2D HIDDEN MARKOV MODEL WITH SPATIALLY ADAPTIVE STATE-SPACE FOR TRACING MANY CELLS IN IMAGE SEQUENCE

Min-Chi Shih, *Renuka, Shenoy,* *Kenneth Rose*

Department of Electrical and Computer Engineering
University of California, Santa Barbara
{minchi_shih, renuka, rose}@ece.ucsb.edu

ABSTRACT

In this paper we propose a two-dimensional hidden Markov model (HMM)-based framework for solving the cell tracing problem in a biological image sequence. Given label initialization in the first frame, we model the problem as pixel labeling for every consequent frame. Common Markov random field-based frameworks for this task require a fixed set of labels $S = \{1, 2, \cdots, L\}$, while in our framework the set of labels or the state-space is spatially adaptive, i.e., available prior information is exploited to identify a smaller state-space that varies from node to node. In the cell tracing problem, specifically, temporal information on cell location in the previous frame is used to reduce the states to a small subset of the complete label set. The substantial reduction in average cardinality of the label set yields benefits not only in terms of computational complexity, but also in the labeling accuracy. The general idea can be broadly applied to many computer vision and image processing problems, where prior knowledge enables local reduction of the state-space. We consider the cell tracing problem on a publicly available challenging biological image dataset, which contains a series of electron microscopy images of high resolution and a large number of objects (neuronal processes) to be traced. Experimental results compare the approach with other recently proposed methods, and show considerable improvement.

***Index Terms*—** Cell tracing, electron microscopy, 2D-hidden Markov model

1. INTRODUCTION

Neuroscientists have been imaging very large volume of neural tissue for the purpose of reconstructing the neuronal circuitry. The identification of post synaptic densities is crucial for this task, and currently electron microscopy (EM) is the only imaging technique that can provide sufficient resolution. Recent advances in sample preparation and imaging processes, including mosaicking and registration, have enabled the acquisition of multi-terabyte data volume. To acquire complete cell and network maps, manual analysis is labor intensive and error prone even by experienced neuroanatomists, and, therefore, automatic high-throughput techniques are essential for the tasks of reconstruction and evaluation.

In this paper, we consider the problem of tracing many cells in a biological image sequence. Since cell 3D reconstruction bears similarity to video tracking, it is intuitive to assign each cell with a unique label and perform pixel-wise label propagation from frame to frame. By assigning a label to each pixel in every frame, the locations and shapes of all cell structures are determined. In this way, the topological changes such as merging, splitting, and overlapping instances are naturally handled. Such approach usually involves a Markov random fields (MRF), a Markov mesh random field (MMRF), or a conditional random field (CRF), where the labeling result can be solved by belief propagation [1], two-dimensional hidden Markov model (2D-HMM) with turbo-decoding [2] and graph cuts with α-expansion [3], respectively. These methods, employing appropriate energy functions and model parameters, are proven to provide accurate labeling results, e.g. the methods of [4] applied a CRF framework with high-order potentials to the multiple cell tracing problem. However, with increase in the size of the image and the label space, the computational complexity limits the usability of these methods.

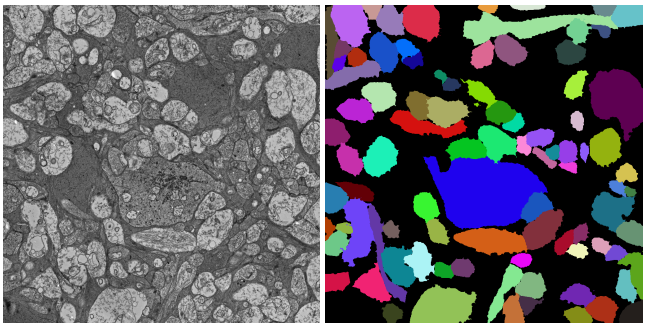

(a) (b)

Fig. 1. (Best in color) (a): an electron microscopy image, (b): the pixel labeling result using the proposed 2D-HMM-SASS framework.

An alternative approach is to convert the atomic units from pixels to superpixels using various over-segmentation methods. This reduces the number of nodes in a graph by at least two orders, and thus dramatically decreases the computational complexity. However, the labeling accuracy heavily depends on the initial over-segmentation. Errors due to the superpixels, i.e. boundary inconsistency between real objects and superpixels, will be propagated to the final labeling results and, typically, there is no correction provision in such formulation. Consequently, superpixel-based methods often yield worse performance than the MRF-like methods but at lower complexity. In this paper we propose a MMRF-based framework, 2D-HMM with spatially adaptive state-space (2D-HMM-SASS), where the set of labels for each node is determined by available prior information on the local area. It offers the capability of handling a large number of

This work was supported by two grants from National Science Foundation, NSF OIA under grant 0941717 and NSF III under grant 0808772. The authors thank Dr. Robert Marc, Dr. Bryan Jones and Dr. James Anderson from the Univ. of Utah for providing data used in experiments.

labels at reasonable complexity without incurring the penalty of superpixel over-segmentation. The idea of SASS could be utilized in many different pixel labeling problems in computer vision. It is particularly suitable for the problem of multiple cell tracing in biological images with a large number of objects of interest. Experimental results show that the proposed framework is of low computational complexity due to the smaller label spaces, and yields better pixel labeling performance due to minimized confusion per node. The states of hidden Markov model are equivalent to the labels in a MRF-based framework, and we use these terms interchangeably in the sequel.

2. RELATED WORK

There exists substantial prior work on 3D model reconstruction from biological image sequences. Existing methods usually formulate the 3D reconstruction problem as one of grouping superpixels or supervoxels using various graphical models. Graph-based segmentation is utilized in [5] to obtain a 2D over-segmentation of every section, and a hypergraph framework is designed to efficiently solve the grouping problem. In [6], a probability map of membrane detection is learned by means of a random forest, and watershed segmentation is performed on top of the probability map to produce supervoxels that are then further grouped to reconstruct the 3D model. In [7], a set of supervoxels is obtained by the SLIC algorithm and then grouped by graph cuts [8].

Other methods assume that all neuronal processes can be segmented well in every frame and then connect those segments in the vertical direction. In [9], a simple thresholding and watershed segmentation is performed on each frame to acquire 2D segments of all neurons. Then a directed graph is constructed and Dijkstra's algorithm is used to find the optimal connectivity for each neuron in the first section. In [10], the intra-section and inter-section segmentation are done in separate procedures, where a merge tree with superpixels as leaf nodes was used to reconstruct the 2D segmentation, and a section classifier was trained for the inter-section neuron reconstruction. In [11], a deep neural network is utilized to learn the neuronal membranes in EM images. In [12], a framework was proposed for tracing neuronal processes over serial sections by using graph-cut optimization over the 3D volume. While using the directional energy in the graph cuts boosts the 2D segmentation performance, this method fails if the structures of interest are not orthogonal to the cutting direction, due to the simple 3D linkage energy function employed. In [13], an active semi-supervised learning method was proposed for EM segmentation. The above 2D-segmentation-based methods perform well on datasets where cell membranes can be easily segmented, but, to the best of our knowledge, no single segmentation method works effectively on the challenging dataset we have, in the presence of highly cluttered background and intracellular structures.

Other than superpixels or supervoxel-based grouping methods, in [14] and [15], a set of over-segmentation results are generated as multiple hypothesis, and graph cuts are utilized for optimizing the merging results. In this paper, we propose a pixel-wise labeling framework, 2D-HMM-SASS, to effectively and efficiently solve the problem of cell tracing on a highly challenging EM image sequence of a rabbit retina.

3. BACKGROUND

3.1. Turbo 2D-HMM

The hidden Markov model has been known to provide optimal solutions for one-dimensional multiple-labeling problems, leveraging the efficient forward-backward algorithm or Viterbi decoding. However, its direct extension to two-dimension is impractical because the complexity grows exponentially with the image size. There are several approximating variants of 2D-HMM algorithms [2], [16] and [17], of which the turbo 2D-HMM [2] is an effective and efficient algorithm that has been applied to a number of computer vision problems [18] [19] [20]. With a modified version of the forward-backward algorithm, the turbo 2D-HMM iteratively decodes each row and column independently as a 1D-HMM, but allows them to "*communicate*" by inducing priors on each other. More specifically, the column-wise occupancy probabilities from the previous iteration are utilized as prior information in current row-wise decoding, and vice versa.

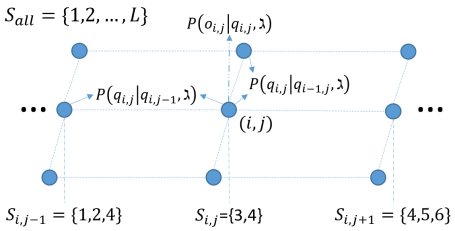

Fig. 2. Illustration of a 2D-HMM with *spatially adaptive states*. Similar to standard 2D-HMM, an observation probability $P(o_{i,j}|q_{i,j},\lambda)$ is emitted at each node (i,j), and the (horizontal) transition probability $P(q_{i,j}|q_{i,j-1},\lambda)$ is defined over nodes (i,j) and $(i,j-1)$. However, the state-space is spatially adaptive and the cardinality is often reduced to a small subset of the full set S_{all}.

Let $O = \{o_{i,j}\}$ and $Q = \{q_{i,j}\}$ be the observations and the states to be predicted of all pixels, and λ be the model parameters. Using the Markovian property, the joint likelihood of all nodes is

$$\begin{aligned} P(O,Q|\lambda) &= P(O|Q,\lambda)P(Q|\lambda) \\ &= \prod_{i,j} P(o_{i,j}|q_{i,j},\lambda)P(q_{i,j}|q_{i-1,j},q_{i,j-1},\lambda), \end{aligned} \tag{1}$$

where $P(o_{i,j}|q_{i,j},\lambda)$ and $P(q_{i,j}|q_{i-1,j},q_{i,j-1},\lambda)$ are observation and transition probabilities respectively. The observation probabilities are parallel to the unary potentials, and the transition probabilities are parallel to the binary potentials in a Markov random field formulation. For clarity, we will omit λ in the following equations. It is assumed that the transition probabilities can be decomposed into the vertical and horizontal components, i.e.

$$P(q_{i,j}|q_{i-1,j},q_{i,j-1}) \propto P(q_{i,j}|q_{i-1,j})P(q_{i,j}|q_{i,j-1}). \tag{2}$$

Let the observations and the state sequence of row i and column j be denoted by o_j^V, o_i^H, q_j^V, q_i^H, where the superscripts V and H denote vertical and horizontal, respectively. The turbo 2D-HMM approximates the overall joint likelihood to be the product of all row-wise (column-wise) joint likelihoods accompanied with some priors from column-wise (row-wise) decoding results, i.e.

$$\begin{aligned} P(O,Q|\lambda) &\approx \prod_j \left[P(o_j^V, q_j^V) \prod_i P(q_{i,j}|o_i^H) \right] \\ &\approx \prod_i \left[P(o_i^H, q_i^H) \prod_j P(q_{i,j}|o_j^V) \right]. \end{aligned}$$

As a result, the optimization can be solved by iteratively decoding the columns and rows, and the computational complexity of turbo

2D-HMM is $O(n|S|^2)$, where n is the number of nodes and $|S|$ the number of all possible labels.

4. PROPOSED METHOD

4.1. 2D-HMM with Spatially Adaptive States

A "standard" 2D-HMM is spatially invariant, i.e., the set of states $S = \{1, 2, \cdots, L\}$ is fixed for all nodes. The proposed 2D-HMM framework, on the other hand, allows each node to have its own label candidates, i.e., the set of allowable states for each node is varied adaptively accounting for the available prior information on each location, as shown in Fig. 2, where the prior could be extracted from temporal, spatial, or other source of information. We will show that, with the appropriate prior, the state-space at each node could be reduced to a small subset of the full set S. The idea of SASS could be utilized in many different pixel labeling problems in computer vision, such as optic flow, scene analysis, and video tracking. Specifically, it is suitable for the problem of multiple cell tracing in biological images, which has very high resolution (a few nanometers per pixel), and a large number of objects of interest.

A typical 8192×8192 image frame in the electron microscopy (EM) image sequence of our dataset [21] contains around 100 neuronal processes. For each node, we could naively use all 100 unique labels to form the state-space. However, with the temporal prior of the locations of all cells in the previous frame and an assumption that the maximum displacement of any cell is within d_{max} pixels, the size of the state-space could be reduced to less than 3 on an average. The necessary cell label candidates of pixel (i, j, t) in frame t are the set of unique labels of those in a $(2d_{max}+1) \times (2d_{max}+1)$ bounding box centered at $(i, j, t-1)$. We illustrate how the label candidates are extracted at each pixel with Fig. 3 (a). Let n be the number of pixels and S be the set of all labels, the computational complexity of turbo HMM is $O(n|S|^2)$, while that of 2D-HMM-SASS is approximately $O(nC_b^2)$, where C_b is the average number of unique label candidates in a bounding box. Since the cardinality of S is proportional to n and C_b is a constant, the proposed algorithm essentially reduces the complexity from $O(n^3)$ to $O(n)$. Note that the efficiency gain increases with the cardinality of the full label set. In addition to the reduction in computational complexity, the pixel labeling precision also benefits from the proposed SASS idea, since only the necessary subset of labels is considered at every pixel and there is less risk of confusion, compared to the original framework. Given all the above features of 2D-HMM-SASS, it is suitable for the problem of multiple cell tracing in biological images.

4.2. Observation Probabilities

The observation probabilities in a hidden Markov model are analogous to the unary potentials in a MRF model. In our proposed framework, we extract three local features at every pixel, the histogram of pixel intensity, local binary patterns, and distance transform. Given the segmentation in frame $t-1$ (manual initialization of the pixel-wise labels of all cells in the first frame), we compute $H_{pi}^{t-1}(l)$, the histograms of pixel intensity of cell l in frame $t-1$, and similarly $H_{lbp}^{t-1}(l)$ for local binary patterns. The histograms are normalized so that they all sum to one, hence representing probability distributions. We extract local histograms of pixel intensity $F_{pi}^{t}(i,j)$ and local binary patterns $F_{lbp}^{t}(i,j)$ with a bounding box of size $B \times B$ at pixel (i, j) in frame t. Let $D_{KL}(P||Q) = \sum_i P(i) \ln \frac{P(i)}{Q(i)}$ be the asymmetric KullbackLeibler (KL) divergence of two probability distributions P and Q. To compute the "*distance*" between each pixel and a label l, we employ a symmetric version of the KL divergence between the local feature and the histogram of the target label, i.e., $d_{pi}^{l}(i,j) = D_{KL}^{sym}\left(F_{pi}^{t}(i,j) \middle\| H_{pi}^{t-1}(l)\right)$ and $d_{lbp}^{l}(i,j) = D_{KL}^{sym}\left(F_{lbp}^{t}(i,j) \middle\| H_{lbp}^{t-1}(l)\right)$, where $D_{KL}^{sym}(P||Q) = D_{KL}(P||Q) + D_{KL}(Q||P)$ is the symmetric KL divergence. Also, we compute $d_{dt}^{l}(i,j)$, the minimum distance from point (i,j) to the target cell l in the previous frame using the efficient distance transform algorithm [22]. Each feature is assumed to be drawn from an independent exponential distribution, and the observation probability is $P(o_{i,j}|q = l) = P_1\left(d_{pi}^{l}(i,j)\right) \cdot P_2\left(d_{lbp}^{l}(i,j)\right) \cdot P_3\left(d_{dt}^{l}(i,j)\right)$, where $P_i(d) = \lambda_i \exp\left(-\lambda_i d\right)$ is the exponential probability density function.

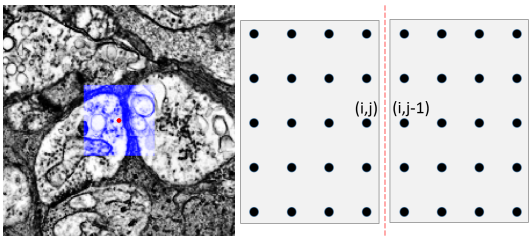

(a) (b)

Fig. 3. (Best in color) (a) an illustration of how label candidates are extracted at each pixel, (b) an illustration of how the transition probabilities from pixel (i, j) to $(i, j+1)$ is computed, we consider the regional difference of both sides..

4.3. Transition Probabilities

The transition probabilities, which are analogous to the binary potentials, control the smoothness of the result. For each pixel (i, j) and its immediate horizontal neighbor $(i, j+1)$, we compute the regional difference of both sides, as shown in Fig. 3 (b). Let $d_{i,j}^{tr,h}$ be the horizontal regional difference between the left and right $k \times k$ pixels, where we forge the simplistic measures, such as the difference of average pixel intensities or the weighted directional derivatives, and use a refined distribution-based measure. Let H_{pi}^{prev} and H_{pi}^{post} be the normalized histograms of the pixel intensities on both sides, and we compute $d_{i,j}^{tr,h} = D_{KL}^{sym}\left(H_{pi}^{prev}, H_{pi}^{post}\right)$. Since we assume the pixel labeling result should be piecewise constant and drastic changes are allowed on the boundaries, the Potts model is applied to the transition probabilities in our framework, i.e.

$$P(q_{i,j+1}|q_{i,j}) = \begin{cases} 0 & : q_{i,j+1} = q_{i,j} \\ \lambda^{tr,h} \exp\left(-\lambda^{tr,h} r(d_{i,j}^{tr,h})\right) & : q_{i,j+1} \neq q_{i,j} \end{cases},$$

where $\lambda^{tr,h}$ is the parameter of an exponential distribution, and $r(d) = -\log\left(1 - \exp(-d)\right)$ is a reciprocal-like function. In this setting, a large $d_{i,j}^{tr,h}$, which shows strong difference of the two regions, will result in a high transition probability from $q_{i.j}$ to $q_{i.j+1}$ if $q_{i.j} \neq q_{i.j+1}$.

5. EXPERIMENTAL VALIDATION

5.1. Datasets and Implementation Details

In this section we demonstrate the capability of the proposed 2D-HMM-SASS for tracing neuronal cells in a challenging electron microscopy image dataset [21], which is in the inner nuclear layer of

Table 1. 2D-HMM Run-time with Multi-threading

No. threads	1	2	4	8
Run time	2.33	1.78	1.39	1.13

Table 2. F-measure and Run-time Comparison

Method	Median	Mean	Std	Time(sec.)
Level Set Tracking [23]	0.55	0.53	0.25	1080
HGraph [5]	0.71	0.78	0.22	34
Graph Cuts [25]	0.67	0.77	0.28	1320
2D-HMM-SASS	**0.83**	**0.81**	**0.12**	**9.3 (4.5)**
2D-HMM-SASS (refine)	**0.87**	**0.84**	**0.12**	32

a rabbit retina. The original imaging resolution is 2.17 nanometers per pixel, and there are 95 neuronal processes to be traced in an 8192×8192 image. The proposed pixel-labeling-based algorithm is tailored for this kind of dataset where the cell profiles, which including brightness intensities and textures, are distinct. We down-sample the image to 512×512 and compare the performance of the proposed framework with the sparse field level set method [23] with the implementation of [24], Graph Cuts-based with higher order potentials [25], and a recent hypergraph-based method [5], where the runtime and pixel labeling accuracy performance of these methods are reported in [5]. With the temporal prior and the assumption that the maximum displacement of a cell d_{max} is 30 pixels, the maximum and average number of states of each node is only 6 and 2.12 respectively. Since the computational complexity of the original turbo 2D-HMM is $O(n|S|^2)$, compared with the original 95 labels, the proposed method reduce the computations by a factor of over 2000.

5.2. Parallel Computing for speed-up

Our proposed 2D-HMM-SASS is based on the turbo 2D-HMM, where the decoding of each column and row is independent. We implemented a C++ code with parallel computation supported by STL multi-threading. With an Intel i7-930 2.8GHz CPU with 4 cores and 8 threads, the run time for computing a 1024×1024 image reduces more than 50%. The run time comparison is shown in Table 1.

5.3. Parameter Learning

An advantage of utilizing a hidden Markov model-based method is its ability to learn the model parameters automatically for both the observation probabilities and transition probabilities. Let us use the notation listed in Table 4. The parameters update by the re-estimation formulas can be derived by maximizing the Baum's auxiliary function $Q(\lambda, \bar{\lambda})$ over $\bar{\lambda}$, where

$$Q(\lambda, \bar{\lambda}) = \sum_{\mathbf{Q}} P(\mathbf{Q}|\mathbf{O}, \lambda) \log[P(\mathbf{O}, \mathbf{Q}|\bar{\lambda})]. \quad (3)$$

This re-estimation procedure can be interpreted as an implementation of the Expectation-Maximization (EM) algorithm, and it can be proved that the likelihood function monotonically increases with the parameter updates. From (1), (2) and (3), we have $Q(\lambda, \bar{\lambda}) = Q(\lambda_b, \mathbf{b}) + Q(\lambda_a, \mathbf{a})$, where

$$Q(\lambda_b, \mathbf{b}) = \sum_{\mathbf{Q}} \sum_{i,j} \gamma_{i,j}(q_{i,j}) \log \left[b_{q_{i,j}}(o_{i,j}) \right]$$

$$Q(\lambda_a, \mathbf{a}) = \sum_{\mathbf{Q}} \sum_{i,j} \xi^V_{q_{i,j}} \log \left(a^V_{q_{i,j}, q_{i-1,j}} \right) + \xi^H_{q_{i,j}} \log \left(a^H_{q_{i,j}, q_{i,j-1}} \right).$$

Table 3. Rand Index Comparison

Method\Frame	1	2	3	4	5
Level Set [23]	0.82	0.78	0.75	0.72	0.69
HGraph [5]	0.88	0.86	0.85	0.83	0.80
Graph Cuts [25]	-	0.81	0.85	0.84	0.80
2D-HMM-SASS	-	**0.93**	**0.91**	**0.89**	**0.87**
2D-HMM-SASS (refine)	-	**0.94**	**0.93**	**0.92**	**0.92**

Table 4. 2D-HMM notation summary

$b_{q_{i,j}}(o_{i,j})$	$P(o_{i,j}\|q_{i,j})$
$a^V_{q_{i,j},q_{i-1,j}}$	$P(q_{i,j}\|q_{i-1,j}, \lambda^{tr,v})$
$\gamma^H_{i,j}(q_{i,j})$	$P(q_{i,j}\|\mathbf{o}^H_i, \lambda_b)$
$\gamma_{i,j}(q_{i,j})$	$(\gamma^H_{i,j}(q_{i,j}) + \gamma^V_{i,j}(q_{i,j}))/2$
$\xi^H_{q_{i,j}}$	$P(q_{i,j}, q_{i,j-1}\|\mathbf{O}, \lambda)$

We can apply maximum likelihood estimation on $Q(\lambda_b, \mathbf{b})$ to obtain the updated parameters λ_{pi}, λ_{lbp} and λ_{dt}, e.g., $\lambda_{pi} = \frac{\sum_{\mathbf{Q}} \sum_{i,j} \gamma_{i,j}(q_{i,j})}{\sum_{\mathbf{Q}} \sum_{i,j} \gamma_{i,j}(q_{i,j}) d^{q_{i,j}}_{pi}(i,j)}$, and similarly $Q(\lambda_a, \mathbf{a})$ for $\lambda^{tr,h}$ and $\lambda^{tr,v}$.

5.4. Experimental Comparisons

We consider the publicly available dataset [21] and compare the results of the proposed framework with other recently published methods, particularly a hypergraph-based method [5] that reported the best performance in this dataset. Two standard metrics, F-measure and Rand Index, are commonly used for measuring the segmentation performance, and the comparisons are shown in Table 2 and 3. While the pixel labeling result by the proposed method, which is shown as "2D-HMM-SASS", is sometime not smooth on the boundary of the cells, we use simple morphological operations, including dilation, filling, and erosion, to refine the output, where the results is shown as "2D-HMM-SASS (refine)". Compared with other recently published methods, our proposed framework provides a 7% improvement in terms of segmentation accuracy in the mean F-measure and a 10% improvement in the average rand index. Also, the run-time for solving the pixel labeling problem by the proposed 2D-HMM-SASS is 140 times faster than using the graph cuts method with higher order potentials (P^n model). It is clear that the proposed framework significantly outperforms the competitor. For visual inspection, we provide the segmentation result of 2D-HMM-SASS (refine) of the second frame in the image sequence in Fig. 1, and note that most neuronal processes are well segmented.

6. CONCLUSION

In this paper we proposed a two-dimensional hidden Markov model-based framework with spatially adaptive states for the cell tracing problem in a biological image sequence, which is posed as a pixel labeling problem. Exploiting available prior information, the concept of spatially adaptive states reduces the state-space to only a small subset of the full label set, which reduces not only the computational complexity but also the risk of label confusion. We compare the proposed framework with other recently published methods on a publicly available and highly challenging dataset of electron microscopy images of the rabbit retina. The experimental validation shows that the proposed framework significantly outperforms the competitors.

7. REFERENCES

[1] P. F. Felzenszwalb and D. P. Huttenlocher, "Efficient belief propagation for early vision.," *Int. J. Comput. Vision*, vol. 70, 2006.

[2] F. Perronnin, J. L. Dugelay, and K. Rose, "Iterative decoding of two-dimensional hidden markov models," in *Proc. IEEE ICASSP*, 2003.

[3] Y. Boykov, Olga Veksler, and R. Zabih, "Fast approximate energy minimization via graph cuts," *IEEE Trans. Pattern Anal. Mach. Intell.*, vol. 23, 1999.

[4] V Jagadeesh, N Vu, and B.S. Manjunath, "Multiple structure tracing in 3d electron micrographs," in *Proc. MICCAI*, Sep 2011.

[5] V. Jagadeesh, M.-C. Shih, B. S. Manjunath, and K. Rose, "Scalable tracing of electron micrographs by fusing top down and bottom up cues using hypergraph diffusion.," in *Proc. MICCAI*. Oct 2012, vol. 7512, pp. 321–328, Springer.

[6] B. Andres, U. Kthe, M. Helmstaedter, W. Denk, and F.A. Hamprecht, "Segmentation of sbfsem volume data of neural tissue by hierarchical classification," in *Proc. DAGM Pattern Recognition*. 2008, vol. 5096, pp. 142–152, Springer.

[7] A. Lucchi, K. Smith, R. Achanta, G. Knott, and P. Fua, "Supervoxel-based segmentation of mitochondria in em image stacks with learned shape features.," *IEEE Trans. Med. Imaging*, vol. 31, no. 2, pp. 474–486, 2012.

[8] Y. Boykov and G. Funka-Lea, "Graph cuts and efficient n-d image segmentation," *Int. J. Comput. Vision*, vol. 70, no. 2, pp. 109–131, Nov. 2006.

[9] E. Jurrus, R. Whitaker, B. W. Jones, R. Marc, and T. Tasdizen, "An optimal-path approach for neural circuit reconstruction," in *Proc. IEEE ISBI*, June 2008, pp. 1609 –1612.

[10] T. Liu, C. Jonesa, M. Seyedhosseinia, and T. Tasdizen, "A modular hierarchical approach to 3d electron microscopy image segmentation," *J. Neuroscience Methods*, vol. 226, pp. 88–102, 2014.

[11] D. C. Ciresan, A. Giusti, L. M. Gambardella, and J. Schmidhuber, "Deep neural networks segment neuronal membranes in electron microscopy images.," in *NIPS*, 2012, pp. 2852–2860.

[12] V. Kaynig, T. Fuchs, and J. M. Buhmann, "Neuron geometry extraction by perceptual grouping in sstem images," in *Proc. IEEE CVPR*, June 2010, pp. 2902–2909.

[13] T. Parag, S. Plaza, and L. Scheffer, "Small sample learning of superpixel classifiers for em segmentation," in *Proc. MICCAI*. 2014, vol. 8673, pp. 389–397, Springer-Verlag.

[14] J. Funke, B. Andres, F. A. Hamprecht, A. Cardona, and M. Cook, "Efficient automatic 3d-reconstruction of branching neurons from em data," in *Proc. IEEE CVPR*, Jun. 2012, pp. 1004–1011.

[15] V. Jagadeesh, J. Anderson, B. Jones, R. Marc, S. Fisher, and B.S. Manjunath, "Robust segmentation based tracing using an adaptive wrapper for inducing priors," *IEEE Transactions on Image Processing*, vol. 22, no. 12, pp. 4952–4963, Dec 2013.

[16] M. E. Sargin, A. Altinok, K. Rose, and B. S. Manjunath, "Conditional iterative decoding of two dimensional hidden markov models," in *Proc. IEEE ICIP*, 2008.

[17] J. Li, A. Najmi, and R.M. Gray, "Image classification by a two dimensional hidden markov model," *IEEE Trans. Signal Processing*, vol. 48, pp. 517–533, Feb. 2000.

[18] F. Perronnin, J.-L. Dugelay, and K. Rose, "A probabilistic model of face mapping with local transformations and its application to person recognition.," *IEEE Trans. Pattern Anal. Mach. Intell.*, vol. 27, no. 7, pp. 1157–1171, 2005.

[19] M. E. Sargin, L. Bertelli, B. S. Manjunath, and K. Rose, "Probabilistic occlusion boundary detection on spatio-temporal lattices.," in *Proc. IEEE ICCV*, 2009, pp. 560–567.

[20] R. Shenoy, M.-C. Shih, and K. Rose, "Hidden markov model-based multi-modal image fusion with efficient training," in *Proc. IEEE ICIP*, Oct. 2014.

[21] J. R. Anderson, B. W. Jones, C. V. Watt, and et al. M. V. Shaw, "Exploring the retinal connectome.," *Molecular Vision*, vol. 17, pp. 355–379, 2011.

[22] C. R. Maurer Jr., R. Qi, and V. Raghavan, "A linear time algorithm for computing exact euclidean distance transforms of binary images in arbitrary dimensions.," *IEEE Trans. Pattern Anal. Mach. Intell.*, vol. 25, no. 2, pp. 265–270, 2003.

[23] R.T. Whitaker, "A level-set approach to 3d reconstruction from range data.," *Int. J. Comput. Vision*, vol. 29, no. 3, pp. 203–231, 1998.

[24] S. Lankton, "Sparse field methods - technical report," Jun 2009.

[25] P. Kohli, L. Ladicky, and P.H.S. Torr, "Robust higher order potentials for enforcing label consistency.," in *CVPR*, 2008.