


Connectomic Reconstruction with Machine Learning



Viren Jain

Research & Machine Intelligence, Google Inc.

- reconstructing neural wiring diagrams from 3d, nanoscale images of brain tissue

1. deep networks for boundary detection

2. metrics & cost functions

3. neurite agglomeration

“connectomics”





“connectomics”




Which neurons form synaptic connections with each other?


- ▶ Associate each dendrite and axon with a parent cell body
- ▶ Find synapses (connections) between particular cells





Synapse Identification



- Conventional image processing techniques and even modern 'off-the-shelf' machine learning techniques make too many mistakes







- How to increase accuracy?

Image Segmentation for Connectomics

- Image may contain trillions of voxels
- Linear time methods
- Need to decompose the problem

Local Boundary Prediction



Global Region Formation


“Sophisticated” Local Boundary Prediction




“Simple” Global Region Formation



input image



local boundaries



nonlocal
segmentation


Convolutional Networks

Multi-layer neural network architecture that can be trained by gradient learning to interpret images in a desired way.


Why Convolutional Networks?

- No prior research on best feature representations for EM data
- Not afraid to gather lots of training data

LeNet-5 Digit/Letter Recognition Network




Yann LeCun and others
Bell Labs (1980's)




Jain et al, ICCV 2007
Turaga et al, NC 2010
Turaga et al, NIPS 2009
Jain et al, CVPR 2010

“Connectomic reconstruction of the IPL in mouse retina”



Helmstaedter et al, Nature 2013

test set performance



Natural Image Denoising with Convolutional Networks
Jain & Seung, NIPS 2008

- Takes a long time to learn the parameters (even on a GPU)
- Computations involved in backpropagation make (further) parallelization difficult


- Exploit more powerful network architectures
- Explore larger parameter space of cost functions, representations, etc
- Quickly retrain based on additional labeled data

- Deep and Wide Multiscale Recursive Networks

Network Architecture “Desiderata”


- (1) wide feature representation
- (2) large field of view
- (3) model statistical structure in label-space

(1) Narrow vs Wide Feature Representation




node = feature map
arrow = filter

(1) Narrow vs Wide Feature Representation



(1) Narrow vs Wide Feature Representation



Filter/Receptive Field/Projection/Feature

Random

Hand engineered

Brute-force search

Supervised Learning

Unsupervised Learning (kmeans, omp, autoencoders)

unsupervised

input

node = feature map
arrow = filter

x_1

x_2

x_3

x_N

y_1

y_2

y_3





y_N

z_1

z_2

z_3

z_N




supervised

mlp_1


mlp_2

output



(2) Field of View


input




classifier



output




(2) Field of View




- Number of layers

- Size of filters

node = feature map
arrow = filter




Foveated Pooling



- (3) Model Statistical Structure in Label-Space

(3) Model Statistical Structure in Label-Space



(3) Model Statistical Structure in Label-Space

- Labels for neighboring image locations are correlated (i.e., structured prediction).
- In supervised image labeling approaches, predictions for neighboring locations typically only correlated through dependence on overlapping input image patches.


(3) Model Statistical Structure in Label-Space

- (Bayesian) MRF approach: model $p(Y)$ and $P(X|Y)$
- Computational expense of probabilistic learning & inference can limit model complexity of MRFs in important ways (Jain 2007, 2009).

unsupervised

input

node = feature map
arrow = filter





supervised

mlp_1



mlp_2


output









DAWMR Architecture






- distributed CPU feature extraction takes advantage of parallel file system for accessing raw image data of training examples spread across large image space
- distributed CPU feature extraction stage allows expensive operations (sparse coding, multiscale downsampling, etc) without having to optimize implementation for GPUs
- GPU stage focused on simple and relatively “stable” multilayer perceptron computations





recursive processing




image



affinity graph




segmented
affinity graph



DAWMR
iteration 1




DAWMR
iteration 2



DAWMR
iteration 3

Boundary Prediction Generalization Performance (AUC)



Huang & Jain, In Review





Figure 5: Results on the test set. For Rand Index, x -axis is scaled by the number of clusters in segmentations generated at various binarization thresholds (the specific thresholds are chosen custom to each methods output, in order to yield one-thousand evenly spaced quantiles from the analog affinity graph values).

Rand Index:

$$R(S, T) = \frac{1}{\binom{n}{2}} \sum_{i, j \neq i} \delta(S_i = S_j, T_i = T_j) + \delta(S_i \neq S_j, T_i \neq T_j)$$

Huang & Jain, In Review

1. deep networks for boundary detection

2. metrics & cost functions

3. neurite agglomeration

- What to measure?

- How to optimize?




- Pixel-wise boundary classification accuracy not directly related to ultimate practical goals involving reconstruction accuracy.

Segmentation*

Labeling*

Image

Ground Truth






Segmentation*

Labeling*

Image



Ground Truth





Segmentation A

Labeling A

Interpretation A







Rand Index:


$$R(S, T) = \frac{1}{\binom{n}{2}} \sum_{i, j \neq i} \delta(S_i = S_j, T_i = T_j) + \delta(S_i \neq S_j, T_i \neq T_j)$$

Maximin Affinity Learning of Image Segmentation (Turaga et al., NIPS 2010)

Topology

- Properties of objects (topological spaces) that are preserved under certain transformations (like bending and twisting)
- Exact shape is not always important to define certain (very important) mathematical properties

topological equivalence



Digital Topology

- Formalizes continuous topological concepts in the digital plane
- Defines connected objects by discrete adjacency relationships
- Yields a rigorous connection between a pixel, a binary labeling, and a segmentation


Warping Error

- ▶ Given images L^* and T ,
 L is a warping of L^* onto T , $L \triangleleft (L^*)_T$, if
 1. L can be obtained by a sequence of topology preserving operations on L^* .
 2. All differences between L and T , $E(L, T)$, occur at non-simple points of L .

Warping Error

The warping error between L^* and T is:


$$D(T||L^*) = \min_{L \triangleleft (L^*)_T} |E(L, T)|$$



dual-optimization with smooth approximation:

$$\min_{\vec{w}} \min_{L \triangleleft (L^*)_{F_I(\vec{w})}} \sum_i [l_i - f_i(\vec{w})]^2$$

alternate





$$\min_{\vec{w}} \sum_i [l_i - f_i(\vec{w})]^2$$

gradient descent

$$\min_{L \triangleleft (L^*)_{F_I(\vec{w})}} \sum_i [l_i - f_i(\vec{w})]^2$$

greedy descent


performance on the test set




1. deep networks for boundary detection

2. metrics & cost functions


3. neurite agglomeration






- ▶ Agglomeration as process determined by the “policy” of a deterministic markov decision process (DMDP)
- ▶ State space: segmentation given by current clusters
Action space: agglomeration of any two clusters OR
query a human OR re-image sample
Objective function: error at end of agglomeration process
according to a specified error metric with a given cost
- ▶ Use techniques from on-policy control in RL to train
function that approximates optimal policy

“Learning to Agglomerate Superpixel Hierarchies”
Jain et al, NIPS 2011



negative



positive

SHREC'12 Track: Generic 3D Shape Retrieval

B. Li^{1,6}, A. Godil¹, M. Aono², X. Bai³, T. Furuya⁴, L. Li³, R. López-Sastre⁵, H. Johan⁶, R. Ohbuchi⁷,
C. Redondo-Cabrera⁵, A. Tatsuma², T. Yanagimachi⁷, S. Zhang³

¹ National Institute of Standards and Technology, Gaithersburg, USA ² Toyohashi University of Technology, Japan

³ Northwestern Polytechnical University, Xi'an, China ⁴ Nisca Corp., Yamanashi, Japan

⁵ Department of Signal Theory and Communications, University of Alcalá, Spain

⁶ School of Computer Engineering, Nanyang Technological University, Singapore ⁷ University of Yamanashi, Yamanashi, Japan





Figure 2: *Feature extraction process.*



$$= [1 \ 0 \ 1 \ 0 \ \dots \ 1 \ 1 \ 1 \ 0 \ 1] \\ \{0,1\}^N$$

↓ dimensionality
reduction

$$[x_1 \ x_2 \ \dots \ x_k] \\ \mathbb{R}^k \quad (k \ll N)$$

Object-Level
Feature Design




```
graph TD; A[Object-Level Feature Design] --> B["'hand-designed' 3d geometry & morphology descriptors"]; A --> C[unsupervised & supervised machine learning];
```

"hand-designed"
3d geometry &
morphology
descriptors

unsupervised &
supervised
machine
learning



Object-Level
Feature Design

“hand-designed”
3d geometry &
morphology
descriptors




	Exp.	Feature Set Description	Training set		Testing set		Dim.	Cost
			ACC(%)	AUC(%)	ACC(%)	AUC(%)		
Boundary Map (bm)	1	bm mean, median, interface len.	83.64	92.44	80.92	90.46	6	5.7
	2	exp 1 + bm moments	85.27	93.39	82.68	91.64	9	5.7
	3	exp 1 + bm quantiles	84.54	93.11	82.04	91.33	8	5.7
	4	exp 1 + bm quantiles, min/max	85.03	93.36	82.69	91.63	10	7.9
	5	exp 1 + bm deriv. mean, median	90.31	96.71	88.64	95.73	14	7.9
	6	exp 1 + all bm deriv. stats	91.85	97.61	89.11	96.05	42	14.0
	7	baseline ($\bigcup$ exp 1:6)	92.30	97.85	89.41	96.05	49	14.0
	8	baseline ($\bigcup$ exp 1:6) boosting	92.17	97.88	88.56	95.36	49	-
Object	9	exp 7 + growth	92.55	98.09	89.67	96.32	51	1.0
	10	exp 7 + proximity	92.18	97.85	89.09	96.02	50	489.1
	11	exp 7 + angles	95.74	99.25	89.65	96.27	82	13.0
	12	exp 7 + size	93.31	98.43	90.28	96.61	53	1.9
	13	exp 7 + rays	94.36	98.92	90.06	96.52	91	44.2
	14	exp 7 + shape diam. quantiles	93.52	98.56	89.82	96.46	59	402.5
	15	exp 7 + shape diam. moments	94.26	98.72	86.32	92.14	57	402.5
	16	exp 7 + shape context	94.71	99.03	89.91	96.50	69	5.5
	17	exp 7 + convex hull	93.47	98.56	89.97	96.74	57	8.7
	18	exp 7 + level sets overlap	92.23	97.90	89.16	96.08	55	464.0
	19	exp 7 + level sets gradient v.f.	93.20	98.28	89.59	96.17	53	35.0
	20	exp 7 + level sets orientation	93.74	98.61	90.13	96.75	55	229.4
Image	21	exp 7 + SIFT soft v.q.	99.04	99.93	88.75	95.67	149	56.0
	22	exp 7 + image moments	93.16	98.29	89.26	96.12	53	4.1
	23	exp 7 + image deriv. stats	95.58	99.22	89.09	95.61	85	5.8
	24	exp 7 + image stats	96.42	99.43	88.85	95.73	94	5.8
	25	all hand-designed ($\bigcup$ exp 1:24)	99.98	99.98	92.33	97.61	363	-

Object-Level Feature Design




unsupervised &
supervised
machine
learning


Precision vs Recall (Testing Set)






Object-Level Feature Design






unsupervised &
supervised
machine
learning



Dynamic
Object
pooling




Dynamic
Boundary
pooling






Segments




Pooling region

Rendering



$$R(S, T) = \frac{1}{\binom{n}{2}} \sum_{i, j \neq i} \delta(S_i = S_j, T_i = T_j) + \delta(S_i \neq S_j, T_i \neq T_j)$$





Next few years...

- ▶ Allen Brain Institute: mouse visual cortex
- ▶ 1 cubic mm at $5 \times 5 \times 30 \text{ nm}$ $\sim$ 1 petabyte image
- ▶ Combined functional and anatomical data.










Whole Mouse Brain?








- ▶ Multibeam SEM: 100 TB/day
- ▶ Entire brain: 50-100 petabytes of raw image data

Google: Connectomic Reconstruction

- ▶ Scalable infrastructure for petabyte scale image storage and manipulation.
- ▶ Scalable image alignment.
- ▶ Deep networks for automated image annotation and reconstruction.
- ▶ Services provided to the scientific community via Google cloud infrastructure.

- analysis of wiring diagram structure





Acknowledgements

Janelia

Goran Ceric
Mitya Chklovskii
Tom Dolafi
Corey Fisher
Reed George
Harald Hess
Gary Huang
Tzumin Lee
Rob Lines
Zhiyuan Lu

Ian Meinertzhagen
Chris Ordish
Ankur Patel
Stephen Plaza
Pat Rivlin
Lou Scheffer
Mat Saunders
Jackie Stephens
Greg Tumolo
Shan Xu
Charlotte Weaver

MIT

Sebastian Seung

UCL

Srinivas Turaga




MPI


Moritz Helmstaedter
Winfried Denk

NIH



Kevin Briggman

- analysis of wiring diagram structure






$N=1$




a) Sequenced DNA fragments


















b) Comparing similar regions



c) Searching for high similarity







- How to analyze and compare neural wiring diagrams to each other?

