SEGMENTATION OF CELLS IN ELECTRON MICROSCOPY IMAGES THROUGH MULTIMODAL LABEL TRANSFER

Renuka Shenoy, Min-Chi Shih, Kenneth Rose

Department of Electrical and Computer Engineering, University of California, Santa Barbara

ABSTRACT

Automated segmentation of electron microcope (EM) images is a challenging problem, but the presence of related images of a different modality can be a valuable resource. This paper describes a method to effectively utilize complementary information, if available, in EM segmentation. Images of both modalities are oversegmented into superpixels. A 2D hidden Markov model (HMM) is set up on the superpixel graph to determine the optimal superpixel mapping between images. This mapping is used to transfer labels and generate preliminary segmentations in the EM domain, whose boundaries are then refined, to eliminate imprecisions due to the superpixel grid, using a 1D HMM based contour refinement method. The performance of the proposed approach is demonstrated on a challenging dataset, and significant improvement is observed over related techniques.

Index Terms— Segmentation, Multimodal, Electron Microscopy

1. INTRODUCTION

Connectomics [1] is a sub-field of bioinformatics which aims at comprehensively mapping neural connections at various scales. Analyzing this map of connections is essential in understanding the the architecture and dynamics of the nervous system. To assist in the understanding and interpretation of connectomes, tissues are often imaged with two or different modalities, resulting in complementary types of information (eg., structural and functional). The task of building the underlying network of connections from this image data relies on accurate segmentation of electron microscopy images. In this paper, we consider the problem of segmentation of EM images by transferring learned segmentations from a complementary source of information.

We test our algorithm on the RC1 connectome [2], an open access volume obtained from a 0.25 mm diameter column of the inner plexiform layer (IPL) of rabbit retina. The entire volume consists of 370 layers of image data, each one corresponding to a physical slice. Each slice is imaged either using Computational Molecular Phenotyping (CMP), a form of multi-marker light microscopy, or automated transmission electron microscopy (ATEM). The "capstone" region in the initial layers of this dataset consists of 6 CMP layers (each probed using a unique marker) followed an ATEM layer. Segmentation of cells in the ATEM images in RC1 is particularly difficult, due to considerable variation in appearance of cells, cluttered background, and, in some cases, low contrast between cell interiors and boundaries. On the other hand, cells in the CMP image can be segmented and classified by performing multivariate analysis on 6 layers taken together [3]. This observation led us to explore methods that utilize the information, available from the CMP modality, to achieve reliable EM segmentation.

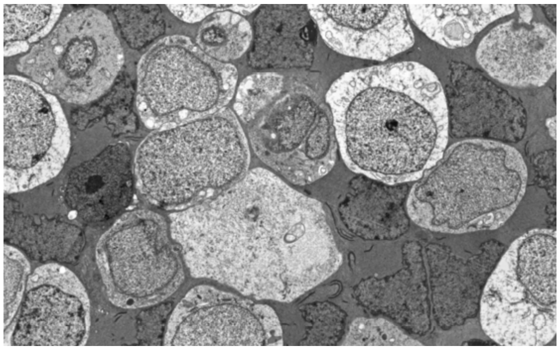

Fig. 1: An example of a region in an ATEM image. We observe considerable variation in appearance within and across cells, low contrast between cells and cell boundaries and background clutter.

1.1. Related work

In the recent past, there has been considerable interest in developing reliable segmentation algorithms for electron microscope (EM) images [4, 5, 6, 7, 8, 9]. In [4], membrane detection is performed using a random forest classifier followed by gap completion. A neural network approach to predict membranes is described in [6]. In [5] and [7], the authors use support vector machines (SVMs) to learn shape-based and context-based features respectively, which are used to segment mitochondria in EM images. Recently introduced methods have used hierarchical clustering with active learning [8] and merge trees [9] to perform EM segmentation. However, most existing EM segmentation algorithms are unable to reliably segment ATEM images from RC1 due to the variation in cell appearance through the image and lack of a clearly defined membrane in some cells. Further, due to background clutter and variation of cell characteristics through the image, traditional region-based segmentation approaches such as graph cuts [10, 11, 12] require careful initialization on each cell to perform well on such data.

Since the CMP data is available as an additional source of information, an alternative approach is to apply multimodal registration to deform CMP segments to the ATEM. There are many recent methods that aim at solving the multimodal alignment problem. In [13], the deformation is pursued using a Markov random field with a mutual information (MI) based matching criterion. The method in [14] maximizes multivariate α-MI between the two images to perform registration. In [15] and [16], the deformation is found by matching modality independent descriptors across images. Modality-specific parameters are learned and embedded in a probabilistic graphical

---

This work was supported by NSF OIA 0941717 and NSF III 0808772.

model in the multimodal fusion technique proposed in [17]. The drawback of using such methods is that they attempt to optimize a criterion that measures the quality of the mapping itself, which may be mismatched with out ultimate objective of optimal EM segmentation.

In this paper, we propose a multi-stage technique for segmentation of EM images. Information from automatically labeled and segmented CMP images is transferred to corresponding ATEM images by finding the optimal mapping between the two images. The initial contours obtained in the ATEM image as a result of this mapping are then further processed to get improved segmentation. The described system has the added advantage of transferring cell type information (which is difficult to determine using ATEM images alone) to each segment. While the discussions in this paper are in the context of the two modalities of the RC1 connectome, the described method is fairly general and may be easily adapted and applied in other multimodal settings.

2. BACKGROUND: HIDDEN MARKOV MODELS

Conventional (1D) hidden Markov models [18] have been widely and succesfully used in diverse domains. A major reason for their popularity is that optimal and computationally efficient techniques exist for parameter estimation and inference of state sequences in HMMs. The direct extension of these algorithms to 2D results in exponential growth in complexity and is usually intractable in practical applications. In our work, we use the "turbo-HMM" (T-HMM) [19], a reduced complexity variant of the 2D HMM which has been shown to perform well in optimization. In this framework, each row and column is a separately decoded 1D HMM. Rows and columns "communicate" by inducing priors on each other, and typically reach a high level of agreement within a few iterations.

Both 1D and 2D HMMs are used for optimization in the proposed approach, and are applied at different stages in the segmentation process.

3. PROPOSED METHOD

The proposed method aims at segmenting cells in ATEM images by leveraging information transferred from CMP to ATEM. We use a multi-stage technique to perform segmentation. The first step involves segmenting and labeling cells in CMP using a superpixel grid. Supersegmentation is then applied to the ATEM image, and a 2D HMM is used to find the optimal mapping between the two images. The cell segmentations are then transferred from CMP to ATEM, and an HMM-based contour refinement method is used to further improve the quality of segmentation.

3.1. CMP Segmentation

Segmentation and labeling of the CMP images, a prerequisite for the mapping step, is done using the probabilistic framework described in [3]. For N cell types, we consider an $M = (N + 1)$ class labeling problem, where the additional "non-cell" class provides for background points in the image which lie between cells. The input images are oversegmented via the SLIC [20] algorithm, using the implementation provided in [21] with a spatial regularization of 1. This results in a regular lattice of superpixels with high boundary recall. Labeling of each superpixel and cell segmentation are performed using a graphical model. This result is used in subsequent steps in out method.

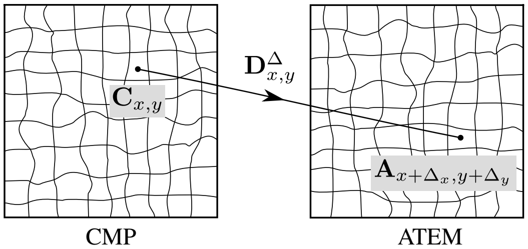

Fig. 2: Mapping a label from a SLIC superpixel in CMP to a SLIC superpixel in ATEM, using a mapping Δ

3.2. Multimodal Segmentation Transfer

Cell segmentations from the CMP image are transferred to the ATEM image using a 2D HMM framework built on a superpixel lattice. Using superpixels rather than a sliding window results in a reduced run time, while at the same time operating on a meaningful neighborhood of similar pixels. We oversegment the ATEM image using SLIC superpixels, in a similar setting to that used for CMP oversegmentation.

We construct a 2D HMM over a first order Markov mesh random field of size $X \times Y$, where X and Y are, respectively, the number of superpixels per row and column in the CMP image. Each superpixel in the CMP image corresponds to a node at a location (x, y) in the 2D HMM, and is denoted by $\mathbf{C}_{x,y}$. The label of $\mathbf{C}_{x,y}$ is denoted by $\omega_{x,y}$. Similarly, $\mathbf{A}_{x,y}$ denotes a superpixel at a location (x, y) in the ATEM superpixel lattice.

Each state q of the 2D HMM corresponds to a unique mapping Δ from a superpixel $\mathbf{C}_{x,y}$ in CMP to $\mathbf{A}_{x+\Delta_x,y+\Delta_y}$ in ATEM. Our aim is to find the optimal state sequence, $Q^* = \{q^*_{x,y}, x = 1, 2, \ldots X, y = 1, 2, \ldots Y\}$, which describes a deformation of the superpixel lattice. A pixel-level deformation is estimated by interpolating the superpixel-level result. This deformation is used to map labels from superpixels in CMP to superpixels in ATEM.

3.2.1. Data Matching Costs

The cost of matching a superpixel from the CMP image to a superpixel from the ATEM is captured by the emission probabilities of the 2D HMM. The emission probability $b^{\Delta}_{x,y}$ represents the probability of matching $\mathbf{C}_{x,y}$ to $\mathbf{A}_{x+\Delta_x,y+\Delta_y}$ (see Fig. 2).

$$b^{\Delta}_{x,y} = P(\mathbf{A}_{x+\Delta_x,y+\Delta_y} | \mathbf{C}_{x,y}) \tag{1}$$

We extract two features from every superpixel $\mathbf{A}_{x,y}$ in the ATEM image, the average and median intensity in the superpixel, and combine them into a single feature vector $\mathbf{f}_{x,y}$. Given the label $\omega_{x,y}$ from $\mathbf{C}_{x,y}$, emission probability can be rewritten as

$$b^{\Delta}_{x,y} = P(\mathbf{f}_{x+\Delta_x,y+\Delta_y} | \omega_{x,y}) \tag{2}$$

We model $P(\mathbf{f}_{x+\Delta_x,y+\Delta_y} | \omega_m)$ for the class ω_m with a mixture of K Gaussians.

$$P(\mathbf{f}_{x+\Delta_x,y+\Delta_y} | \omega_{x,y}) = \sum_{k=1}^{K} w^k_m P(\mathbf{f}_{x+\Delta_x,y+\Delta_y} | \omega^k_{x,y}) \tag{3}$$

where the set of mixture component weights for each class must satisfy $\sum_{k=1}^{K} w^k_m = 1 \ \forall m \in \{1, 2, \ldots M\}$. Each individual component density is a Gaussian of dimension 2.

$$P(\mathbf{f}_{x+\Delta_x,y+\Delta_y}|\omega_m^k) = \frac{\exp\{-\frac{1}{2}(\mathbf{f}_{x+\Delta_x,y+\Delta_y} - \boldsymbol{\mu}_m^k)^T {\boldsymbol{\Sigma}_m^k}^{-1}(\mathbf{f}_{x+\Delta_x,y+\Delta_y} - \boldsymbol{\mu}_m^k)\}}{(2\pi)|\boldsymbol{\Sigma}_m^k|^{\frac{1}{2}}} \tag{4}$$

where $\boldsymbol{\mu}_m^k$ and $\boldsymbol{\Sigma}_m^k$ are the mean and covariance, respectively.

3.2.2. Neighborhood Consistency

Consistency between translations of neighboring superpixels is ensured by the transition probabilities of the 2D HMM. In a first order 2D HMM, the state of a node depends on the state of its immediate neighbors. In our experiments, we define a spatially varying transition probability matrix which depends on the centroids of superpixels in both images.

A state corresponding to a mapping Δ maps the superpixel $\mathbf{C}_{x,y}$ in CMP to $\mathbf{A}_{x+\Delta_x,y+\Delta_y}$ in ATEM. Let $\mathbf{D}^{\Delta}_{x,y}$ be a 2D vector representing the distance between the centroids of superpixels $\mathbf{C}_{x,y}$ and $\mathbf{A}_{x+\Delta_x,y+\Delta_y}$. Each element of the horizontal transition probability matrix, $a^H_{\Delta',\Delta}(x,y)$, represents the probability of moving from state Δ at a location $(x-1,y)$ to the state Δ' at (x,y) in the 2D HMM, and is modeled by a Gaussian given by

$$a^H_{\Delta',\Delta}(x,y) \propto \exp\left\{-\frac{1}{2}\left(\frac{\|\mathbf{D}^{\Delta'}_{x,y} - \mathbf{D}^{\Delta}_{x-1,y}\|^2}{\sigma^2}\right)\right\} \tag{5}$$

where $\|\cdot\|$ denotes Euclidean distance. This model ensures smoothness in the resulting deformation by encouraging neighboring superpixels to take similar translations. The T-HMM framework allows for different translation probability matrices to be used for horizontal and vertical component HMMs. However, since cellular microscopy images do not typically exhibit different behavior in different directions, we assume a similar model for the vertical transition probability matrix.

3.2.3. Parameter Estimation and Inference

The parameters of the 2D HMM are learned using Baum-Welch training [22], which is an efficient implementation of maximum likelihood estimation. The optimal state sequence is inferred using the Viterbi decoding algorithm with modified forward-backward iterations, as described in [23].

3.3. Contour Refinement

We apply a basic version of the tracking algorithm from [24] to refine the obtained segmentation. Each segmented cell is taken individually, and an initial contour is set up on perimeter of the cell. Nodes for the HMM are initialized by sampling points at constant spacing along this contour, resulting in N_ϕ total nodes. A normal line is constructed at each node, and N_ψ equally spaced points (each corresponding to a state) are placed symmetrically along the normal (see Fig.3(a)), resulting in a deformable trellis. A given sequence of states, $Q = \{q_\phi, \phi = 1, 2, \ldots N_\phi\}$, corresponds to a path through the trellis, which forms a contour (see Fig.3(b)). Our aim is to estimate the optimal cell boundary by deforming the initial contour.

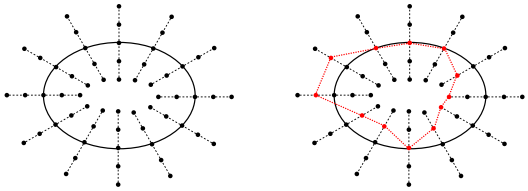
(a) (b)

Fig. 3: (Best in color) (a) Construction of a deformable trellis around the initial contour. The black curve is the initial contour, the dashed lines are the constructed normals, and each black dot corresponds to a state. (b) An example of an estimated contour (in red) obtained by joining a given sequence of states (red dots).

3.3.1. Emission Probability

The emission probabilities of the HMM represent the cost of the contour passing through a given point on the trellis. We extract three features to measure the emission probability b^ψ_ϕ for a state ψ at a location ϕ on the contour - one region-based feature which captures local object characteristics around the point of interest and two edge features which involve the first and second order average gradients along the normal line. The three features are concatenated into a single feature vector $\mathbf{f}^\psi_\phi$, and the emission probability is modeled with a 3-dimensional Gaussian with mean $\boldsymbol{\mu}_c$ and variance $\boldsymbol{\Sigma}_c$.

$$b^\psi_\phi = \frac{\exp\{-\frac{1}{2}(\mathbf{f}^\psi_\phi - \boldsymbol{\mu}_c)^T {\boldsymbol{\Sigma}_c}^{-1}(\mathbf{f}^\psi_\phi - \boldsymbol{\mu}_c)\}}{(2\pi)^{\frac{3}{2}}|\boldsymbol{\Sigma}_c|^{\frac{1}{2}}} \tag{6}$$

3.3.2. Transition Probability

The transition probabilities of the HMM ensure smoothness in the estimated contour, and are modeled with a tilted Gibbs distribution. Each element of the transition probability matrix, $a_{\psi',\psi}$, is the probability of moving to state ψ' from state ψ, and is given by

$$a_{\psi',\psi} \propto \exp\left\{-\frac{(\psi-\psi')^2}{2\theta_1^2}\right\} \cdot \exp\left\{-\frac{(\psi'-\psi_0)^2}{2\theta_2^2}\right\} \tag{7}$$

where ψ_0 is the middle state, which corresponds to the initial contour. The term containing θ_1 ensures consistency between consecutive points on the contour, whereas the term containing θ_2 penalizes large deviations from the initial contour.

3.3.3. Parameter Estimation and Inference

The emission parameters for the HMM used for contour refinement are learned using a support vector machine (SVM) [25] on labeled training data. Transition parameters are trained using maximum likelihood (Baum-Welch) on the training dataset. The optimal sequence is inferred using the Viterbi algorithm, and the resulting contour forms the refined segmentation boundary of the cell.

4. EXPERIMENTAL VALIDATION

We test the performance of several methods on images obtained from the RC1 connectome [2]. The capstone section consists of 6 channels of CMP data and a single channel of ATEM data and contains a total of 581 cells of 7 major types. The data is divided into training (90%) and testing (10%) portions. The test portion of the data consists of 85 total cells. To evaluate the accuracy of segmentation

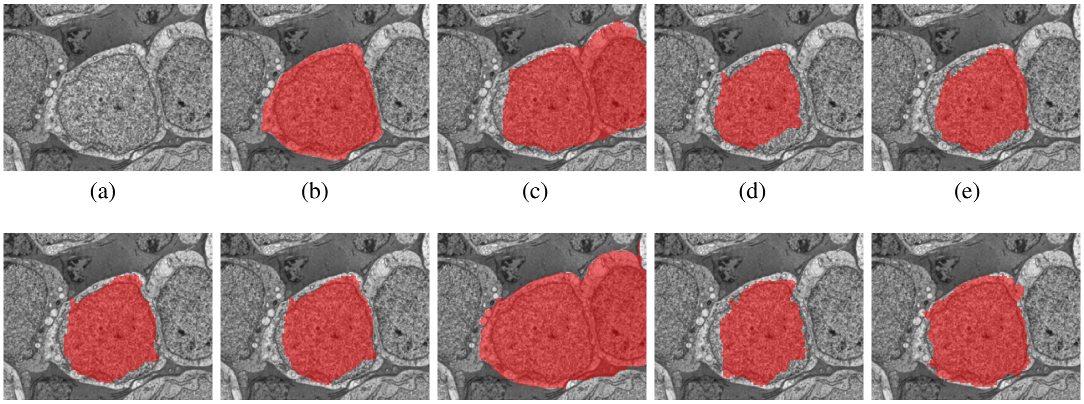
(a) (b) (c) (d) (e)

(f) (g) (h) (i) (j)

Fig. 4: (Best in color) Visual results of cell segmentation in a challenging scenario. (a) Region of interest in the ATEM image (b) Ground truth of the cell overlaid in red (c-i) Image with overlaid segmentation results from (c) the membrane detection method in [4] (d) DROP (e) α-MI (f) MIND (g) SSC (h) Graph cuts (i) the multimodal fusion method from [17] (j) the proposed approach

of each method, we compare the result with manually guided expert annotated ground truth. Segmented cells are optimally matched with ground truth cells using magnitude of overlap. The F-measure F is used as a measure of similarity between each ground truth cell (S_{GT}) and its corresponding segmentation (S_{SEG}).

$$F = \frac{2 \cdot PR \cdot RC}{PR + RC} \tag{8}$$

where the precision PR is given by $PR = \frac{|S_{GT} \cap S_{SEG}|}{|S_{SEG}|}$ and recall RC is given by $RC = \frac{|S_{GT} \cap S_{SEG}|}{|S_{GT}|}$. The operator $|\cdot|$ denotes number of pixels. The area (in pixels) of each ground truth cell is used as weight in the calculation of F-measure statistics.

We compare the performance of the proposed approach with related approaches that fall in two categories, direct segmentation methods and segmentation transfer via multimodal registration. The methods are listed below along with their respective optimal setting (each empirically found).

Direct ATEM Segmentation Methods: (i) Membrane detection using random forest classification followed by gap completion [4] with manually selected cells (since only contours are detected) (ii) Graph cut segmentation [10] in a multiple object setting, with parameters learned via manual seeding of a small subset of cells.

Label transfer-based Methods: (i) Graph cut segmentation [10], seeded with the results of CMP segmentation (ii) DROP [13], with $\lambda = 0.1$ (iii) α-MI-based registration [14], with $\alpha = 0.99$ and $k = 7$ (iv) MIND [15], with $\alpha = 0.1$ (v) SSC [16], with $\alpha = 0.1$ (vi) Our recent approach for multimodal fusion [17] (vii) the proposed method.

A comparison of quantitative results is provided in Table 1. We see that the proposed approach shows considerable improvement over competing methods in terms of segmentation accuracy. Visual results for a challenging scenario (touching cells without a well-defined membrane separating them) are shown in Fig. 4. We observe that direct EM segmentation methods are unable to separate the cell of interest from the visually similar adjoining cell, while registration-based methods utilize the label information from CMP to mitigate this problem. The additional contour refinement step of the proposed method results in further improved cell segmentation.

Method	Mean	Std Dev
Membrane detection [4]	0.7016	0.2732
Graph Cuts [10] (automatic)	0.6765	0.2889
DROP [13]	0.7412	0.1674
α-MI [14]	0.7288	0.1620
MIND [15]	0.7537	0.1688
SSC [16]	0.7523	0.1696
Graph Cuts [10] (seeded from CMP)	0.7808	0.2014
Multimodal fusion [17]	0.7978	0.1480
Proposed method	**0.8651**	**0.1287**

Table 1: F-measure statistics of segmentation results from various methods.

5. CONCLUSION

This paper presents a novel approach to segment objects in a multimodal setting, where information from a supplementary source is used to facilitate the segmentation of a challenging dataset. The labels from segmented cells in light microscopy images are transferred to electron micrographs using a 2D HMM-based mechanism built over a superpixel lattice. The obtained cell segmentations are refined using a HMM-based contour refinement technique. Experimental results show the capability of the proposed approach to effectively segment cells in ATEM images.

6. REFERENCES

[1] T.E.J. Behrens and O. Sporns, "Human connectomics," *Current opinion in neurobiology*, vol. 22, no. 1, pp. 144–153, 2012.

[2] J.R. Anderson, B.W. Jones, C.B. Watt, M.V. Shaw, J.-H. Yang, D. DeMill, J.S. Lauritzen, Y. Lin, K.D. Rapp, D. Mastronarde, et al., "Exploring the retinal connectome," *Molecular vision*, vol. 17, pp. 355, 2011.

[3] R. Shenoy, M.-C. Shih, and K. Rose, "A probabilistic framework for simultaneous segmentation and classification of multiple cells in multi-marker microscopy systems," in *IEEE International Symposium on Biomedical Imaging*, 2015.

[4] V. Kaynig, T. Fuchs, and J.M. Buhmann, "Neuron geometry extraction by perceptual grouping in sstem images," in *IEEE Conference on Computer Vision and Pattern Recognition*, 2010, pp. 2902–2909.

[5] 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 Transactions on Medical Imaging*, vol. 31, no. 2, pp. 474–486, 2012.

[6] D. Ciresan, A. Giusti, L.M. Gambardella, and J. Schmidhuber, "Deep neural networks segment neuronal membranes in electron microscopy images," in *Advances in neural information processing systems*, 2012, pp. 2843–2851.

[7] A. Lucchi, C.J. Becker, P. Marquez Neila, and P. Fua, "Exploiting enclosing membranes and contextual cues for mitochondria segmentation," in *Medical Image Computing and Computer Assisted Intervention*, 2014, number 199450.

[8] J. Nunez-Iglesias, R. Kennedy, T. Parag, J. Shi, and D.B. Chklovskii, "Machine learning of hierarchical clustering to segment 2d and 3d images," *PloS one*, vol. 8, no. 8, pp. e71715, 2013.

[9] T. Liu, C. Jones, M. Seyedhosseini, and T. Tasdizen, "A modular hierarchical approach to 3d electron microscopy image segmentation," *Journal of Neuroscience Methods*, vol. 226, pp. 88–102, 2014.

[10] Y. Boykov and G. Funka-Lea, "Graph cuts and efficient nd image segmentation," *International journal of computer vision*, vol. 70, no. 2, pp. 109–131, 2006.

[11] Y. Boykov and V. Kolmogorov, "An experimental comparison of min-cut/max-flow algorithms for energy minimization in vision," *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. 26, no. 9, pp. 1124–1137, 2004.

[12] V. Kolmogorov and R. Zabin, "What energy functions can be minimized via graph cuts?," *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. 26, no. 2, pp. 147–159, 2004.

[13] B. Glocker, A. Sotiras, N. Komodakis, and N. Paragios, "Deformable medical image registration: Setting the state of the art with discrete methods," *Annual review of biomedical engineering*, vol. 13, pp. 219–244, 2011.

[14] M. Staring, U.A. van der Heide, S. Klein, M.A. Viergever, and J. Pluim, "Registration of cervical mri using multifeature mutual information," *IEEE Transactions on Medical Imaging*, vol. 28, no. 9, pp. 1412–1421, 2009.

[15] M.P. Heinrich, M. Jenkinson, M. Bhushan, T. Matin, F.V. Gleeson, Sir M. Brady, and J.A. Schnabel, "Mind: Modality independent neighbourhood descriptor for multi-modal deformable registration," *Medical Image Analysis*, vol. 16, no. 7, pp. 1423–1435, 2012.

[16] M.P. Heinrich, M. Jenkinson, B.W. Papież, M. Brady, and J.A. Schnabel, "Towards realtime multimodal fusion for image-guided interventions using self-similarities," in *Medical Image Computing and Computer-Assisted Intervention–MICCAI 2013*, pp. 187–194. Springer, 2013.

[17] R. Shenoy, M.-C. Shih, and K. Rose, "Hidden markov model-based multi-modal image fusion with efficient training," in *IEEE International Conference on Image Processing*, 2014.

[18] L.R. Rabiner, "A tutorial on hidden markov models and selected applications in speech recognition," *Proceedings of the IEEE*, vol. 77, no. 2, pp. 257–286, 1989.

[19] F. Perronnin, J.-L. Dugelay, and K. Rose, "Iterative decoding of two-dimensional hidden markov models," in *IEEE International Conference on Acoustics, Speech, and Signal Processing*, 2003, vol. 3, pp. III–329.

[20] R. Achanta et al., "Slic superpixels compared to state-of-the-art superpixel methods," *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. 34, no. 11, pp. 2274–2282, 2012.

[21] A. Vedaldi and B. Fulkerson, "VLFeat: An open and portable library of computer vision algorithms," `http://www.vlfeat.org/`, 2008.

[22] L.E. Baum, T. Petrie, G. Soules, and N. Weiss, "A maximization technique occurring in the statistical analysis of probabilistic functions of markov chains," *The annals of mathematical statistics*, vol. 41, no. 1, pp. 164–171, 1970.

[23] F. Perronnin, J.-L. Dugelay, and K. Rose, "A probabilistic model of face mapping with local transformations and its application to person recognition," *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. 27, no. 7, pp. 1157–1171, 2005.

[24] M.-C. Shih and K. Rose, "Hidden markov models for tracking neuronal structure contours in electron micrograph stacks," in *IEEE International Symposium on Biomedical Imaging*, 2012, pp. 1377–1380.

[25] C.-C. Chang and C.-J. Lin, "LIBSVM: A library for support vector machines," *ACM Transactions on Intelligent Systems and Technology*, vol. 2, pp. 27:1–27:27, 2011, Software available at `http://www.csie.ntu.edu.tw/~cjlin/libsvm`.