

DeepOrgan: Multi-level Deep Convolutional Networks for Automated Pancreas Segmentation

Holger R. Roth, Le Lu, Amal Farag, Hoo-Chang Shin, Jiamin Liu,
Evrin B. Turkbey, and Ronald M. Summers

Imaging Biomarkers and Computer-Aided Diagnosis Laboratory, Radiology and
Imaging Sciences, National Institutes of Health Clinical Center, Bethesda, MD
20892-1182, USA

Abstract. Automatic organ segmentation is an important yet challenging problem for medical image analysis. The pancreas is an abdominal organ with very high anatomical variability. This inhibits previous segmentation methods from achieving high accuracies, especially compared to other organs such as the liver, heart or kidneys. In this paper, we present a probabilistic bottom-up approach for pancreas segmentation in abdominal computed tomography (CT) scans, using multi-level deep convolutional networks (ConvNets). We propose and evaluate several variations of deep ConvNets in the context of hierarchical, coarse-to-fine classification on image patches and regions, i.e. superpixels. We first present a dense labeling of local image patches via P-ConvNet and nearest neighbor fusion. Then we describe a regional ConvNet (R_1 -ConvNet) that samples a set of bounding boxes around each image superpixel at different scales of contexts in a “zoom-out” fashion. Our ConvNets learn to assign class probabilities for each superpixel region of being pancreas. Last, we study a stacked R_2 -ConvNet leveraging the joint space of CT intensities and the P -ConvNet dense probability maps. Both 3D Gaussian smoothing and 2D conditional random fields are exploited as structured predictions for post-processing. We evaluate on CT images of 82 patients in 4-fold cross-validation. We achieve a Dice Similarity Coefficient of $83.6 \pm 6.3\%$ in training and $71.8 \pm 10.7\%$ in testing.

1 Introduction

Segmentation of the pancreas can be a prerequisite for computer aided diagnosis (CADx) systems that provide quantitative organ volume analysis, e.g. for diabetic patients. Accurate segmentation could also necessary for computer aided detection (CADE) methods to detect pancreatic cancer. Automatic segmentation of numerous organs in computed tomography (CT) scans is well studied with good performance for organs such as liver, heart or kidneys, where Dice Similarity Coefficients (DSC) of $>90\%$ are typically achieved [1,2,3,4]. However, achieving high accuracies in automatic pancreas segmentation is still a challenging task. The pancreas’ shape, size and location in the abdomen can vary drastically between patients. Visceral fat around the pancreas can cause large variations in contrast along its boundaries in CT (see Fig. 3). Previous methods

report only 46.6% to 69.1% DSCs [1,2,3,5]. Recently, the availability of large annotated datasets and the accessibility of affordable parallel computing resources via GPUs have made it feasible to train deep convolutional networks (ConvNets) for image classification. Great advances in natural image classification have been achieved [6]. However, deep ConvNets for semantic image segmentation have not been well studied [7]. Studies that applied ConvNets to medical imaging applications also show good promise on detection tasks [8,9]. In this paper, we extend and exploit ConvNets for a challenging organ segmentation problem.

2 Methods

We present a coarse-to-fine classification scheme with progressive pruning for pancreas segmentation. Compared with previous top-down multi-atlas registration and label fusion methods, our models approach the problem in a bottom-up fashion: from dense labeling of image patches, to regions, and the entire organ. Given an input abdomen CT, an initial set of superpixel regions is generated by a coarse cascade process of random forests based pancreas segmentation as proposed by [5]. These pre-segmented superpixels serve as regional candidates with high sensitivity (>97%) but low precision. The resulting initial DSC is $\sim 27\%$ on average. Next, we propose and evaluate several variations of ConvNets for segmentation refinement (or pruning). A dense local image patch labeling using an axial-coronal-sagittal viewed patch (P -ConvNet) is employed in a sliding window manner. This generates a per-location probability response map P . A regional ConvNet (R_1 -ConvNet) samples a set of bounding boxes covering each image superpixel at multiple spatial scales in a “zoom-out” fashion [7,10] and assigns probabilities of being pancreatic tissue. This means that we not only look at the close-up view of superpixels, but gradually add more contexts to each candidate region. R_1 -ConvNet operates directly on the CT intensity. Finally, a stacked regional R_2 -ConvNet is learned to leverage the joint convolutional features of CT intensities and probability maps P . Both 3D Gaussian smoothing and 2D conditional random fields for structured prediction are exploited as post-processing. Our methods are evaluated on CT scans of 82 patients in 4-fold cross-validation (rather than “leave-one-out” evaluation [1,2,3]). We propose several new ConvNet models and advance the current state-of-the-art performance to a DSC of 71.8 in testing. To the best of our knowledge, this is the highest DSC reported in the literature to date.

2.1 Candidate Region Generation

We describe a coarse-to-fine pancreas segmentation method employing multi-level deep ConvNet models. Our hierarchical segmentation method decomposes any input CT into a set of local image superpixels $S = \{S_1, \dots, S_N\}$. After evaluation of several image region generation methods [11], we chose *entropy rate* [12] to extract N superpixels on axial slices. This process is based on the criterion of DSCs given optimal superpixel labels, in part inspired by the PASCAL

semantic segmentation challenge [13]. The optimal superpixel labels achieve a DSC upper-bound and are used for supervised learning below. Next, we use a two-level cascade of random forest (RF) classifiers as in [5]. We only operate the RF labeling at a low class-probability cut >0.5 which is sufficient to reject the vast amount of non-pancreas superpixels. This retains a set of superpixels $\{S_{\text{RF}}\}$ with high recall ($>97\%$) but low precision. After initial candidate generation, over-segmentation is expected and observed with low DSCs of $\sim 27\%$. The optimal superpixel labeling is limited by the ability of superpixels to capture the true pancreas boundaries at the per-pixel level with $DSC_{\text{max}} = 80.5\%$, but is still much above previous state-of-the-art [1,2,3,5]. These superpixel labels are used for assessing ‘positive’ and ‘negative’ superpixel examples for training. Assigning image regions drastically reduces the amount of ConvNet observations needed per CT volume compared to a purely patch-based approach and leads to more balanced training data sets. Our multi-level deep ConvNets will effectively prune the coarse pancreas over-segmentation to increase the final DSC measurements.

2.2 Convolutional Neural Network (ConvNet) Setup

We use ConvNets with an architecture for binary image classification. Five layers of *convolutional* filters compute and aggregate image features. Other layers of the ConvNets perform *max-pooling* operations or consist of fully-connected neural networks. Our ConvNet ends with a final two-way layer with *softmax* probability for ‘pancreas’ and ‘non-pancreas’ classification (see Fig. 1). The *fully-connected* layers are constrained using “DropOut” in order to avoid over-fitting by acting as a regularizer in training [14]. GPU acceleration allows efficient training (we use *cuda-convnet2*¹).

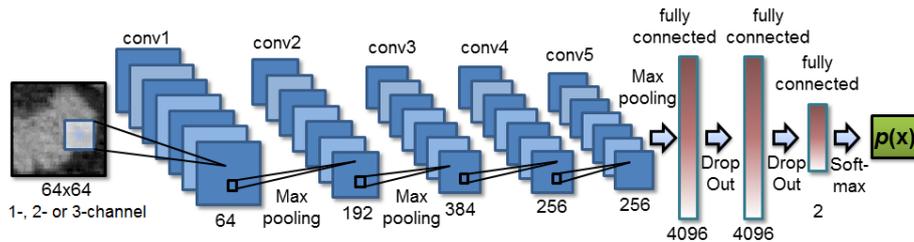


Fig. 1. The proposed ConvNet architecture. The number of convolutional filters and neural network connections for each layer are as shown. This architecture is constant for all ConvNet variations presented in this paper (apart from the number of input channels): P -ConvNet, R_1 -ConvNet, and R_2 -ConvNet.

¹ <https://code.google.com/p/cuda-convnet2>

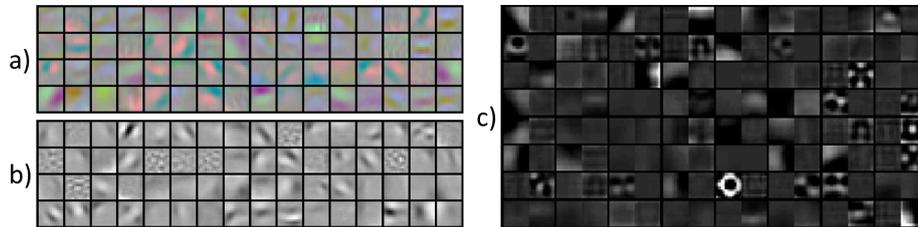


Fig. 2. The first layer of learned convolutional kernels using three representations: a) 2.5D sliding-window patches (P -ConvNet), b) CT intensity superpixel regions (R_1 -ConvNet), and c) CT intensity + P_0 map over superpixel regions (R_2 -ConvNet).

2.3 P -ConvNet: Deep Patch Classification

We use a sliding window approach that extracts 2.5D image patches composed of axial, coronal and sagittal planes within all voxels of the initial set of superpixel regions $\{S_{RF}\}$ (see Fig. 3). The resulting ConvNet probabilities are denoted as P_0 hereafter. For efficiency reasons, we extract patches every n voxels and then apply nearest neighbor interpolation. This seems sufficient due to the already high quality of P_0 and the use of overlapping patches to estimate the values at skipped voxels.

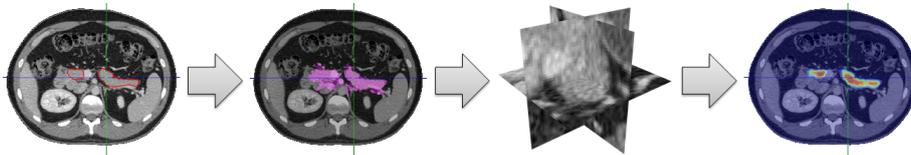


Fig. 3. Axial CT slice of a manual (gold standard) segmentation of the pancreas. From *left to right*, there are the ground-truth segmentation contours (in red), RF based coarse segmentation $\{S_{RF}\}$, and the deep patch labeling result using P -ConvNet.

2.4 R -ConvNet: Deep Region Classification

We employ the region candidates as inputs. Each superpixel $\in \{S_{RF}\}$ will be observed at several scales N_s with an increasing amount of surrounding contexts (see Fig. 4). Multi-scale contexts are important to disambiguate the complex anatomy in the abdomen. We explore two approaches: R_1 -ConvNet only looks at the CT intensity images extracted from multi-scale superpixel regions, and a stacked R_2 -ConvNet integrates an additional channel of patch-level response maps P_0 for each region as input. As a superpixel can have irregular shapes, we warp each region into a regular square (similar to RCNN [10]) as is required by most ConvNet implementations to date. The ConvNets automatically train their convolutional filter kernels from the available training data. Examples of trained

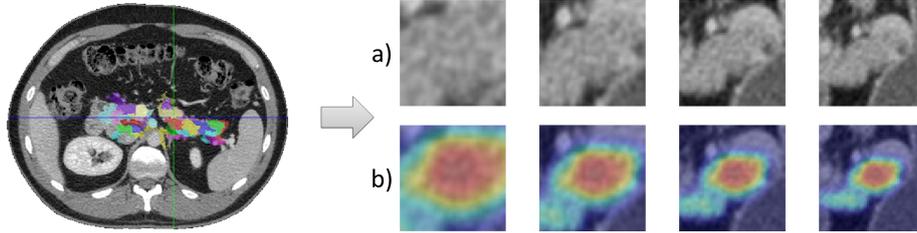


Fig. 4. Region classification using R -ConvNet at different scales: a) one-channel input based on the intensity image only, and b) two-channel input with additional patch-based P -ConvNet response.

first-layer convolutional filters for P -ConvNet, R_1 -ConvNet, R_2 -ConvNet are shown in Fig. 2. Deep ConvNets behave as effective image feature extractors that summarize multi-scale image regions for classification.

2.5 Data Augmentation

Our ConvNet models (R_1 -ConvNet, R_2 -ConvNet) sample the bounding boxes of each superpixel $\in \{S_{RF}\}$ at different scales s . During training, we randomly apply non-rigid deformations t to generate more data instances. The degree of deformation is chosen so that the resulting warped images resemble plausible physical variations of the medical images. This approach is commonly referred to as data augmentation and can help avoid over-fitting [6,8]. Each non-rigid training deformation t is computed by fitting a thin-plate-spline (TPS) to a regular grid of 2D control points $\{\omega_i; i = 1, 2, \dots, k\}$. These control points are randomly transformed within the sampling window and a deformed image is generated using a radial basic function $\phi(r)$, where $t(x) = \sum_{i=1}^k c_i \phi(\|x - \omega_i\|)$ is the transformed location of x and $\{c_i\}$ is a set of mapping coefficients.

2.6 Cross-Scale and 3D Probability Aggregation

At testing, we evaluate each superpixel at N_s different scales. The probability scores for each superpixel being pancreas are averaged across scales: $p(x) = \frac{1}{N_s} \sum_{i=1}^{N_s} p_i(x)$. Then the resulting per-superpixel ConvNet classification values $\{p_1(x)\}$ and $\{p_2(x)\}$ (according to R_1 -ConvNet and R_2 -ConvNet, respectively), are directly assigned to every pixel or voxel residing within any superpixel $\in \{S_{RF}\}$. This process forms two per-voxel probability maps $P_1(x)$ and $P_2(x)$. Subsequently, we perform 3D Gaussian filtering in order to average and smooth the ConvNet probability scores across CT slices and within-slice neighboring regions. 3D isotropic Gaussian filtering can be applied to any $P_k(x)$ with $k = 0, 1, 2$ to form smoothed $G(P_k(x))$. This is a simple way to propagate the 2D slice-based probabilities to 3D by taking local 3D neighborhoods into account. In this paper, we do not work on 3D supervoxels due to computational efficiency²

² Supervoxel based regional ConvNets need at least one-order-of-magnitude wider input layers and thus have significantly more parameters to train.

and generality issues. We also explore conditional random fields (CRF) using an additional ConvNet trained between pairs of neighboring superpixels in order to detect the pancreas edge (defined by pairs of superpixels having the same or different object labels). This acts as the *boundary term* together with the *regional term* given by R_2 -ConvNet in order to perform a min-cut/max-flow segmentation [15]. Here, the CRF is implemented as a 2D graph with connections between directly neighboring superpixels. The CRF weighting coefficient between the boundary and the unary regional term is calibrated by grid-search.

3 Results and Discussion

Data: Manual tracings of the pancreas for 82 contrast-enhanced abdominal CT volumes were provided by an experienced radiologist. Our experiments are conducted using 4-fold cross-validation in a random hard-split of 82 patients for training and testing folds with 21, 21, 20, and 20 patients for each testing fold. We report both training and testing segmentation accuracy results. Most previous work [1,2,3] uses leave-one-patient-out cross-validation protocols which are computationally expensive (e.g., ~ 15 hours to process one case using a powerful workstation [1]) and may not scale up efficiently towards larger patient populations. More patients (i.e. 20) per testing fold make the results more representative for larger population groups.

Evaluation: The ground truth superpixel labels are derived as described in Sec. 2.1. The optimally achievable DSC for superpixel classification (if classified perfectly) is 80.5%. Furthermore, the training data is artificially increased by a factor $N_s \times N_t$ using the data augmentation approach with both scale and random TPS deformations at the R -ConvNet level (Sec. 2.5). Here, we train on augmented data using $N_s = 4$, $N_t = 8$. In testing we use $N_s = 4$ (without deformation based data augmentation) and $\sigma = 3$ voxels (as 3D Gaussian filtering kernel width) to compute smoothed probability maps $G(P(x))$. By tuning our implementation of [5] at a low operating point, the initial superpixel candidate labeling achieves the average DSCs of only 26.1% in testing; but has a 97% sensitivity covering all pancreas voxels. Fig. 5 shows the plots of average DSCs using the proposed ConvNet approaches, as a function of $P_k(x)$ and $G(P_k(x))$ in both training and testing for one fold of cross-validation. Simple Gaussian 3D smoothing (Sec. 2.6) markedly improved the average DSCs in all cases. Maximum average DSCs can be observed at $p_0 = 0.2$, $p_1 = 0.5$, and $p_2 = 0.6$ in our training evaluation after 3D Gaussian smoothing for this fold. These calibrated operation points are then fixed and used in testing cross-validation to obtain the results in Table 1. Utilizing R_2 -ConvNet (stacked on P -ConvNet) and Gaussian smoothing ($G(P_2(x))$), we achieve a final average DSC of 71.8% in testing, an improvement of 45.7% compared to the candidate region generation stage at 26.1%. $G(P_0(x))$ also performs well with 69.5% mean DSC and is more efficient since only dense deep patch labeling is needed. Even though the absolute difference in DSC between $G(P_0(x))$ and $G(P_2(x))$ is small, the surface-to-surface

distance improves significantly from $1.46 \pm 1.5\text{mm}$ to $0.94 \pm 0.6\text{mm}$, ($p < 0.01$). An example of pancreas segmentation at this operation point is shown in Fig. 6. Training of a typical R -ConvNet with $N \times N_s \times N_t \approx 850k$ superpixel examples of size 64×64 pixels (after warping) takes ~ 55 hours for 100 epochs on a modern GPU (Nvidia GTX Titan-Z). However, execution run-time in testing is in the order of only 1 to 3 minutes per CT volume, depending on the number of scales N_s . Candidate region generation in Sec. 2.1 consumes another 5 minutes per case.

To the best of our knowledge, this work reports the highest average DSC with 71.8% in testing. Note that a direct comparison to previous methods is not possible due to lack of publicly available benchmark datasets. We will share our data and code implementation for future comparisons³⁴. Previous state-of-the-art results are at $\sim 68\%$ to $\sim 69\%$ [1,2,3,5]. In particular, DSC drops from 68% (150 patients) to 58% (50 patients) under the leave-one-out protocol [3]. Our results are based on a 4-fold cross-validation. The performance degrades gracefully from training ($83.6 \pm 6.3\%$) to testing ($71.8 \pm 10.7\%$) which demonstrates the good generality of learned deep ConvNets on unseen data. This difference is expected to diminish with more annotated datasets. Our methods also perform with better stability (i.e., comparing 10.7% versus 18.6% [1], 15.3% [2] in the standard deviation of DSCs). Our maximum test performance is 86.9% DSC with 10%, 30%, 50%, 70%, 80%, and 90% of cases being above 81.4%, 77.6%, 74.2%, 69.4%, 65.2% and 58.9%, respectively. Only 2 outlier cases lie below 40% DSC (mainly caused by over-segmentation into other organs). The remaining 80 testing cases are all above 50%. The minimal DSC value of these outliers is 25.0% for $G(P_2(x))$. However [1,2,3,5] all report gross segmentation failure cases with DSC even below 10%. Lastly, the variation $CRF(P_2(x))$ of enforcing $P_2(x)$ within a structured prediction CRF model achieves only $68.2\% \pm 4.1\%$. This is probably due to the already high quality of $G(P_0)$ and $G(P_2)$ in comparison.

Table 1. 4-fold cross-validation: optimally achievable DSCs, our initial candidate region labeling using S_{RF} , DSCs on $P(x)$ and using smoothed $G(P(x))$, and a CRF model for structured prediction (best performance in bold).

DSC (%)	Opt.	$S_{RF(x)}$	$P_0(x)$	$G(P_0(x))$	$P_1(x)$	$G(P_1(x))$	$P_2(x)$	$G(P_2(x))$	$CRF(P_2(x))$
Mean	80.5	26.1	60.9	69.5	56.8	62.9	64.9	71.8	68.2
Std	3.6	7.1	10.4	9.3	11.4	16.1	8.1	10.7	4.1
Min	70.9	14.2	22.9	35.3	1.3	0.0	33.1	25.0	59.6
Max	85.9	45.8	80.1	84.4	77.4	87.3	77.9	86.9	74.2

³ http://www.cc.nih.gov/about/SeniorStaff/ronald_summers.html

⁴ <http://www.holgerroth.com/>

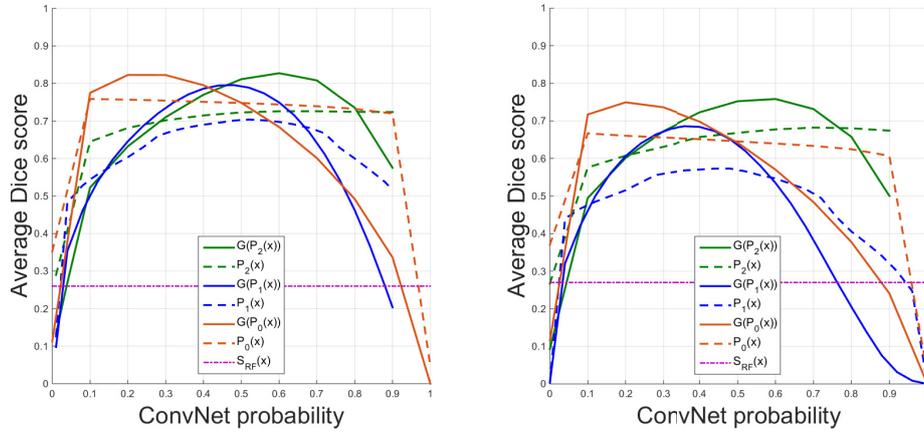


Fig. 5. Average DSCs as a function of un-smoothed $P_k(x)$, $k = 0, 1, 2$, and 3D smoothed $G(P_k(x))$, $k = 0, 1, 2$, ConvNet probability maps in training (left) and testing (right) in one cross-validation fold.

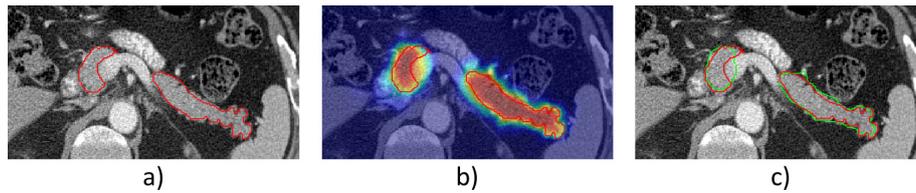


Fig. 6. Example of pancreas segmentation using the proposed R_2 -ConvNet approach in testing. a) The manual ground truth annotation (in red outline); b) the $G(P_2(x))$ probability map; c) the final segmentation (in green outline) at $p_2 = 0.6$ (DSC=82.7%).

4 Conclusion

We present a bottom-up, coarse-to-fine approach for pancreas segmentation in abdominal CT scans. Multi-level deep ConvNets are employed on both image patches and regions. We achieve the highest reported DSCs of $71.8 \pm 10.7\%$ in testing and $83.6 \pm 6.3\%$ in training, at the computational cost of a few minutes, not hours as in [1,2,3]. The proposed approach can be incorporated into multi-organ segmentation frameworks by specifying more tissue types since ConvNet naturally supports multi-class classifications [6]. Our deep learning based organ segmentation approach could be generalizable to other segmentation problems with large variations and pathologies, e.g. tumors.

Acknowledgments. This work was supported by the Intramural Research Program of the NIH Clinical Center.

References

1. Wang, Z., Bhatia, K.K., Glocker, B., Marvao, A., Dawes, T., Misawa, K., Mori, K., Rueckert, D.: Geodesic patch-based segmentation. In: Golland, P., Hata, N., Barillot, C., Hornegger, J., Howe, R. (eds.) MICCAI 2014, Part I. LNCS, vol. MICCAI, pp. 666–673. Springer, Heidelberg (2014)
2. Chu, C., et al.: Multi-organ segmentation based on spatially-divided probabilistic atlas from 3D abdominal CT images. In: Mori, K., Sakuma, I., Sato, Y., Barillot, C., Navab, N. (eds.) MICCAI 2013, Part II. LNCS, vol. 8150, pp. 165–172. Springer, Heidelberg (2013)
3. Wolz, R., Chu, C., Misawa, K., Fujiwara, M., Mori, K., Rueckert, D.: Automated abdominal multi-organ segmentation with subject-specific atlas generation. *TMI* 32(9), 1723–1730 (2013)
4. Ling, H., Zhou, S.K., Zheng, Y., Georgescu, B., Suehling, M., Comaniciu, D.: Hierarchical, learning-based automatic liver segmentation. In: *IEEE CVPR*, pp. 1–8 (2008)
5. Farag, A., Lu, L., Turkbey, E., Liu, J., Summers, R.M.: A bottom-up approach for automatic pancreas segmentation in abdominal CT scans. *MICCAI Abdominal Imaging workshop*, arXiv preprint arXiv:1407.8497 (2014)
6. Krizhevsky, A., Sutskever, I., Hinton, G.E.: Imagenet classification with deep convolutional neural networks. In: *NIPS*, pp. 1097–1105 (2012)
7. Mostajabi, M., Yadollahpour, P., Shakhnarovich, G.: Feedforward semantic segmentation with zoom-out features. arXiv preprint arXiv:1412.0774 (2014)
8. Cireşan, D.C., Giusti, A., Gambardella, L.M., Schmidhuber, J.: Mitosis detection in breast cancer histology images with deep neural networks. In: Mori, K., Sakuma, I., Sato, Y., Barillot, C., Navab, N. (eds.) MICCAI 2013, Part II. LNCS, vol. 8150, pp. 411–418. Springer, Heidelberg (2013)
9. Roth, H.R., et al.: A new 2.5D representation for lymph node detection using random sets of deep convolutional neural network observations. In: Golland, P., Hata, N., Barillot, C., Hornegger, J., Howe, R. (eds.) MICCAI 2014, Part I. LNCS, vol. MICCAI, pp. 520–527. Springer, Heidelberg (2014)
10. Girshick, R., Donahue, J., Darrell, T., Malik, J.: Rich feature hierarchies for accurate object detection and semantic segmentation. In: *IEEE CVPR*, pp. 580–587 (2014)
11. Achanta, R., Shaji, A., Smith, K., Lucchi, A., Fua, P., Susstrunk, S.: Slic superpixels compared to state-of-the-art superpixel methods. *PAMI* 34(11) (2012)
12. Liu, M.Y., Tuzel, O., Ramalingam, S., Chellappa, R.: Entropy rate superpixel segmentation. In: *IEEE CVPR*, pp. 2097–2104 (2011)
13. Everingham, M., Eslami, S.A., Gool, L.V., Williams, C.K., Winn, J., Zisserman, A.: The PASCAL visual object classes challenge: A retrospective. *IJCV* 111(1), 98–136 (2014)
14. Srivastava, N., Hinton, G., Krizhevsky, A., Sutskever, I., Salakhutdinov, R.: Dropout: A simple way to prevent neural networks from overfitting. *JMLR* 15(1), 1929–1958 (2014)
15. Boykov, Y., Funka-Lea, G.: Graph cuts and efficient ND image segmentation. *IJCV* 70(2), 109–131 (2006)