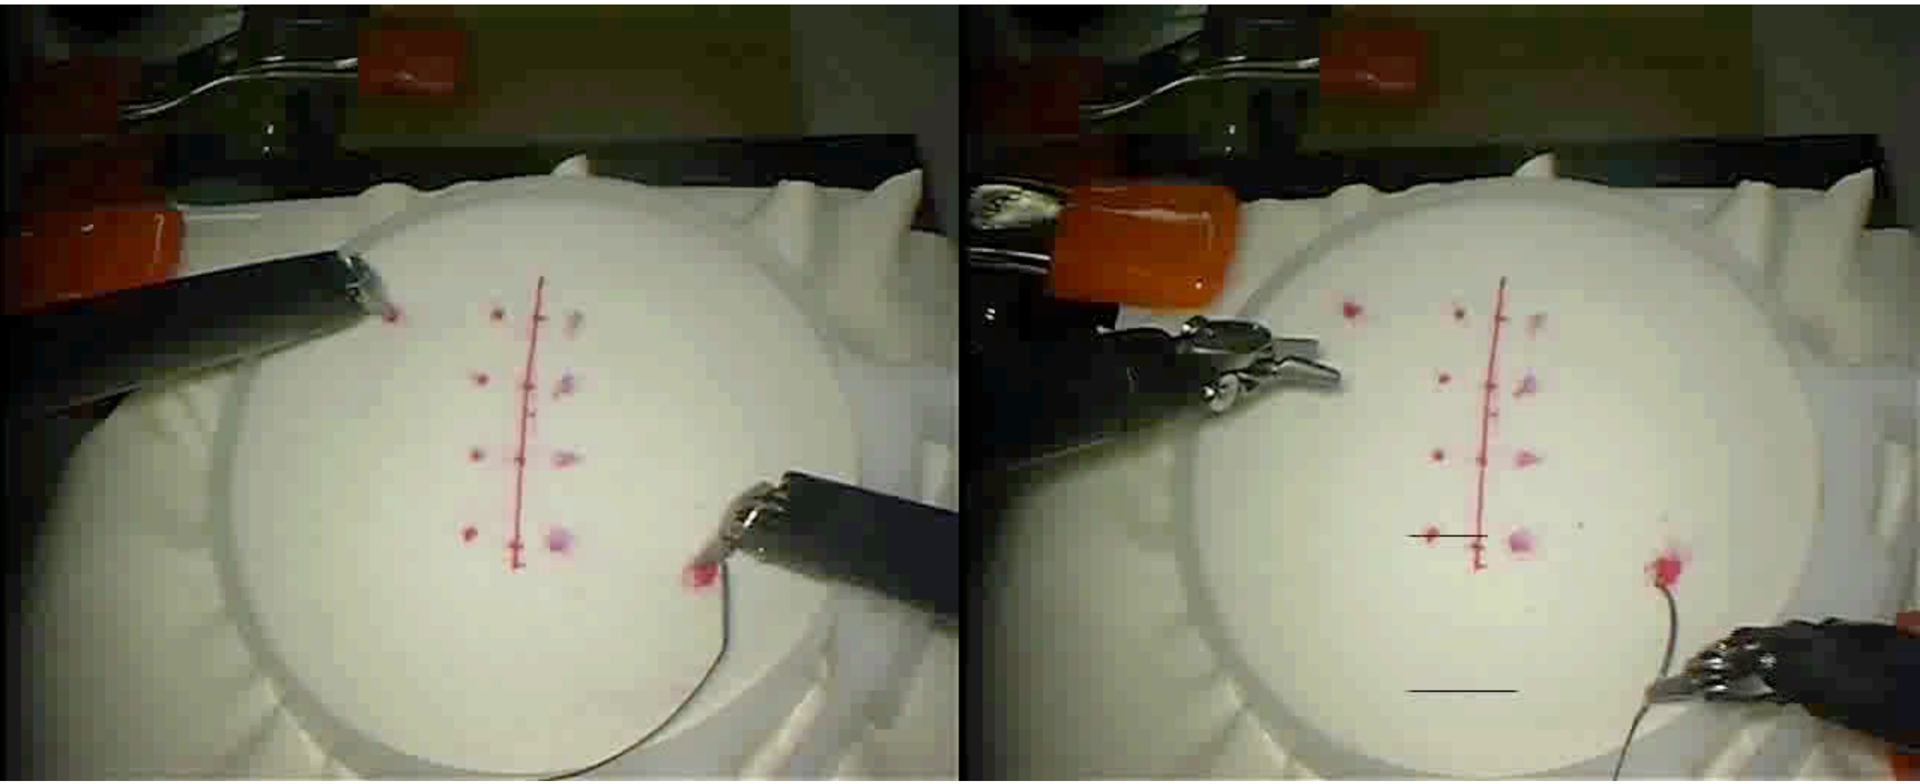

JOHNS HOPKINS
MATHEMATICAL INSTITUTE
for DATA SCIENCE

Bejar, Zappella, Vidal. Surgical Gesture Classification from Video Data, MICCAI12 (**Best paper Award)**
Zappella, Bejar, Hager, Vidal. Surgical Gesture Classification from Kinematic and Video Data, MedIA13
Tao, Zappella, Hager, Vidal. Surgical Gesture Segmentation and Recognition, MICCAI13
Tao, Elhamifar, Khudanpur, Hager, Vidal. Sparse HMMs for Surgical Gesture Classification and Skill Evaluation, IPCAI12

Modeling the Language of Surgery

- Similar to speech, surgical motion is not random:
  - A procedure is composed of tasks (incision, suturing, knot tying, etc.)
  - A task is composed of gestures (insert needle, pull needle, etc.)
  - Procedures, tasks and gestures follow a grammar.

- Our goal: de[redacted]emes and the surgical gra[redacted]

Reiley, C.E., Lin, H.C., Varadarajan, B., Vagolgyi, B., Khudanpur, S., Yuh, D.D., Hager, G.D.: Automatic recognition of surgical motions using statistical modeling for capturing variability. In: Medicine Meets Virtual Reality. (2008) 396–401

Markov/Semi-Markov Random Field Models

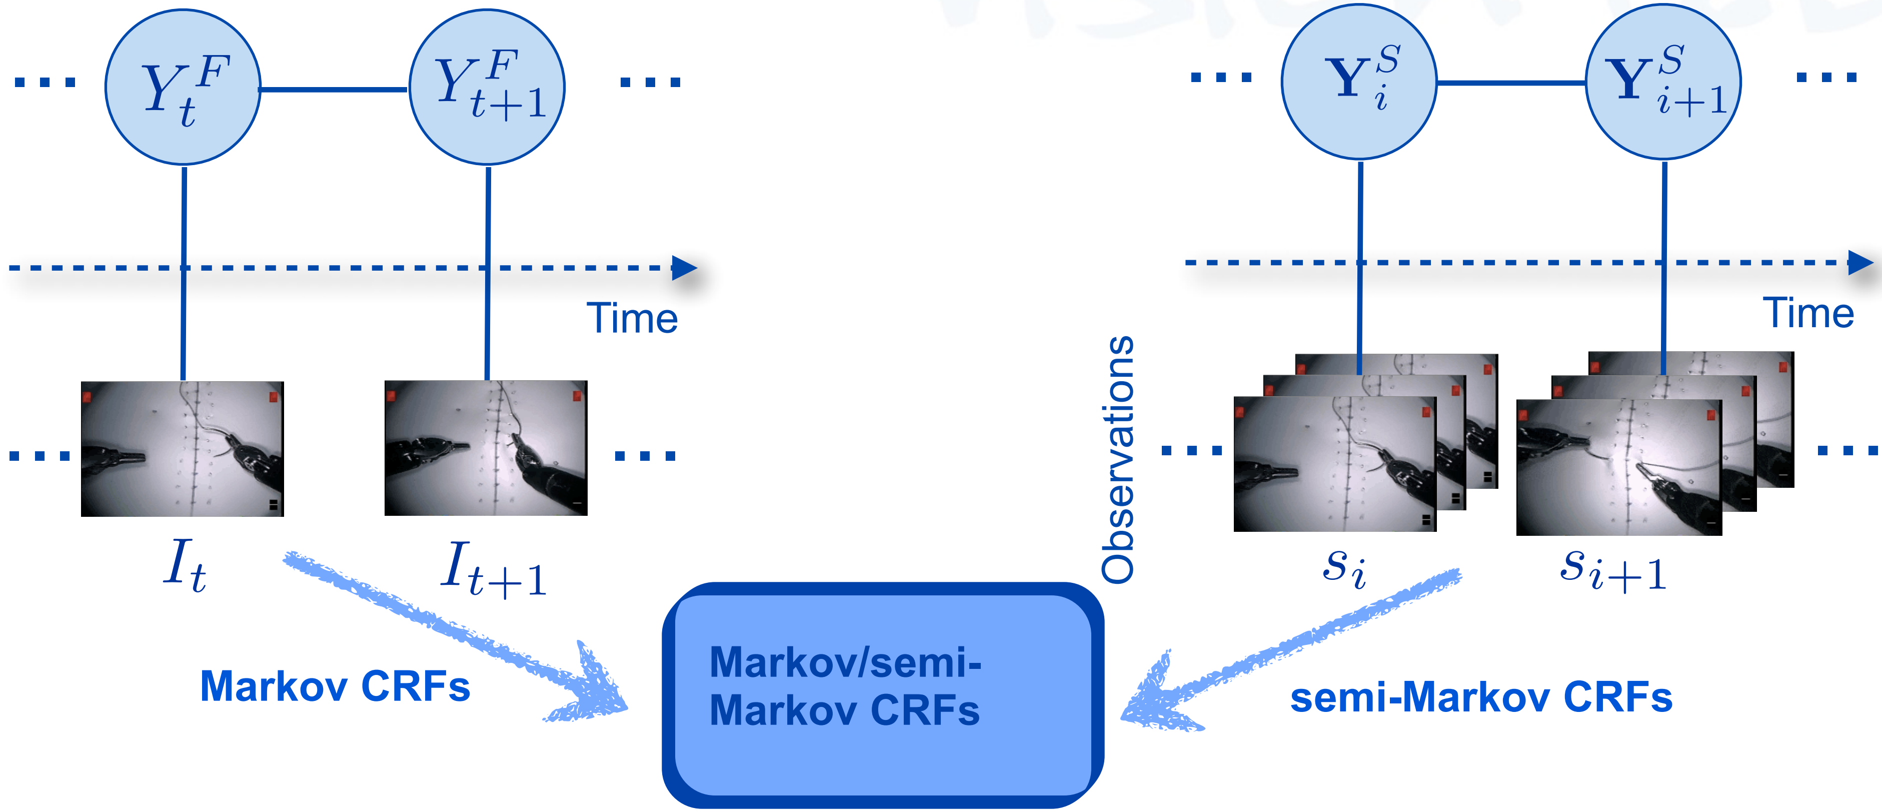

- Inference: find sequence of gestures using a modified Viterbi
- Learning: find parameters using structural output SVMs

[1] Bejar, Zappella, Vidal. Surgical Gesture Classification from Video Data, MICCAI12 (Best paper Award)
[2] Zappella, Bejar, Hager, Vidal. Surgical Gesture Classification from Kinematic and Video Data, MedIA13
[3] Tao, Zappella, Hager, Vidal. Surgical Gesture Segmentation and Recognition, MICCAI13
[4] Tao, Elhamifar, Khudanpur, Hager, Vidal. Sparse HMMs for Surgical Gesture Classification and Skill Evaluation, IPCAI12
[5] Sefati, Cowan,Vidal. Learning Shared, Discriminative Dictionaries for Surgical Gesture Segmentation and Classification, M2CAI15
[6] Lea, Hager, Vidal. An improved model for segmentation and recognition of fine-grained activities with application to surgical training tasks, WACV15

Segmental Spatio-Temporal CNNs

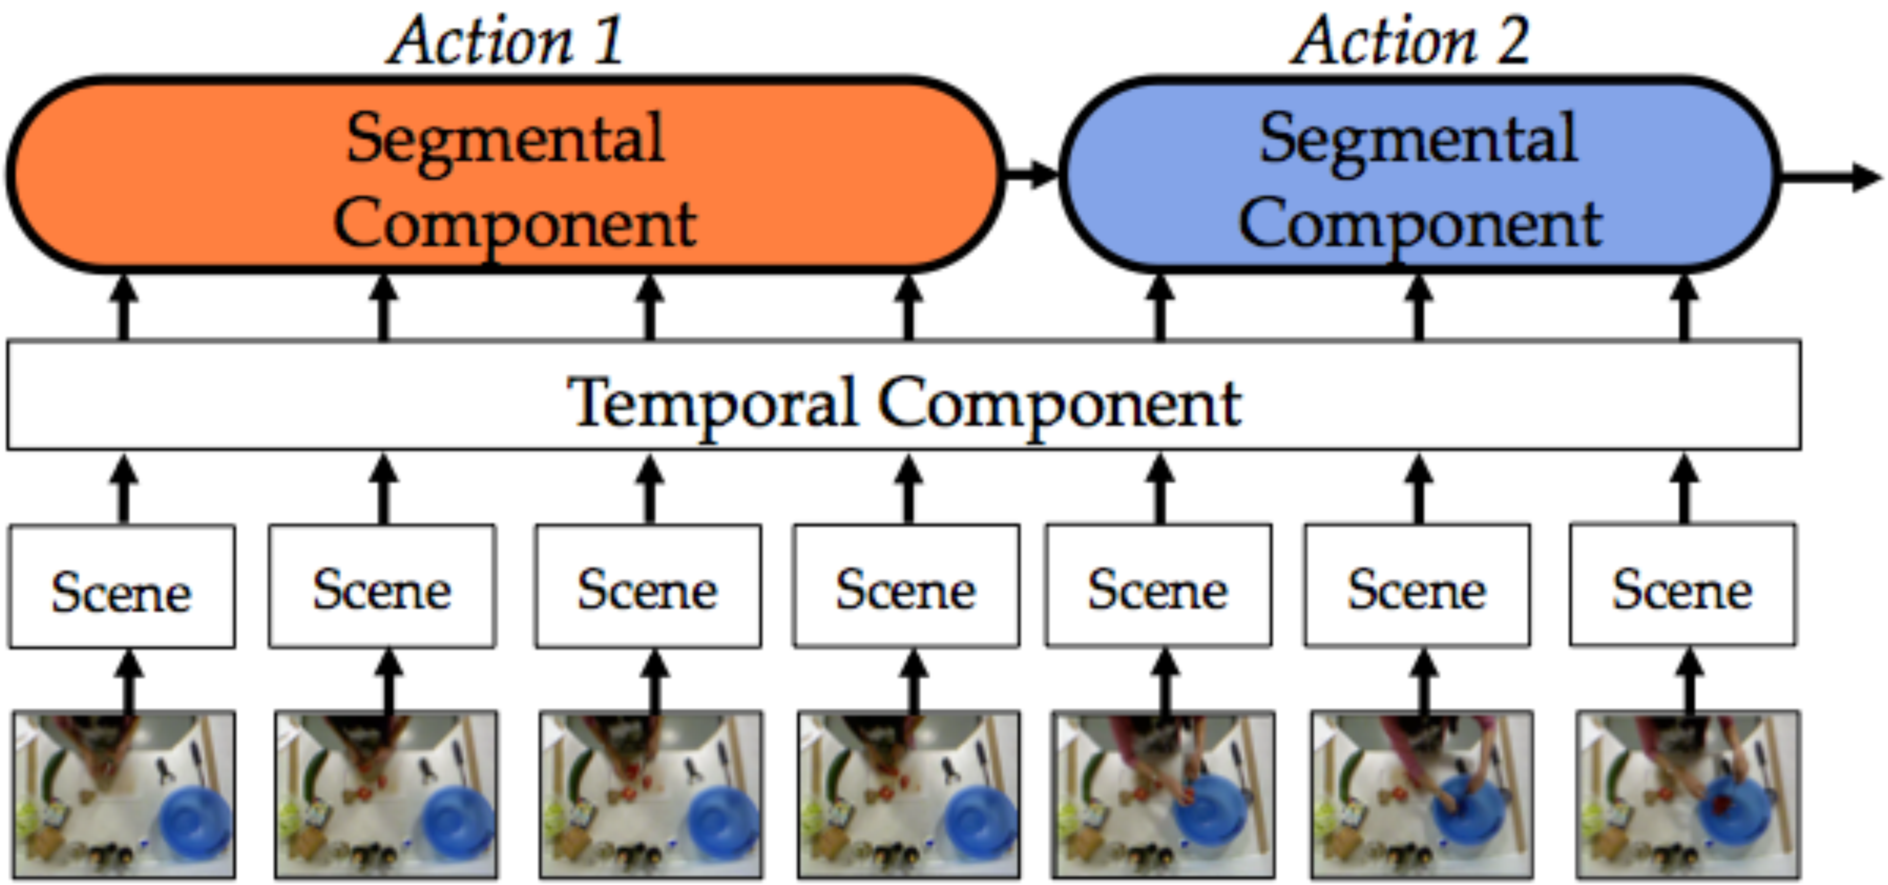

C. Lea, A. Reiter, R. Vidal, G. Hager. Segmental Spatiotemporal CNNs for Fine-grained Action Segmentation. ECCV 2016

Encoder-Decoder Temporal Conv Nets

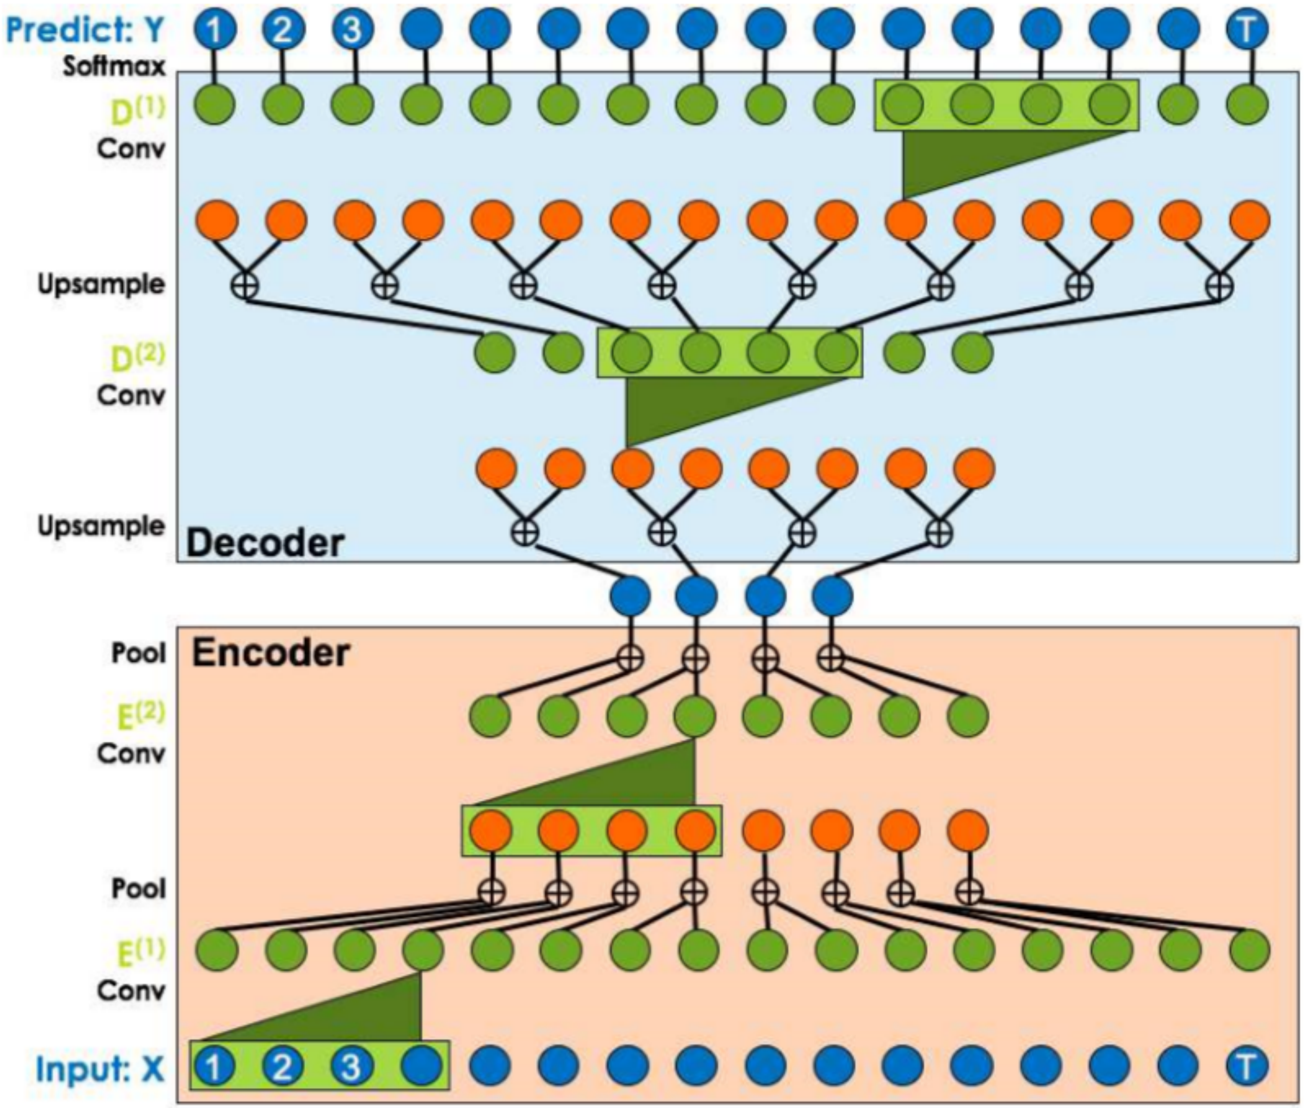

C. Lea, M. Flynn, R. Vidal, A. Reiter, G. Hager. Temporal Convolutional Networks for Action Segmentation and Detection. CVPR 2017

Accuracy of Surgical Gesture Segmentation

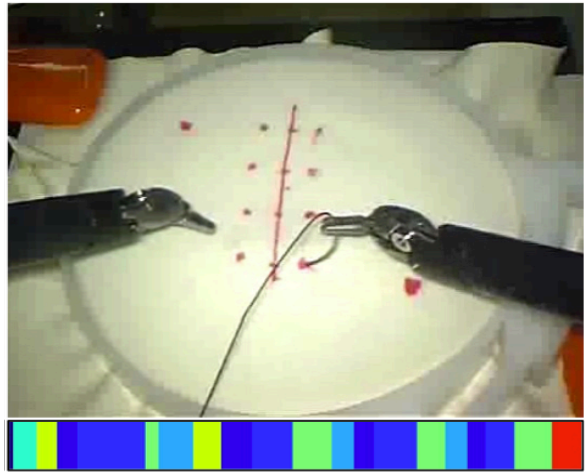

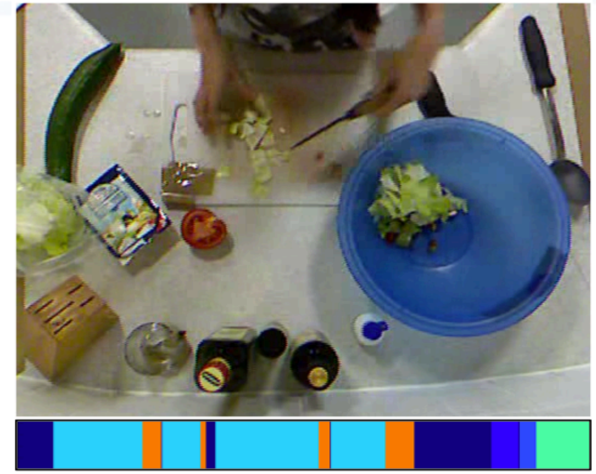

	Model	Accuracy	Edit
Sensors	LC-SC-CRF	81.9*	78.4*
	S-ST-CNN	79.2	82.6
	ED-TCN	82.4	**89.3**
Video	LC-SC-CRF	-	-
	S-ST-CNN	74.2	66.6
	ED-TCN	**78.3**	**85.6**

	Model	Accuracy	Edit
Sensors	LC-SC-CRF	81.8	58.5
	S-ST-CNN	82.1	55.5
	T-CNN	**84.8**	**76.9**
Video	LC-SC-CRF	-	-
	S-ST-CNN	**72.0**	**62.0**
	ED-TCN	71.0*	**62.0***

C. Lea, G. Hager, R. Vidal. An improved model for segmentation and recognition of fine-grained activities with application to surgical training tasks, WACV15
C. Lea, A. Reiter, R. Vidal, G. Hager. Segmental Spatiotemporal CNNs for Fine-grained Action Segmentation. ECCV 2016
C. Lea, M. Flynn, R. Vidal, A. Reiter, G. Hager. Temporal Convolutional Networks for Action Segmentation and Detection. CVPR 2017

Machine Learning in Rehabilitation Therapy

René Vidal
Herschel L. Seder Professor of Biomedical Engineering
Director, Mathematical Institute for Data Science
Johns Hopkins University

Machine Learning in Rehabilitation Therapy

Pediatric rehabilitation based on human-robot interaction

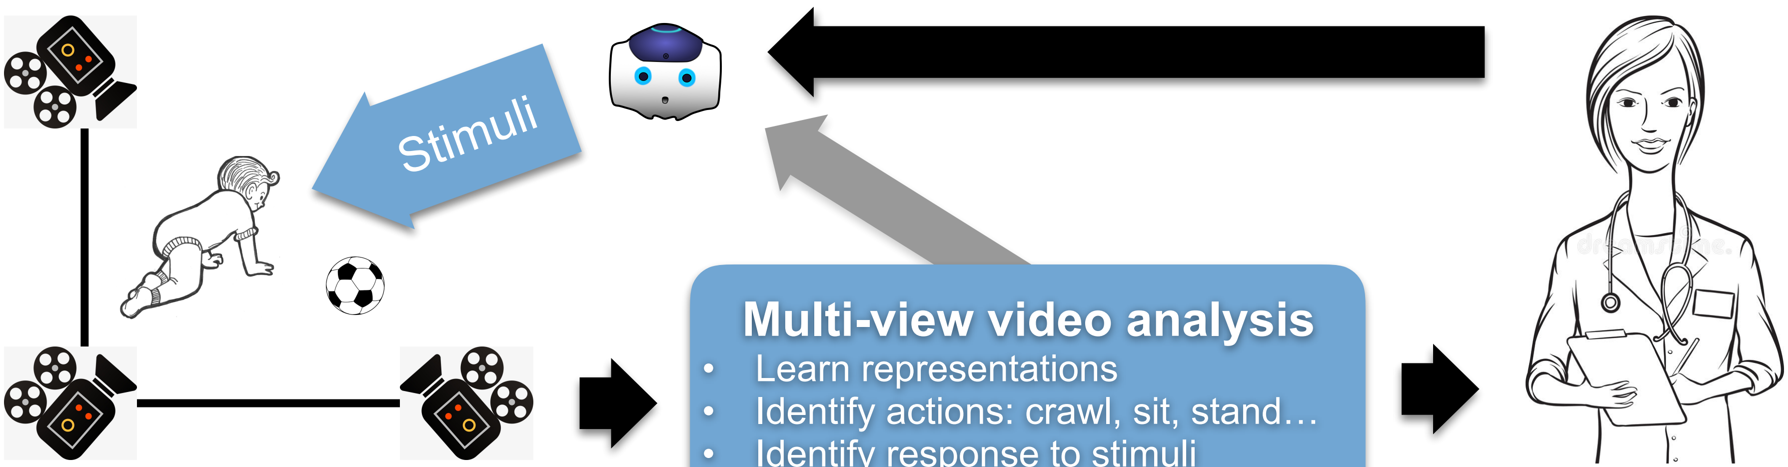

Multi-view video analysis
- Learn representations
- Identify actions: crawl, sit, stand…
- Identify response to stimuli
- Provide feedback to therapists

Effrosyni Mavroudi, Lingling Tao, Rene Vidal, Deep Moving Poselets for Video Based Action Recognition, WACV 2017
Lingling Tao and Rene Vidal, Moving Poselets: A Discriminative and Interpretable Skeletal Motion Representation for Action Recognition, ICCVW 2015

Summary

- **Cellular Level**
  - Dictionary learning for blood cell detection, classification, counting
  - Structured matrix factorization for segmentation of neural activity in calcium imaging
  - Metamorphosis for classification of embryonic cardio-myocytes
- **Organ Level**
  - Compressed sensing and Riemannian geometry for processing diffusion MRI
- **Patient/Surgeon Level**
  - Assessing surgical skill
  - Assessing children’s motions

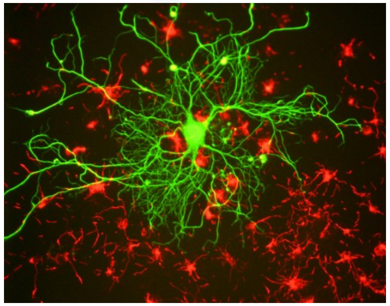

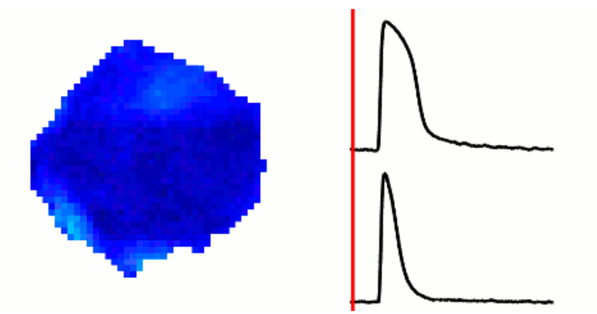

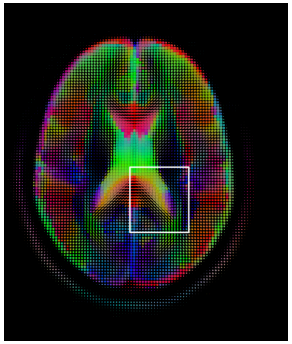

Ghoreyshi ISBI07; Singaraju CVPR08, CVPR09, TPAMI11; Goh ECCV06, ISBI06, ECCV08, CVPR09; MICCAI09, Neuroimage 12; Cetingul PPMIA09, ISBI09, TBME11, ISBI11, ISBI12, TBME14; Schwab IPMI13; Gorospe MICCAI13, TBME13; Haeffele ICML14; Tao IPCAI12, MICCAI13; Bejar-Zapella MICCAI12, Media13, MICCAI13; Lea: WACV15, ECCV16

Mathematical Institute for Data Science (MINDS)

- Establish fundamental principles behind the analysis and interpretation of massive amounts of complex high-dimensional data

- We are creating new
  - Masters
  - PhD

- We are hiring
  - Faculty
  - Postdocs
  - Students

More Information,

Vision Lab @ Johns Hopkins University

http://www.vision.jhu.edu

Center for Imaging Science @ Johns Hopkins University

http://www.cis.jhu.edu

Johns Hopkins Mathematical Institute for Data Science

http://www.minds.jhu.edu

Thank You!