Deep Active Lesion Segmentation

Ali Hatamizadeh^{1*}, Assaf Hoogi^{2*}, Debleena Sengupta¹, Wuyue Lu¹,
Brian Wilcox², Daniel Rubin^{2†}, and Demetri Terzopoulos^{1†}

¹ Computer Science Department, University of California, Los Angeles, CA, USA

² Department of Biomedical Data Science, Stanford University, Stanford, CA, USA

Abstract. Lesion segmentation is an important problem in computer-assisted diagnosis that remains challenging due to the prevalence of low contrast, irregular boundaries that are unamenable to shape priors. We introduce Deep Active Lesion Segmentation (DALs), a fully automated segmentation framework that leverages the powerful nonlinear feature extraction abilities of fully Convolutional Neural Networks (CNNs) and the precise boundary delineation abilities of Active Contour Models (ACMs). Our DALs framework benefits from an improved level-set ACM formulation with a per-pixel-parameterized energy functional and a novel multiscale encoder-decoder CNN that learns an initialization probability map along with parameter maps for the ACM. We evaluate our lesion segmentation model on a new Multiorgan Lesion Segmentation (MLS) dataset that contains images of various organs, including brain, liver, and lung, across different imaging modalities—MR and CT. Our results demonstrate favorable performance compared to competing methods, especially for small training datasets.

Keywords: Lesion Segmentation · Active Contour Model · Level sets · Deep Learning

1 Introduction

Active Contour Models (ACMs) [6] have been extensively applied to computer vision tasks such as image segmentation, especially for medical image analysis. ACMs leverage parametric (“snake”) or implicit (level-set) formulations in which the contour evolves by minimizing an associated energy functional, typically using a gradient descent procedure. In the level-set formulation, this amounts to solving a partial differential equation (PDE) to evolve object boundaries that are able to handle large shape variations, topological changes, and intensity inhomogeneities. Alternative approaches to image segmentation that are based on deep learning have recently been gaining in popularity. Fully Convolutional Neural Networks (CNNs) can perform well in segmenting images within datasets on which they have been trained [9, 5, 2], but they may lack robustness when cross-validated

*Co-primary authorship

†Co-senior authorship

2 A. Hatamizadeh et al.

Fig. 1: Segmentation comparison of (a) medical expert manual with (b) our DALS and (c) U-Net [9], in (1) Brain MR, (2) Liver MR, (3) Liver CT, and (4) Lung CT images.

on other datasets. Moreover, in medical image segmentation, CNNs tend to be less precise in boundary delineation than ACMs.

In recent years, some researchers have sought to combine ACMs and deep learning approaches. Hu et al. [4] proposed a model in which the network learns a level-set function for salient objects; however, they predefined a fixed weighting parameter λ with no expectation of optimality over all cases in the analyzed set of images. Marcos et al. [8] combined CNNs and parametric ACMs for the segmentation of buildings in aerial images; however, their method requires manual contour initialization, fails to precisely delineate the boundary of complex shapes, and segments only single instances, all of which limit its applicability to lesion segmentation due to the irregular shapes of lesion boundaries and widespread cases of multiple lesions (e.g., liver lesions).

We introduce a fully automatic framework for medical image segmentation that combines the strengths of CNNs and level-set ACMs to overcome their respective weaknesses. We apply our proposed Deep Active Lesion Segmentation (DALS) framework to the challenging problem of lesion segmentation in MR

and CT medical images (Fig. 1), dealing with lesions of substantially different sizes within a single framework. In particular, our proposed encoder-decoder architecture learns to localize the lesion and generates an initial attention map along with associated parameter maps, thus instantiating a level-set ACM in which every location on the contour has local parameter values. We evaluate our lesion segmentation model on a new Multiorgan Lesion Segmentation (MLS) dataset that contains images of various organs, including brain, liver, and lung, across different imaging modalities—MR and CT. By automatically initializing and tuning the segmentation process of the level-set ACM, our DALs yields significantly more accurate boundaries in comparison to conventional CNNs and can reliably segment lesions of various sizes.

2 Method

2.1 Level-Set Active Contour Model With Parameter Functions

We introduce a generalization of the level-set ACM proposed by Chan and Vese [1]. Given an image $I(x, y)$, let $C(t) = \{(x, y) | \phi(x, y, t) = 0\}$ be a closed time-varying contour represented in $\Omega \in R^2$ by the zero level set of the signed distance map $\phi(x, y, t)$. We select regions within a square window of size s with a characteristic function W_s . The interior and exterior regions of C are specified by the smoothed Heaviside function $H_\epsilon^I(\phi)$ and $H_\epsilon^E(\phi) = 1 - H_\epsilon^I(\phi)$, and the narrow band near C is specified by the smoothed Dirac function $\delta_\epsilon(\phi)$. Assuming a uniform internal energy model [1], we follow Lankton *et al.* [7] and define m_1 and m_2 as the mean intensities of $I(x, y)$ inside and outside C and within W_s . Then, the energy functional associated with C can be written as

$$E(\phi) = \int_{\Omega} \delta_\epsilon(\phi(x, y, t)) \left(\mu |\nabla \phi(x, y, t)| + \int_{\Omega} W_s F(\phi(u, v, t)) du dv \right) dx dy, \quad (1)$$

where μ penalizes the length of C (we set $\mu = 0.1$) and the energy density is

$$F(\phi) = \lambda_1(u, v)(I(u, v) - m_1(x, y))^2 H_\epsilon^I(\phi) + \lambda_2(u, v)(I(u, v) - m_2(x, y))^2 H_\epsilon^E(\phi). \quad (2)$$

Note that to afford greater control over C , in (2) we have generalized the scalar parameter constants λ_1 and λ_2 used in [1] to *parameter functions* $\lambda_1(x, y)$ and $\lambda_2(x, y)$ over the image domain. Given an initial distance map $\phi(x, y, 0)$ and parameter maps $\lambda_1(x, y)$ and $\lambda_2(x, y)$, the contour is evolved by numerically time-integrating, within a narrow band around C for computational efficiency, the finite difference discretized Euler-Lagrange PDE for $\phi(x, y, t)$ (refer to [1] and [7] for the details).

2.2 CNN Backbone

Our encoder-decoder is a fully convolutional architecture (Fig. 2) that is tailored and trained to estimate a probability map from which the initial distance function $\phi(x, y, 0)$ of the level-set ACM and the functions $\lambda_1(x, y)$ and $\lambda_2(x, y)$ are

Fig. 2: The proposed DALS architecture. DALS is a fully automatic framework without the need for human supervision. The CNN initializes and guides the ACM by its learning local weighted parameters.

computed. In each dense block of the encoder, a composite function of batch normalization, convolution, and ReLU is applied to the concatenation of all the feature maps $[x_0, x_1, \dots, x_{l-1}]$ from layers 0 to $l-1$ with the feature maps produced by the current block. This concatenated result is passed through a transition layer before being fed to successive dense blocks. The last dense block in the encoder is fed into a custom multiscale dilation block with 4 parallel convolutional layers with dilation rates of 2, 4, 8, and 16. Before being passed to the decoder, the output of the dilated convolutions are then concatenated to create a multiscale representation of the input image thanks to the enlarged receptive field of its dilated convolutions. This, along with dense connectivity, assists in capturing local and global context for highly accurate lesion localization.

2.3 The DALS Framework

Our DALS framework is illustrated in Fig. 2. The boundaries of the segmentation map generated by the encoder-decoder are fine-tuned by the level-set ACM that takes advantage of information in the CNN maps to set the per-pixel parameters and initialize the contour.

The input image is fed into the encoder-decoder, which localizes the lesion and, after 1×1 convolutional and sigmoid layers, produces the initial segmentation probability map $Y_{prob}(x, y)$, which specifies the probability that any point (x, y) lies in the interior of the lesion. The Transformer converts Y_{prob} to a Signed Distance Map (SDM) $\phi(x, y, 0)$ that initializes the level-set ACM. Map Y_{prob} is also utilized to estimate the parameter functions $\lambda_1(x, y)$ and $\lambda_2(x, y)$ in the

Table 1: MLS dataset statistics. GC: Global Contrast; GH: Global Heterogeneity.

Organ	Modality	# Samples	Mean _{GC}	Var _{GC}	Mean _{GH}	Var _{GH}	Lesion Radius (pixels)
Brain	MRI	369	0.56	0.029	0.907	0.003	17.42 ± 9.516
Lung	CT	87	0.315	0.002	0.901	0.004	15.15 ± 5.777
Liver	CT	112	0.825	0.072	0.838	0.002	20.483 ± 10.37
Liver	MRI	164	0.448	0.041	0.891	0.003	5.459 ± 2.027

energy functional (1). Extending the approach of Hoogi et al. [3], the λ functions in Fig. 2 are chosen as follows:

$$\lambda_1(x, y) = \exp\left(\frac{2 - Y_{prob}(x, y)}{1 + Y_{prob}(x, y)}\right); \quad \lambda_2(x, y) = \exp\left(\frac{1 + Y_{prob}(x, y)}{2 - Y_{prob}(x, y)}\right). \quad (3)$$

The exponential amplifies the range of values that the functions can take. These computations are performed for each point on the zero level-set contour C . During training, Y_{prob} and the ground truth map $Y_{gt}(x, y)$ are fed into a Dice loss function and the error is back-propagated accordingly. During inference, a forward pass through the encoder-decoder and level-set ACM results in a final SDM, which is converted back into a probability map by a sigmoid layer, thus producing the final segmentation map $Y_{out}(x, y)$.

Implementation Details: DALIS is implemented in Tensorflow. We trained it on an NVIDIA Titan XP GPU and an Intel Core i7-7700K CPU @ 4.20GHz. All the input images were first normalized and resized to a predefined size of 256×256 pixels. The size of the mini-batches is set to 4, and the Adam optimization algorithm was used with an initial learning rate of 0.001 that decays by a factor of 10 every 10 epochs. The entire inference time for DALIS takes 1.5 seconds. All model performances were evaluated by using the Dice coefficient, Hausdorff distance, and BoundF.

3 Multiorgan Lesion Segmentation (MLS) Dataset

As shown in Table 1, the MLS dataset includes images of highly diverse lesions in terms of size and spatial characteristics such as contrast and homogeneity. The liver component of the dataset consists of 112 contrast-enhanced CT images of liver lesions (43 hemangiomas, 45 cysts, and 24 metastases) with a mean lesion radius of 20.483 ± 10.37 pixels and 164 liver lesions from 3T gadoxetic acid enhanced MRI scans (one or more LI-RADS (LR), LR-3, or LR-4 lesions) with a mean lesion radius of 5.459 ± 2.027 pixels. The brain component consists of 369 preoperative and pretherapy perfusion MR images with a mean lesion radius of 17.42 ± 9.516 pixels. The lung component consists of 87 CT images with a mean lesion radius of 15.15 ± 5.777 pixels. For each component of the MLS dataset, we used 85% of its images for training, 10% for testing, and 5% for validation.

Table 2: Segmentation metrics for model evaluations. Box and whisker plots are shown in Fig. 3. CI denotes the confidence interval.

Dataset:	Brain MR					Lung CT				
Model	Dice	CI	Hausdorff	CI	BoundF	Dice	CI	Hausdorff	CI	BoundF
U-Net	0.776 ± 0.214	0.090	2.988 ± 1.238	0.521	0.826	0.817 ± 0.098	0.0803	2.289 ± 0.650	0.53	0.898
CNN Backbone	0.824 ± 0.193	0.078	2.755 ± 1.216	0.49	0.891	0.822 ± 0.115	0.0944	2.254 ± 0.762	0.6218	0.900
Level-set	0.796 ± 0.095	0.038	2.927 ± 0.992	0.400	0.841	0.789 ± 0.078	0.064	3.27 ± 0.553	0.4514	0.879
DALS	0.888 ± 0.0755	0.03	2.322 ± 0.824	0.332	0.944	0.869 ± 0.113	0.092	2.095 ± 0.623	0.508	0.937

Dataset:	Liver MR					Liver CT				
Model	Dice	CI	Hausdorff	CI	BoundF	Dice	CI	Hausdorff	CI	BoundF
U-Net	0.769 ± 0.162	0.093	1.645 ± 0.598	0.343	0.92	0.698 ± 0.149	0.133	4.422 ± 0.969	0.866	0.662
CNN Backbone	0.805 ± 0.193	0.11	1.347 ± 0.671	0.385	0.939	0.801 ± 0.178	0.159	3.813 ± 1.791	1.6	0.697
Level-set	0.739 ± 0.102	0.056	2.227 ± 0.576	0.317	0.954	0.765 ± 0.039	0.034	3.153 ± 0.825	0.737	0.761
DALS	0.894 ± 0.0654	0.036	1.298 ± 0.434	0.239	0.987	0.846 ± 0.090	0.0806	3.113 ± 0.747	0.667	0.773

Fig. 3: Box and whisker plots of: (a) Dice score; (b) Hausdorff distance.

4 Results and Discussion

Algorithm Comparison: We have quantitatively compared our DALS against U-Net [9] and manually-initialized level-set ACM with scalar λ parameter constants as well as its backbone CNN. The evaluation metrics for each organ are reported in Table 2 and box and whisker plots are shown in Fig. 3. Our DALS achieves superior accuracies under all metrics and in all datasets. Furthermore, we evaluated the statistical significance of our method by applying a Wilcoxon paired test on the calculated Dice results. Our DALS performed significantly better than the U-Net ($p < 0.001$), the manually-initialized ACM ($p < 0.001$), and DALS's backbone CNN on its own ($p < 0.005$).

Boundary Delineation: As shown in Fig. 4, the DALS segmentation contours conform appropriately to the irregular shapes of the lesion boundaries, since the learned parameter maps, $\lambda_1(x, y)$ and $\lambda_2(x, y)$, provide the flexibility needed to accommodate the irregularities. In most cases, the DALS has also successfully

Fig. 4: Comparison of the output segmentation of our DALS (red) against the U-Net [9] (yellow) and manual “ground truth” (green) segmentations on images of Brain MR, Liver CT, Liver MR, and Lung CT on the MLS test set.

Fig. 5: (a) Labeled image. (b) Level-set (analogous to scalar λ parameter constants). (c) DALS output. (d), (e) Learned parameter maps $\lambda_1(x, y)$ and $\lambda_2(x, y)$.

avoided local minima and converged onto the true lesion boundaries, thus enhancing segmentation accuracy. DALS performs well for different image characteristics, including low contrast lesions, heterogeneous lesions, and noise.

Parameter functions and backbone CNN: The contribution of the parameter functions was validated by comparing the DALS against a manually initialized level-set ACM with scalar parameters constants as well as with DALS’s backbone

CNN on its own. As shown in Fig. 5, the encoder-decoder has predicted the $\lambda_1(x, y)$ and $\lambda_2(x, y)$ feature maps to guide the contour evolution. The learned maps serve as an attention mechanism that provides additional degrees of freedom for the contour to adjust itself precisely to regions of interest. The segmentation outputs of our DALs and the manual level-set ACM in Fig. 5 demonstrate the benefits of using parameter functions to accommodate significant boundary complexities. Moreover, our DALs outperformed the manually-initialized ACM and its backbone CNN in all metrics across all evaluations on every organ.

5 Conclusion

We have presented Deep Active Lesion Segmentation (DALs), a novel framework that combines the capabilities of the CNN and the level-set ACM to yield a robust, fully automatic medical image segmentation method that produces more accurate and detailed boundaries compared to competing state-of-the-art methods. The DALs framework includes an encoder-decoder that feeds a level-set ACM with per-pixel parameter functions. We evaluated our framework in the challenging task of lesion segmentation with a new dataset, MLS, which includes a variety of images of lesions of various sizes and textures in different organs acquired through multiple imaging modalities. Our results affirm the effectiveness of our DALs framework.

References

- [1] Chan, T.F., Vese, L.A.: Active contours without edges. *IEEE Transactions on Image Processing* 10(2), 266–277 (2001)
- [2] Hatamizadeh, A., Hosseini, H., Liu, Z., Schwartz, S.D., Terzopoulos, D.: Deep dilated convolutional nets for the automatic segmentation of retinal vessels. *arXiv preprint arXiv:1905.12120* (2019)
- [3] Hoogi, A., Subramaniam, A., Veerapaneni, R., Rubin, D.L.: Adaptive estimation of active contour parameters using convolutional neural networks and texture analysis. *IEEE Transactions on Medical Imaging* 36(3), 781–791 (2017)
- [4] Hu, P., Shuai, B., Liu, J., Wang, G.: Deep level sets for salient object detection. In: *Proc. IEEE Conf. on Computer Vision and Pattern Recognition* (2017)
- [5] Imran, A.A.Z., Hatamizadeh, A., Ananth, S.P., Ding, X., Terzopoulos, D., Tajbakhsh, N.: Automatic segmentation of pulmonary lobes using a progressive dense V-network. In: *Deep Learning in Medical Image Analysis, Lecture Notes in Computer Science*, vol. 11045, pp. 282–290. Springer (2018)
- [6] Kass, M., Witkin, A., Terzopoulos, D.: Snakes: Active contour models. *International Journal of Computer Vision* 1(4), 321–331 (1988)
- [7] Lankton, S., Tannenbaum, A.: Localizing region-based active contours. *IEEE Transactions on Image Processing* 17(11), 2029–2039 (2008)
- [8] Marcos, D., Tuia, D., Kellenberger, B., Zhang, L., Bai, M., Liao, R., Urtasun, R.: Learning deep structured active contours end-to-end. In: *Proc. IEEE Conf. on Computer Vision and Pattern Recognition (CVPR)*. pp. 8877–8885 (2018)
- [9] Ronneberger, O., Fischer, P., Brox, T.: U-net: Convolutional networks for biomedical image segmentation. In: *LNCS vol. 9351 (Proc. MICCAI)*. pp. 234–241 (2015)