

# Learned Watershed: End-to-End Learning of Seeded Segmentation

Steffen Wolf\*, Lukas Schott\*, Ullrich Köthe and Fred Hamprecht

HCI/IWR, Heidelberg University, 69120 Germany

firstname.lastname@iwr.uni-heidelberg.de

## Abstract

*Learned boundary maps are known to outperform hand-crafted ones as a basis for the watershed algorithm. We show, for the first time, how to train watershed computation jointly with boundary map prediction. The estimator for the merging priorities is cast as a neural network that is convolutional (over space) and recurrent (over iterations). The latter allows learning of complex shape priors. The method gives the best known seeded segmentation results on the CREMI segmentation challenge.*

## 1. Introduction

The watershed algorithm is an important computational primitive in low-level computer vision. Since it does not penalize segment boundary length, it exhibits no shrinkage bias like multi-terminal cuts or (conditional) random fields and is especially suited to segment objects with high surface-to-volume ratio, e.g. neurons in biological images.

In its classic form, the watershed algorithm comprises three basic steps: altitude computation, seed definition, and region assignment. These steps are designed manually for each application of interest. In a typical setup, the altitude is the output of an edge detector (e.g. the gradient magnitude or the gPb detector [2]), the seeds are located at the local minima of the altitude image, and pixels are assigned to seeds according to the drop-of-water principle [9].

In light of the very successful trend towards learning-based image analysis, it is desirable to eliminate hand-crafted heuristics from the watershed algorithm as well. Existing work shows that learned edge detectors significantly improve segmentation quality, especially when convolutional neural networks (CNNs) are used [7, 27, 33, 4]. We take this idea one step further and propose to learn altitude estimation and region assignment *jointly*, in an end-to-end fashion: Our approach no longer employs an auxiliary objective (e.g. accurate boundary strength prediction), but trains the altitude function together with the subsequent region assignment decisions so that the final segmentation

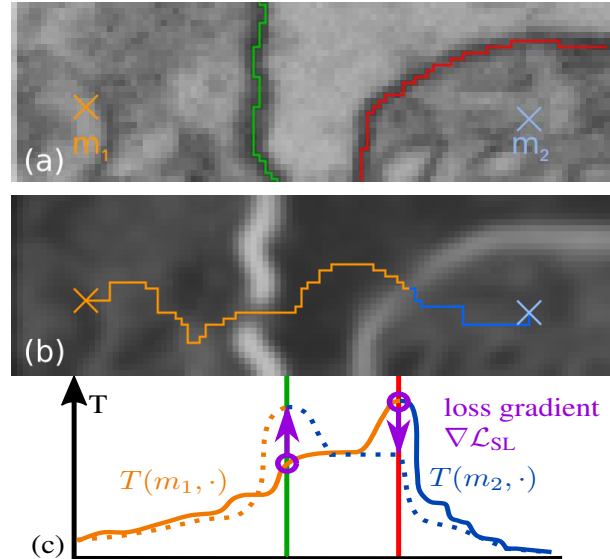


Figure 1. Illustration of watershed learning. (a) Raw data with two seeds  $m_1, m_2$ , ground truth boundary (green), and false boundary (red) at an early training stage. (b) A hole in the altitude map causes the left region to bleed out along the orange path, meeting the blue path at a false watershed. (c) Profile of topographic distances  $T$  along the two paths. Training increases the altitude at the missed edge (green) and decreases it at the false edge (red). Upon convergence, dotted paths meet at the correct location.

error is minimized directly. The resulting training algorithm is closely related to reinforcement learning.

Our method keeps the basic structure of the watershed algorithm intact: Starting from given seeds<sup>1</sup>, we maintain a priority queue storing the topographic distance of candidate pixels to their nearest seed. Each iteration assigns the currently best candidate to “its” region and updates the queue. The topographic distance is induced by an altitude function estimated with a CNN. Crucially, and deviating from prior work, we compute altitudes *on demand*, allowing their conditioning on prior decisions, i.e. partial segmentations. The CNN thus gets the opportunity to learn priors for likely region shapes in the present data. We show how these models

\* Authors contributed equally

<sup>1</sup>Incorporating seed definition into end-to-end learning is a future goal of our research, but beyond the scope of this paper.

can be trained end-to-end from given ground truth segmentations using *structured learning*. Our experiments show that the resulting segmentations are better than those from hand-crafted algorithms or unstructured learning.

## 2. Related Work

Various authors demonstrated that learned boundary probabilities (or, more generally, boundary strengths) are superior to designed ones. In the most common setting, these probabilities are defined on the pixel grid, i.e. on the nodes of a grid graph, and serve as input of a *node-based* watershed algorithm. Training minimizes a suitable loss (e.g. squared or cross-entropy loss) between the predicted probabilities and manually generated ground truth boundary maps in an *unstructured* manner, i.e. over all pixels independently. This approach works especially well with powerful models like CNNs. In the important application of connectomics (see section 6.3), this was first demonstrated by [14]. A much deeper network [7] was the winning entry of the ISBI 2012 Neuro-Segmentation Challenge [3]. Results could be improved further by progress in CNN architectures and more sophisticated data augmentation, using e.g. U-Nets [27], FusionNets [25] or networks based on inception modules [5]. Clustering of the resulting watershed superpixels by means of the GALA algorithm [24, 16] (using altitudes from [3] resp. [27]) or the lifted multicut [5] (using altitudes from their own CNN) lead to additional performance gains.

When ground truth is provided in terms of region labels rather than boundary maps, a suitable boundary map must be created first. Simple morphological operations were found sufficient in [27], while [5] preferred smooth probabilities derived from a distance transform starting at the true boundaries. Outside connectomics, [4] achieved superior results by defining the ground truth altitude map in terms of the *vector* distance transform, which allows optimizing the prediction's gradient direction and height separately.

Alternatively, one can employ the *edge-based* watershed algorithm and learn boundary probabilities for the grid graph's edges. The corresponding ground truth simply indicates if the end points of each edge are supposed to be in different segments or not. From a theoretical perspective, the distinction between node- and edge-based watersheds is not very significant because both can be transformed into each other [22]. However, the algorithmic details differ considerably. Edge-based altitude learning was first proposed in [12], who used hand-crafted features and logistic regression. Subsequently, [31] employed a CNN to learn features and boundary probabilities simultaneously. Watershed superpixel generation and clustering on the basis of these altitudes was investigated in [35].

Learning with unstructured loss functions has the disadvantage that an error at a single point (node or edge) has little effect on the loss, but may lead to large segmenta-

tion errors: A single missed boundary pixel can cause a big false merger. Learning with *structured* loss functions, as advocated in this paper, avoids this by considering the boundaries in each image jointly, so that the loss can be defined in terms of segmentation accuracy rather than point-wise differences. Holistically-nested edge detection [33, 17] achieves a weak form of this by coupling the loss at multiple resolutions using deep supervision. Such a network was successfully used as a basis for watershed segmentation in [6]. The MALIS algorithm [30] computes shortest paths between pairs of nodes and applies a correction to the highest edge along paths affected by false splits or mergers. This is similar to our training, but we apply corrections to root error edges as defined below. Learned, sparse reconstruction methods such as MaskExtend [20] and Flood-filling networks [15] predict region membership for all nodes in a patch jointly, performing region growing for one seed at a time in a one-against-the-rest fashion. In contrast, our algorithm grows all seeds simultaneously and competitively.

## 3. Mathematical Framework

The watershed algorithm is especially suitable when regions are primarily defined by their boundaries, not by appearance differences. This is often the case when the goal is *instance* segmentation (one neuron vs. its neighbors) as opposed to semantic segmentation (neurons vs. blood vessels). In graphical model terms, pairwise potentials between adjacent nodes are crucial in this situation, whereas unary potentials are of lesser importance or missing altogether. Many real-world applications have these characteristics, see [8] and section 6 for examples.

We consider 4-connected grid graphs  $G = (V, E)$ . The input image  $I : V \rightarrow \mathbb{R}^D$  maps all nodes to D-dimensional vectors of raw data. A segmentation is defined by a label image  $S : V \rightarrow \{1, 2, \dots, n\}$  specifying the region index or label of each node. The ground truth segmentation is called  $S^*$ . Pairwise potentials (i.e. edge weights) are defined by an altitude function over the graph's edges

$$f : E \rightarrow \mathbb{R} \quad (1)$$

where higher values indicate stronger boundary evidence. Since this paper focuses on how to learn  $f$ , we assume that a set of seed nodes  $M = \{m_1, \dots, m_n\} \subset V$  is provided by a suitable oracle (see section 6 for details). The watershed algorithm determines  $S$  by finding a mapping  $\sigma : V \rightarrow M$  that assigns each node to the best seed so that

$$\sigma(w) = m_i \quad \Rightarrow \quad S(w) = i \quad (2)$$

Initially, node assignments are unknown (designated by  $\lambda$ ) except at the seeds, where they are assumed to be correct:

$$\sigma_0(w) = \begin{cases} m_i & \text{if } w = m_i \text{ with } S^*(m_i) = i \\ \lambda & \text{otherwise} \end{cases} \quad (3)$$

In this paper, we build upon the *edge-based* variant of the watershed algorithm [21, 9]. This variant is also known as *watershed cuts* because segment boundaries are defined by cuts in the graph, i.e. by the set of edges whose incident nodes have different labels. We denote the cuts in our solution as  $\partial S$  and in the ground truth as  $\partial S^*$ .

Let  $\Phi(m, w)$  denote the set of all paths from seed  $m$  to node  $w$ . Then the *max-arc topographic distance* between  $m$  and  $w$  is defined as [11]

$$T(m, w) = \min_{\phi \in \Phi(m, w)} \max_{e \in \phi} f(e) \quad (4)$$

In words, the highest edge in a path  $\phi$  determines the path's altitude, and the path of lowest altitude determines the topographic distance. The watershed algorithm assigns each node to the topographically closest seed [26]:

$$\sigma(w) = \arg \min_{m \in M} T(m, w) \quad (5)$$

The minimum distance path from seed  $m$  to node  $w$  shall be denoted by  $\phi_m(w)$ . This path is not necessarily unique, but ties are rare and can be broken arbitrarily when  $f(e)$  is a real-valued function of noisy input data.

It was shown in [9] that the resulting partitioning is equivalent to the minimum spanning forest (MSF) over seeds  $M$  and edge weights  $f(e)$ . Thus, we can compute the watershed segmentation incrementally using Prim's algorithm: Starting from initial seeds  $\sigma_0$ , each iteration  $k$  finds the lowest edge whose start point  $u_k$  is already assigned, but whose end point  $v_k$  is not

$$u_k, v_k = \arg \min_{\substack{(u, v) \in E \\ \sigma_{k-1}(u) \neq \lambda, \sigma_{k-1}(v) = \lambda}} f(e = (u, v)) \quad (6)$$

and propagates the seed assignment from  $u_k$  to  $v_k$ :

$$\sigma_k(w) = \begin{cases} \sigma_{k-1}(u_k) & \text{if } w = v_k \\ \sigma_{k-1}(w) & \text{otherwise} \end{cases} \quad (7)$$

In a traditional watershed implementation, the altitude  $f(e)$  is a fixed, hand-designed function of the input data

$$f(e) = f_{\text{fixed}}(e|I) \quad (8)$$

for example, the image's Gaussian gradient magnitude or the "global Probability of boundary" (gPb) detector [2].

#### 4. Joint Structured Learning of Altitude and Region Assignment

We propose to use structured learning to train an altitude regressor  $f(e)$  jointly with the region assignment procedure defined by Prim's algorithm. We will discuss two types of learnable altitude functions:  $f_{\text{static}}$  comprises models that, once trained, only depend on the input image  $I$ , whereas  $f_{\text{dyn}}$  additionally incorporates dynamically changing information about the current state of Prim's algorithm.

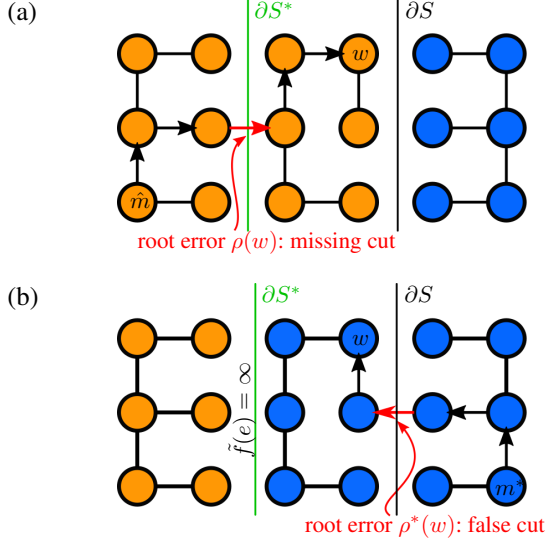


Figure 2. Example of root errors in the minimal spanning forest (a) and the constrained MSF (b) of a grid graph. Orange and blue indicate the segmentation  $S$  in (a) and  $S^*$  in (b). The root errors  $\rho(w)$  and  $\rho^*(w)$  of a wrongly labeled node  $w$  are marked red, with corresponding paths  $\phi_{\hat{m}}(w)$  and  $\psi_{m^*}(w)$  depicted by arrows.

##### 4.1. Static Altitude Prediction

To find optimal parameters  $\theta$  of a model  $f_{\text{static}}(e|I; \theta)$ , consider how Prim's algorithm proceeds: It builds a MSF which assigns each node  $w$  to the closest seed  $\hat{m} = \sigma(w)$  by identifying the shortest path  $\phi_{\hat{m}}(w)$  from  $\hat{m}$  to  $w$ . Such a path can be wrong in two ways: it may cross  $\partial S^*$  and thus miss a ground truth cut edge, or it may end at a false cut edge, placing  $\partial S$  in the interior of a ground truth region (see figure 1). More formally, we have to distinguish two failure modes: (i) A node was assigned to the wrong seed, i.e.  $\hat{m} \neq m^* = m_{S^*(w)}$  or (ii) it was assigned to the correct seed via a non-admissible path, i.e. a path taking a detour across a different region. To treat both cases uniformly, we construct the corresponding ground truth paths  $\psi_{m^*}(w)$ .

These paths can be found by running Prim's algorithm with a modified altitude

$$\tilde{f}(e) = \begin{cases} \infty & \text{if } e \in \partial S^* \\ f_{\text{static}}(e|I; \theta) & \text{otherwise} \end{cases} \quad (9)$$

forcing cuts in the resulting constrained MSF to coincide with  $\partial S^*$  (see figure 2). We denote the topographic distances along  $\phi_{\hat{m}}(w)$  and  $\psi_{m^*}(w)$  as  $T(\hat{m}, w)$  and  $T^*(m^*, w)$  respectively. By construction of the MSF,  $\phi$  and  $\psi$  are equal for all correct nodes. Conversely, they differ for incorrect nodes, causing distance  $T^*$  to exceed distance  $T$ . This property defines the set  $V_-$  of incorrect nodes:

$$V_- = \{w : T^*(m^*, w) > T(\hat{m}, w)\} \quad (10)$$

Every incorrect path  $\phi_{\hat{m}}(w)$  contains at least one erroneous

cut edge. The first such edge shall be called the path's *root error edge*  $\rho(w)$  and is always a missing cut. Training should increase its altitude until it becomes part of the cut set  $\partial S$ . The root error edge  $\rho^*(w)$  of a ground truth path  $\psi_{m^*}(w)$  is the first false cut edge in  $\psi$  in failure mode (i) and the first edge where  $\psi$  deviates from  $\phi$  in mode (ii). Here, the altitude should be decreased to make the edge part of the MSF, see figure 2. Accordingly, we denote the sets of root edges as  $E_{\uparrow} := \{\rho(w) : w \in V_{-}\}$  and  $E_{\downarrow} := \{\rho^*(w) : w \in V_{-}\}$ .

Since all assignment decisions in Prim's algorithm are conditioned on decisions taken earlier, the errors in any path also depend on the path's root error. Structured learning must therefore consider these errors jointly, and we argue that training updates must be derived solely from the root edges: They are the only locations whose required correction direction is unambiguously known. In contrast, we cannot even tell if subsequent errors will simply disappear once the root error has been fixed, or need updates of their own. When the latter applies, however, these edges will eventually become root errors in later training epochs, and we delay updating them to that point.

Since we need a differentiable loss to perform gradient descent, we use the perceptron loss of distance differences:

$$\mathcal{L} = \sum_w T^*(m^*, w) - T(\hat{m}, w) \quad (11)$$

Correct nodes have zero contribution since  $T^* = T$  holds for them. To serve as a basis for structured learning, we transform this into a loss over altitude differences at root edges. Since topographic distances equal the highest altitude along the shortest path, we have

$$T(\hat{m}, w) \geq f_{\text{static}}(\rho(w)) \quad (12)$$

To derive similar relations for  $T^*$ , consider how the constrained MSF is constructed from the unconstrained one: First, edges crossing  $\partial S^*$  are removed from the MSF. Each of the resulting orphaned subgraphs is then reconnected into the constrained MSF via the lowest edge not crossing  $\partial S^*$ . The newly inserted edges are the root edges  $\rho^*$  of all their child nodes, i.e. all nodes in the respective subgraph. Since these root edges did not belong to the original MSF, their altitude cannot be less than the maximum altitude in the corresponding child subgraph. For  $w \in V_{-}$ , it follows that

$$T^*(m^*, w) = f_{\text{static}}(\rho^*(w)) \quad (13)$$

We can therefore upper-bound the perceptron loss by

$$\mathcal{L}_{\text{SL}} = \sum_{w \in V_{-}} f_{\text{static}}(\rho^*(w)) - f_{\text{static}}(\rho(w)) \geq \mathcal{L} \quad (14)$$

and minimize this upper bound. By rearranging the sum, the loss can be simplified into

$$\mathcal{L}_{\text{SL}}(\theta) = \sum_{e \in E} R(e) f_{\text{static}}(e|I; \theta) \quad (15)$$

where we introduced a weight function counting the children of each root edge

$$R(e) := \begin{cases} \sum_{w: e=\rho^*(w)} 1 & \text{if } e \in E_{\downarrow} \\ -\sum_{w: e=\rho(w)} 1 & \text{if } e \in E_{\uparrow} \\ 0 & \text{otherwise} \end{cases} \quad (16)$$

A training epoch of structured learning thus consists of the following steps:

1. Compute  $f_{\text{static}}(e|I; \theta^{(t)})$  and  $\tilde{f}^{(t)}(e)$  with current model parameters  $\theta^{(t)}$  and determine the MSF and the constrained MSF.
2. Identify root edges and define the weights  $R^{(t)}(e)$  and the loss  $\mathcal{L}_{\text{SL}}^{(t)}(\theta)$ .
3. Obtain an updated parameter vector  $\theta^{(t+1)}$  via gradient descent on  $\nabla_{\theta} \mathcal{L}^{(t)}(\theta)$  at  $\theta = \theta^{(t)}$ .

These steps are iterated until convergence, and the resulting final parameter vector is denoted as  $\theta_{\text{SL}}$ .

## 4.2. Relation to Reinforcement Learning

In this section we compare the structured loss function  $\mathcal{L}_{\text{SL}}$  with policy gradient reinforcement learning, which will serve as motivation for a refinement of the weighting function  $R(e)$ . To see the analogy, we refer to continuous control deep reinforcement learning as proposed by [29, 28, 18].

Looking at the region growing procedure from a reinforcement learning perspective, we define states as tuples  $s = (e, I)$  where  $e \in E$  is the edge under consideration, and the action space  $\mathcal{A} := \mathbb{R}$  is the altitude to be predicted by  $f_{\text{static}}(e|I; \theta)$ . The *Policy Gradient Theorem* [29] defines the appropriate update direction of the parameter vector  $\theta$ . In a continuous action space, it reads

$$\nabla_{\theta} J = \nabla_{\theta} \sum_s d^{\pi}(s) \int_{\mathcal{A}} \pi(a|s; \theta) Q^{\pi}(s, a) da \quad (17)$$

where  $J$  is the performance to be optimized,  $d^{\pi}(s)$  the discounted state distribution,  $\pi$  the policy to be learned, and  $Q$  the action-value function estimating the discounted expected future reward

$$Q^{\pi}(s, a) = \mathbb{E}_{\pi} \left[ \sum_{t=0}^T \gamma^t r^t \middle| a_0 = a, s_0 = s; \pi \right] \quad (18)$$

In our case, the state distribution reduces to  $d^{\pi}(s) = \frac{1}{|V|}$  because Prim's algorithm reaches each edge exactly once. Inserting our deterministic altitude prediction

$$\pi(a|s) := f_{\text{static}}(s|I; \theta) \delta(a - f_{\text{static}}(s|I; \theta)), \quad (19)$$



where  $\delta$  is the Dirac distribution, we get

$$\nabla_{\theta} J = \frac{1}{|V|} \nabla_{\theta} \sum_s f_{\text{static}}(s|I; \theta) Q^{\pi}(s, a). \quad (20)$$

Comparing equation (20) with equation (15), we observe that  $\nabla_{\theta} J \sim \nabla_{\theta} \mathcal{L}_{\text{SL}}$ , where our weights  $R(e)$  essentially play the role of the action-value function  $Q$ . This suggests to introduce a *discount factor* in  $R(e)$ . To do so, we replace the temporal differences  $t$  between states in (18) with tree distances  $\text{dist}(w, \rho(w))$  or  $\text{dist}(w, \rho^*(w))$  counting the number of edges between node  $w$  and its root edge. This gives the discounted weights

$$R_{\text{RL}}(e) := \begin{cases} \sum_{w: e=\rho^*(w)} \gamma^{\text{dist}(w, \rho^*(w))} & \text{if } e \in E_{\downarrow} \\ \sum_{w: e=\rho(w)} -\gamma^{\text{dist}(w, \rho(w))} & \text{if } e \in E_{\uparrow} \\ 0 & \text{otherwise} \end{cases} \quad (21)$$

with discount factor  $0 \leq \gamma \leq 1$  to be chosen such that  $\gamma^{\text{dist}}$  decays roughly according to the size of the CNNs receptive field. Substituting  $R_{\text{RL}}(e)$  for  $R(e)$  in (15) significantly improves convergence in our experiments. This analogy further motivates the application of current deep reinforcement training methods as described in section 5.2.

### 4.3. Dynamic Altitude Prediction

In every iteration, region growing according to Prim’s algorithm only considers edges with exactly one end node located in the already assigned set. This offers the possibility to delay altitude estimation to the time when the values are actually needed. On-demand altitude computations can take advantage of the partial segmentations already available to derive additional *shape* clues that help resolving difficult assignment decisions.

**Relative Assignments:** To incorporate partial segmentations, we remove their dependence on the incidental choice of label values by means of *label-independent projection*. Consider an edge  $e = (u, v)$  where node  $u$  is assigned to seed  $m$  and node  $v$  is unassigned. We now construct a labeling relative to  $m$ , distinguishing nodes assigned to  $m$  (“me” region), to another seed (“them”) and unassigned (“nobody”). Relative labelings are represented by a standard 1-of-3 coding:

$$\mathcal{P}(w | m, \sigma) = \begin{cases} (1, 0, 0) & \text{if } \sigma(w) = m \text{ (me)} \\ (0, 1, 0) & \text{if } \sigma(w) = \lambda \text{ (nobody)} \\ (0, 0, 1) & \text{otherwise (them)} \end{cases} \quad (22)$$

In practice, we process relative labelings by adding a new branch to our neural network that receives  $\mathcal{P}$  as an input, see section 5.1 for details.

**Non-Markovian modeling:** Another potentially useful cue is afforded by the fact that Prim’s algorithm propagates

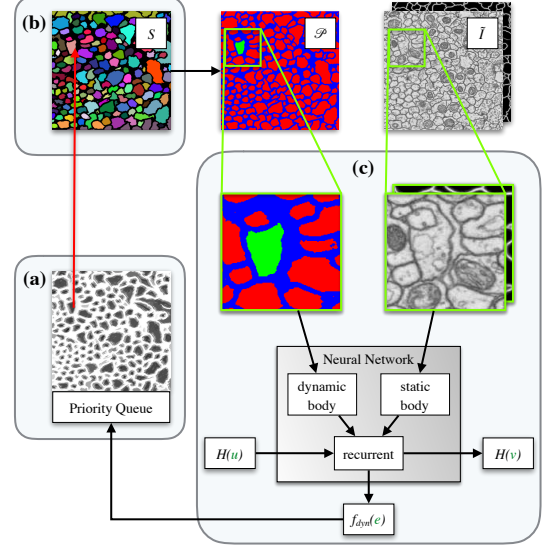


Figure 3. Overview implementation of learned watershed algorithm with neural network and priority queue. In each iteration the minimal edge according to equation (6) is found using a priority queue (a) and the region label is propagated (b), which updates the projection  $\mathcal{P}$ . For all unassigned edges that are not in the priority queue and need to be considered by Prim’s algorithm in the next iteration, the altitude  $f_{\text{dyn}}(e)$  is evaluated using the dynamic edge prediction network (c).

the assignments recursively. Thus during every evaluation of  $f$  the complete history from previous iterations along the growth paths  $\phi_m$  can be incorporated. We encode the history  $\mathcal{H} : V \rightarrow \mathbb{R}^r$  about past assignment decisions as an  $r$ -dimensional vector in each node. In practice, we incorporate history by adding a recurrent layer to our neural network.

We introduce the *dynamic* altitude predictions:

$$f(e = (u, v)), \mathcal{H}(v) = f_{\text{dyn}}(e | I, \mathcal{P}(\cdot | \sigma(u), \sigma), \mathcal{H}(u); \theta_{\text{dyn}}) \quad (23)$$

that receives the relative assignments projection  $\mathcal{P}$  and  $u$ ’s hidden state  $\mathcal{H}(u)$  as an additional input and outputs both the edge’s altitude  $f(e)$  and  $v$ ’s hidden state  $\mathcal{H}(v)$ : This variant of the altitude estimator performs best in our experiments. The emergent behavior of our models suggests that the algorithm uses history to adjust local diffusion parameters such as preferred direction and “viscosity”, similar to the viscous watershed transform described in [32].

## 5. Methods

### 5.1. Neural Network Architecture

Our network architecture builds mainly on the work of Yu and Koltun [34] who introduced dilated convolutions to achieve dense segmentations and systematically aggregate multi-scale contextual information without pooling

operations. We split our network into two convolutional branches (see Figure 4): The upper branch processes the static input  $I$ , and the lower one the dynamic input  $\mathcal{P}(\cdot)$ . Since the input of the upper branch doesn't change during prediction, its network activations can be precomputed for all edges, leading to a significant speed-up. We choose gated recurrent units (GRU) instead of long short-term memory (LSTM) in the recurrent network part, because GRUs have no internal state and get all history from the hidden state vector  $\mathcal{H}(\cdot)$ , saving on memory and bookkeeping.

## 5.2. Training Methods

**Augmenting the Input Image:** We noted above that structured learning is superior because it considers edges jointly. However, it can only rely on the sparse training sets  $E_{\downarrow} \cup E_{\uparrow}$ . In contrast, unstructured learning can make use of all edges and thus has a much bigger training set. This means that more powerful predictors, e.g. much deeper CNNs, can be trained, leading to more robust predictions and bigger receptive fields.

To combine the advantages of both approaches, we propose to augment the input image  $I$  with an additional channel holding node boundary probabilities predicted by an unstructured model  $g(w|I; \theta_{UL})$ :

$$\tilde{I} := [I \quad g] : V \rightarrow \mathbb{R}^{D+1} \quad (24)$$

We train the CNN  $g$  separately beforehand and replace  $I$  with the *augmented input*  $\tilde{I}$  everywhere in  $f_{\text{static}}$  and  $f_{\text{dyn}}$ . This simplifies structured learning because the predictor only needs to learn a refinement of the already reasonable altitudes in  $g$ . In principle, one could even train  $f$  and  $g$  jointly, but the combined model is too big for reliable optimization.

**Training Schedule:** Taking advantage of the close relationship with reinforcement learning, we adopt the asynchronous update procedure proposed by [23]. Here, independent workers fetch the current CNN parameters  $\theta$  from the master and compute loss gradients for randomly selected training images in parallel. The master then applies these updates to the parameters in an asynchronous fashion. We found experimentally, that this lead to faster and more stable convergence than sequential training.

In order to train the recurrent network part, we replace the standard temporal input ordering with the succession of edges defined by the paths  $\phi$  and  $\psi$ . In a sense, backpropagation in time thus becomes backpropagation along the minimum spanning forest.

## 6. Experiments and Results

Our experiments illustrate the performance of our proposed end-to-end trainable watershed in combination with static

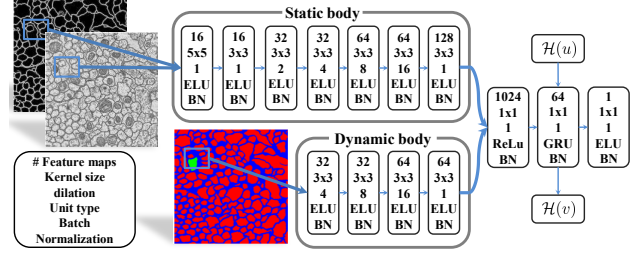


Figure 4. Network architecture: The static convolutional body extracts features from the raw input and edge detector output. The more shallow dynamic body processes the interactions of different region projections  $\mathcal{P}$ . These informations are combined in a fully connected layer and set into a temporal context using a recurrent GRU layer. The network output is the priority of the edge towards the pixel at the center of the field of view.

and dynamic altitude prediction. To this end, we compare with standard watershed and power watershed algorithms [8] on statically trained CNNs according to [5], see section 6.2. Furthermore we show in section 6.3 that the learned watershed surpasses the state-of-the-art segmentation in an adapted version of the CREMI Neuron Segmentation Challenge [10].

### 6.1. Experimental Setup and Evaluation Metrics

**Seed Generation Oracle:** All segmentation algorithms start at initial seeds  $M$  which are here provided by a “perfect” oracle. In our experiments, this oracle uses the ground truth segmentation to select one pixel with maximal  $L_2$  distance to the region boundary per ground truth region.

**Segmentation Metrics:** In accordance with the CREMI challenge [10], we use the following segmentation metrics: The Rand score  $V^{\text{Rand}}$  measures the probability of agreement between segmentation  $S$  and ground truth  $S^*$  w.r.t. a randomly chosen node pair  $w, w'$ . Two segmentations agree if both assign  $w$  and  $w'$  to the same region or to different regions. The Rand error  $\text{ARAND} = 1 - V^{\text{Rand}}$  is the opposite, so that smaller values are better.

The *Variation of Information (VOI)* between  $S$  and  $S^*$  is defined as  $\text{VOI}(S; S^*) = H(S|S^*) + H(S^*|S)$ , where  $H$  is the conditional entropy [19]. To distinguish split errors from merge errors, we report the summands separately as  $\text{VOI}^{\text{SPLIT}} = H(S|S^*)$  and  $\text{VOI}^{\text{MERGE}} = H(S^*|S)$ .

### 6.2. Artificial Data

**Dataset:** In order to compare our models  $f_{\text{static}}$  and  $f_{\text{dyn}}$  with solutions based on unstructured learning, we create an artificial segmentation benchmark dataset with variable difficulty. First, we generate an edge image via the zero crossing of a 2D Gaussian process. This image is then smoothed with a Gaussian filter and corrupted with Gaussian noise at  $\sigma_{\text{noise}} \in \{0.3, 0.6, 0.9\}$ . For each  $\sigma_{\text{noise}}$ , we generate 1900 training images and 100 test images of size 252x252. One

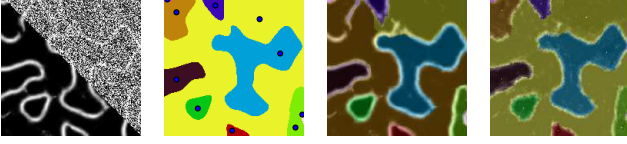


Figure 5. Artificial data example. a) Raw with  $\sigma_{\text{noise}} = 0.6$  and prediction of baseline CNN. b) Ground truth. c) The brown region leaks out when standard watershed runs on top of baseline CNN. d) Our algorithm uses learned shape priors to close boundary gaps.

test image with corresponding ground truth and results is shown in figure 5.

**Baseline:** We choose a recent edge detection network from [34] to predict boundaries between different instances in combination with standard watershed (WS) and Power Watershed (PWS) [8] to generate an instance segmentation. Since these algorithms work best on slightly smoothed inputs, we apply Gaussian smoothing to the CNN output. The optimal smoothing parameters are determined by grid search on the training dataset. Additionally, we apply all watershed methods directly to smoothed raw image and report their overall best result as *RAW + WS*.

**Performance:** The measured segmentation errors of all algorithms are shown in table 1. Observed differences in performance mainly indicate how well each method handles low-contrast edges and narrow gaps between regions. The structurally trained watersheds outperform the unstructured baselines, because our loss function  $\mathcal{L}_{\text{SL}}$  heavily penalizes the resulting segmentation errors. In all experiments, the *dynamic prediction* function  $f_{\text{dyn}}$  has the best performance, due to its superior modeling power. It can identify holes and close most contours correctly because it learns to derive shape and contingency clues from monitoring intermediate results during the flooding process. A representative example of this effect is shown in figure 5.

ARAND	$\sigma_{\text{noise}} = 0.3$	$\sigma_{\text{noise}} = 0.6$	$\sigma_{\text{noise}} = 0.9$
$\mathcal{L}_{\text{SL}} + f_{\text{dyn}}$	<b>5.8 <math>\pm</math> 0.8</b>	<b>12.5 <math>\pm</math> 1.7</b>	<b>32.2 <math>\pm</math> 1.8</b>
$\mathcal{L}_{\text{SL}} + f_{\text{static}}$	6.4 $\pm$ 0.9	13.8 $\pm$ 1.6	32.4 $\pm$ 2.2
NN + WS	6.5 $\pm$ 0.8	14.9 $\pm$ 3.6	33.4 $\pm$ 1.7
NN + PWS	6.5 $\pm$ 0.8	14.9 $\pm$ 1.7	33.2 $\pm$ 1.7
RAW + WS	24.0 $\pm$ 1.6	41.9 $\pm$ 1.8	55.0 $\pm$ 1.8

Table 1. Quality of the segmentation results on our artificial dataset. Reported lowest error for all parameters of baseline watersheds based on the rand error and a 2 pixel boundary distance tolerance.

### 6.3. Neurite Segmentation

**Dataset:** The MICCAI Challenge on Circuit Reconstruction from Electron Microscopy Images [10] contains 375 fully annotated slices of electron microscopy images  $I$  (of resolution 1250x1250 pixels). Part of a data slice is displayed in figure 7 top. Since the test ground truth segmentation has not been disclosed, we generate a new train/test

split from the 3 original challenge training datasets by splitting them into 3x75 z-continuous training- and 3x50 z-continuous test blocks.

Ideally, we would compare with [5] whose results define the state-of-the-art on the CREMI Challenge at time of submission. However, their pipeline, as described in their supplementary material, optimizes 2D segmentations jointly across multiple slices with a complex graphical model, which is beyond the scope of this paper.

Instead, we isolate the 2D segmentation aspect by adapting the challenge in the following manner: We run each segmentation algorithm with fixed ground truth seeds (see section 6.1) and evaluate their results on each  $z$ -slice separately. The restriction to 2D evaluation requires a slight manual correction of the ground truth: The ground truth accuracy in  $z$ -direction is just  $\pm 1$  slice. The official 3D evaluation scores compensate for this by ignoring deviations of  $\pm 1$  pixels in  $z$ -direction. Since this trick doesn't work in 2D, we remove 4 regions with no visual evidence in the image and all segments smaller than 5 pixels. Boundary tolerances in the x-y plane are treated as in the official CREMI scores where deviations from the true boundary are ignored if they do not exceed 6.25 pixels.

**Baseline:** We compare the Learned Watershed performance against the Power Watershed[8], Viscous Watershed [32], RandomWalker[13], Stochastic Watershed[1] and Distance Transform Watershed[5]. The boundary probability prediction  $g$  (the same  $g$  as in equation (24)) was provided by a deep CNN trained with an unstructured loss-function. In particular, the Distance Transform Watershed (DTWS) and the prediction  $g$  were used to produce the current state-of-the-art on the CREMI challenge. To obtain the DTWS, one thresholds  $g$ , computes a distance transform of the background, i.e. the non-boundary pixels and runs the watershed algorithm on the inverted distances. According to [5], this is the best known heuristic to close boundary gaps in these data, but requires manual parameter tuning. We found the parameters of all baseline algorithms by grid search using the training dataset. To ensure fair comparison, we start region growing from ground truth seeds in all cases. Our algorithm takes the augmented image  $\tilde{I}$  from equation (24) as input and learns how to close boundary gaps.

**Comparison to state-of-the-art:** We show the 2D CREMI segmentation scores in table 2. It is evident that the learned watershed transform significantly outperforms DTWS in both ARAND and VOI score. Quantitatively, we find that the flooding patterns and therefore the region shapes of the learned watershed prefer to adhere to biologically sensible structures. We illustrate this with our results on one CREMI test slice in figure 7, as well as specific examples in figure 6. We find throughout the dataset that especially *thin processes*, as depicted in fig. 6(a) left, are a strength of our algorithm. Biologically sensible *shape*



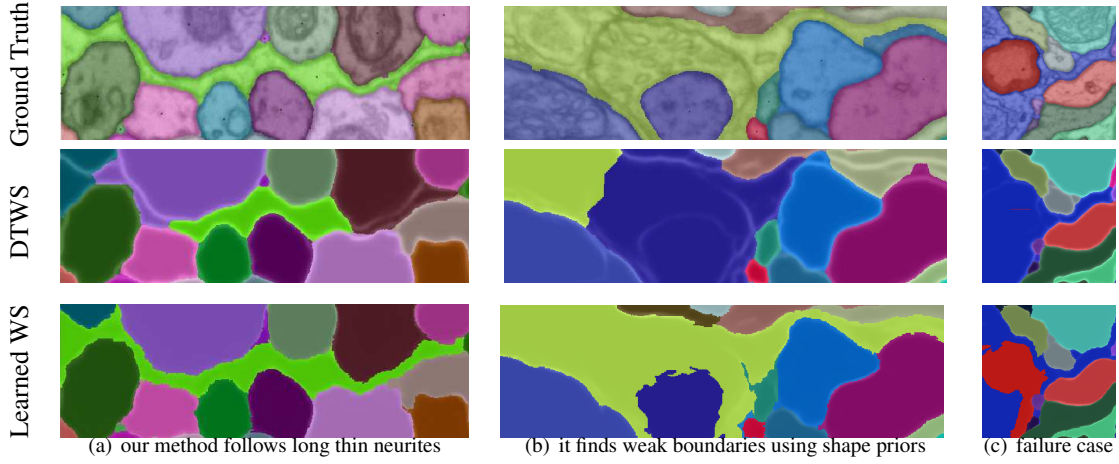


Figure 6. Detailed success and failure cases of our method.

*completions* can also be found for roundish objects and is particularly noticeable when boundary evidence is weak, as shown in Fig. 6(b) center. However, in rare cases, we find incorrect shape completions (see Fig. 6(c) right), mainly in areas of weak boundary evidence. It stands to reason that these errors could be fixed by providing more training data.

## 7. Conclusion

This paper proposes an end-to-end learnable seeded watershed algorithm that performs well on artificial data and neurosegmentation EM images. We found the following aspects to be critical success factors: First, we train a very powerful CNN to control region growing. Second, the CNN is trained in a structured fashion, allowing it to optimize segmentation performance directly, instead of treating pixels independently. Third, we improve modeling power by incorporating dynamic information about the current state of the segmentation. Specifically, feeding the current partial segmentation into the CNN provides shape clues for the next assignment decision, and maintaining a latent history along assignment paths allows to adjust growing parameters locally. We demonstrate experimentally that the resulting algorithm successfully solves difficult configurations like narrow region parts and low-contrast boundaries, where previous algorithms fail. In future work, we plan to include seed generation into the end-to-end learning scheme.

	ARAND	VOI split	VOI merge
PowerWS	$0.122 \pm 0.003$	$0.340 \pm 0.031$	$0.180 \pm 0.019$
ViscousWS	$0.093 \pm 0.003$	$0.328 \pm 0.030$	$0.069 \pm 0.003$
RandomWalker	$0.103 \pm 0.004$	$0.355 \pm 0.037$	$0.060 \pm 0.004$
Stochastic WS	$0.193 \pm 0.012$	$0.612 \pm 0.080$	$0.077 \pm 0.004$
DTWS	$0.085 \pm 0.001$	$0.320 \pm 0.029$	$0.070 \pm 0.005$
Learned WS	<b><math>0.082 \pm 0.001</math></b>	<b><math>0.319 \pm 0.030</math></b>	<b><math>0.057 \pm 0.004</math></b>

Table 2. CREMI segmentation metrics evaluated on 2D slices: The Variation of Information between a predicted segmentation and ground truth (lower is better) and the Adapted Rand Error (lower is better) [3].

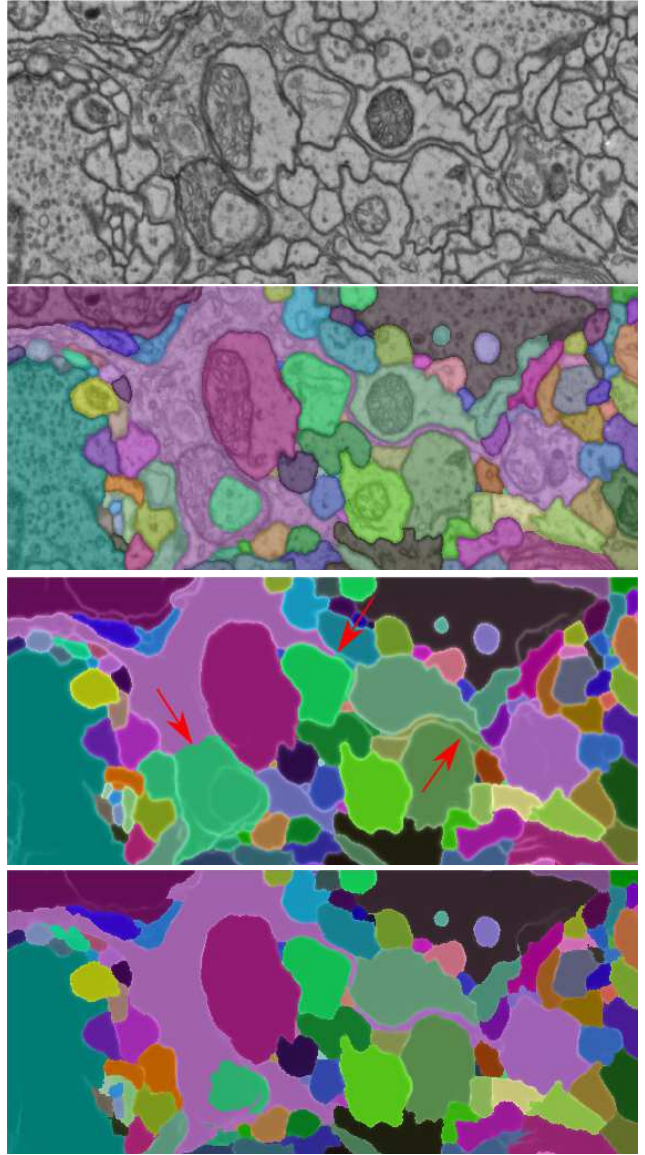


Figure 7. From top: Raw data. Ground truth. Result of distance transform WS (arrows point out major errors). Our algorithm.



## References

- [1] J. Angulo and D. Jeulin. Stochastic watershed segmentation. In *PROC. of the 8th International Symposium on Mathematical Morphology*, pages 265–276, 2007. [7](#)
- [2] P. Arbelaez, M. Maire, C. Fowlkes, and J. Malik. Contour detection and hierarchical image segmentation. *IEEE Trans. Patt. Anal. Mach. Intell.*, 33(5):898–916, 2011. [1](#), [3](#)
- [3] I. Arganda-Carreras, S. Turaga, D. Berger, et al. Crowdsourcing the creation of image segmentation algorithms for connectomics. *Front. Neuroanatomy*, 9:142, 2015. [2](#), [8](#)
- [4] M. Bai and R. Urtasun. Deep watershed transform for instance segmentation. *arXiv:1611.08303*, 2016. [1](#), [2](#)
- [5] T. Beier, C. Pape, N. Rahaman, T. Prange, et al. Multicut brings automated neurite segmentation closer to human performance. *Nature Methods*, 14(2):101–102, 2017. [2](#), [6](#), [7](#)
- [6] J. Cai, L. Lu, Z. Zhang, F. Xing, L. Yang, and Q. Yin. Pancreas segmentation in MRI using graph-based decision fusion on convolutional neural networks. In *Proc. MICCAI*, 2016. [2](#)
- [7] D. C. Ciresan, A. Giusti, L. M. Gambardella, and J. Schmidhuber. Deep neural networks segment neuronal membranes in electron microscopy images. *Proc. NIPS'12*, 2012. [1](#), [2](#)
- [8] C. Couprie, L. Grady, L. Najman, and H. Talbot. Power watershed: A unifying graph-based optimization framework. *IEEE Trans. Patt. Anal. Mach. Intell.*, 33(7), 2011. [2](#), [6](#), [7](#)
- [9] J. Cousty, G. Bertrand, L. Najman, and M. Couprie. Watershed cuts: Minimum spanning forests and the drop of water principle. *IEEE Trans. Patt. Anal. Mach. Intell.*, 2009. [1](#), [3](#)
- [10] C. CREMI. Miccai challenge on circuit reconstruction from electron microscopy images, 2017. [6](#), [7](#)
- [11] A. X. Falcão, J. Stolfi, and R. de Alencar Lotufo. The image foresting transform: Theory, algorithms, and applications. *IEEE Trans. Patt. Anal. Mach. Intell.*, 26(1):19–29, 2004. [3](#)
- [12] C. Fowlkes, D. Martin, and J. Malik. Learning affinity functions for image segmentation: combining patch-based and gradient-based approaches. In *Proc. CVPR*, 2003. [2](#)
- [13] L. Grady. Random walks for image segmentation. *IEEE Trans. Patt. Anal. Mach. Intell.*, 28(11):1768–1783, 2006. [7](#)
- [14] V. Jain, J. F. Murray, F. Roth, S. Turaga, V. Zhigulin, K. L. Briggman, M. N. Helmstaedter, W. Denk, and H. S. Seung. Supervised learning of image restoration with convolutional networks. *Proc. ICCV'07*, pages 1–8, 2007. [2](#)
- [15] M. Januszewski, J. Maitin-Shepard, P. Li, J. Kornfeld, W. Denk, and V. Jain. Flood-filling networks. *arXiv:1611.00421*, 2016. [2](#)
- [16] S. Knowles-Barley, V. Kaynig, T. R. Jones, A. Wilson, J. Morgan, D. Lee, D. Berger, N. Kasthuri, J. W. Lichtman, and H. Pfister. RhoanaNet pipeline: Dense automatic neural annotation. *arXiv:1611.06973*, 2016. [2](#)
- [17] I. Kokkinos. Pushing the boundaries of boundary detection using deep learning. *arXiv:1511.07386*, 2015. [2](#)
- [18] T. P. Lillicrap, J. J. Hunt, A. Pritzel, N. Heess, T. Erez, Y. Tassa, D. Silver, and D. Wierstra. Continuous control with deep reinforcement learning. *arXiv:1509.02971*, 2015. [4](#)
- [19] M. Meila. Comparing clusterings: an axiomatic view. In *Proc. ICML'05*, pages 577–584, 2005. [6](#)
- [20] Y. Meirovitch, A. Matveev, H. Saribekyan, D. Budden, D. Rolnick, G. Odor, S. K.-B. T. R. Jones, H. Pfister, J. W. Lichtman, and N. Shavit. A multi-pass approach to large-scale connectomics. *arXiv preprint:1612.02120*, 2016. [2](#)
- [21] F. Meyer. Minimum spanning forests for morphological segmentation. In *Mathematical morphology and its applications to image processing*, pages 77–84. 1994. [3](#)
- [22] F. Meyer. Watersheds on weighted graphs. *Pattern Recognition Letters*, 47:72 – 79, 2014. [2](#)
- [23] V. Mnih, A. P. Badia, M. Mirza, A. Graves, T. P. Lillicrap, et al. Asynchronous methods for deep reinforcement learning. In *Proc. ICML'16*, 2016. [6](#)
- [24] J. Nunez-Iglesias, R. Kennedy, T. Parag, J. Shi, and D. Chklovskii. Machine learning of hierarchical clustering to segment 2D and 3D images. *PLoS one*, 8:e71715, 2013. [2](#)
- [25] T. M. Quan, D. G. Hilderbrand, and W.-K. Jeong. FusionNet: a deep fully residual convolutional neural network for image segmentation in connectomics. *arXiv:1612.05360*, 2016. [2](#)
- [26] J. B. Roerdink and A. Meijster. The watershed transform: Definitions, algorithms and parallelization strategies. *Fundamenta informaticae*, 41(1, 2):187–228, 2000. [3](#)
- [27] O. Ronneberger, P. Fischer, and T. Brox. U-Net: convolutional networks for biomedical image segmentation. *Proc. MICCAI'15*, pages 234–241, 2015. [1](#), [2](#)
- [28] D. Silver, G. Lever, N. Heess, T. Degris, et al. Deterministic policy gradient algorithms. In *Proc. ICML'14*, 2014. [4](#)
- [29] R. S. Sutton, D. A. McAllester, S. P. Singh, Y. Mansour, et al. Policy gradient methods for reinforcement learning with function approximation. In *Proc. NIPS'99*, 1999. [4](#)
- [30] S. C. Turaga, K. L. Briggman, M. Helmstaedter, W. Denk, and H. S. Seung. Maximin affinity learning of image segmentation. *arXiv:0911.5372*, 2009. [2](#)
- [31] S. C. Turaga, J. F. Murray, V. Jain, F. Roth, M. Helmstaedter, K. Briggman, W. Denk, and H. S. Seung. Convolutional networks can learn to generate affinity graphs for image segmentation. *Neural Computation*, 22(2):511–538, 2010. [2](#)
- [32] C. Vachier and F. Meyer. The viscous watershed transform. *J. Math. Imaging and Vision*, 22(2):251–267, 2005. [5](#), [7](#)
- [33] S. Xie and Z. Tu. Holistically-nested edge detection. In *Proc. ICCV'15*, pages 1395–1403, 2015. [1](#), [2](#)
- [34] F. Yu and V. Koltun. Multi-scale context aggregation by dilated convolutions. *arXiv:1511.07122*, 2015. [5](#), [7](#)
- [35] A. Zlateski and H. S. Seung. Image segmentation by size-dependent single linkage clustering of a watershed basin graph. *arXiv:1505.00249*, 2015. [2](#)