Edge-Gated CNNs for Volumetric Semantic Segmentation of Medical Images

Ali Hatamizadeh^{1,2}

AHATAMIZ@CS.UCLA.EDU

Demetri Terzopoulos¹

DT@CS.UCLA.EDU

Andriy Myronenko²

AMYRONENKO@NVIDIA.COM

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

² *NVIDIA, Santa Clara, CA, USA*

Abstract

Textures and edges contribute different information to image recognition. Edges and boundaries encode shape information, while textures manifest the appearance of regions. Despite the success of Convolutional Neural Networks (CNNs) in computer vision and medical image analysis applications, predominantly only texture abstractions are learned, which often leads to imprecise boundary delineations. In medical imaging, expert manual segmentation often relies on organ boundaries; for example, to manually segment a liver, a medical practitioner usually identifies edges first and subsequently fills in the segmentation mask. Motivated by these observations, we propose a plug-and-play module, dubbed Edge-Gated CNNs (EG-CNNs), that can be used with existing encoder-decoder architectures to process both edge and texture information. The EG-CNN learns to emphasize the edges in the encoder, to predict crisp boundaries by an auxiliary edge supervision, and to fuse its output with the original CNN output. We evaluate the effectiveness of the EG-CNN with various mainstream CNNs on two publicly available datasets, BraTS 19 and KiTS 19 for brain tumor and kidney semantic segmentation. We demonstrate how the addition of EG-CNN consistently improves segmentation accuracy and generalization performance.

1. Introduction

Image segmentation plays an important role in medical image analysis as accurate delineation of anomalies is crucial for computer aided diagnosis and treatment planning. With the advent of deep learning, Convolutional Neural Networks (CNNs) have been successfully adopted in various medical semantic segmentation applications^[7;5;16;9;22;17]. In particular, the seminal U-Net architecture^[19] demonstrated the effectiveness of down-sampling and up-sampling paths for multi-scale feature representation learning, and many encoder-decoder CNNs have since been introduced based on the same principles. Medical images contain different types of representations: edges encode shape information, while the appearance of regions is manifested by textures. As such, a single processing pipeline may lead to the loss of shape information and may result in imprecise boundary definitions. It has been empirically demonstrated^[6] that unlike the human visual system, common CNN architectures are biased towards recognizing texture content. As a result, CNN predictions often need to be post-processed^[13;8;10] to compensate for the shape details that are lost during training.

The current paradigm of processing different abstractions within a single pipeline is sub-optimal. It can be remedied by utilizing effective processing of information in a structured

EG-CNN

manner, similar to the human visual perception system. In medical imaging, radiologists usually rely on identifying the boundaries of the organ/lesion of interest as a first step in manual delineation. For instance, segmenting brain tumors from MR images would entail following lesion edges and subsequently deducing the interior region.

Instead of proposing a new CNN architecture, we propose an novel 3D plug-and-play module that we call the Edge-Gated CNN (EG-CNN), which can be incorporated with any encoder-decoder architecture to disentangle the learning of texture and edge representations. The contribution of the proposed EG-CNN is two-fold. First, EG-CNN proposes an effective way to progressively learn to highlight the edge semantics from multiple scales of feature maps in the main encoder-decoder architecture by a novel and efficient layer denoted the edge-gated layer. Second, instead of separately supervising the edge and texture outputs, the EG-CNN introduces a dual-task learning scheme, in which these representations are jointly learned by a consistency loss. Therefore, without increasing the cost of data annotation and by exploiting the duality between edge and texture predictions, the EG-CNN improves the overall segmentation performance with highly detailed boundaries. Figure 1 illustrates the integration of the EG-CNN with an existing encoder-decoder CNN architecture.

We validate the effectiveness of our EG-CNN on two publicly available datasets, BraTS 2019^[1] and KiTS 2019^[11] for brain and kidney tumor segmentation, respectively. For this purpose, we utilize as backbones three popular 3D CNN architectures, U-Net^[19], V-Net^[15], and Seg-Net^[16], and our results demonstrate substantial improvement when EG-CNN is leveraged with these architectures.

2. Methodology

2.1. EG-CNN

We first introduce the architecture of EG-CNN. A generic CNN encoder-decoder, as we denote the main stream, learns feature representations that span multiple resolutions. Our EG-CNN receives each of the feature maps in the main stream and learns to highlight the edge representations. In particular, the EG-CNN consists of a sequence of residual blocks followed by tailored layers, as we denote the edge-gated layers, to progressively extract the edge representations. The output of the EG-CNN is then concatenated with the output of the main stream in order to produce the final segmentation output. Furthermore, the main stream and the EG-CNN are supervised by their own dedicated loss layers as well as a consistent loss function which jointly learns the output of both streams. The edge ground-truth is generated online by applying a 3D Sobel filter to the original ground truth masks.

Each edge-gated layer requires two inputs that originate from the main stream and the EG-CNN stream that we denote the edge stream. The intermediate feature maps from every resolution of the main stream as well as the first up-sampled feature maps in the decoder are fed to the EG-CNN as inputs. The latter is first fed into a residual block followed by bilinear upsampling before being fed into the edge-gated layer along with the input from its previous resolution in the encoder. The output of each edge-gated layer (except for the last one) is fed into another residual block followed by bilinear upsampling before being fed to the next edge-gated layer along with its corresponding input from the encoder.

EG-CNN

Figure 1: The proposed EG-CNN module can be integrated with any generic encoder-decoder architecture and highlight the edge representations of the intermediate feature maps.

2.2. Edge-Gated Layer

Edge-gated layers highlight the edge features and connect the feature maps learned in the main and edge streams. They receive inputs from the previous edge-gated layers as well as the main stream at its corresponding resolution. Let $e_{r,in}$ and m_r denote the inputs coming from edge and main streams, respectively, at resolution r . First, an attention map, α_r is obtained by feeding each input into a $1 \times 1 \times 1$ convolutional layer, $C_{1 \times 1 \times 1}$, fusing the outputs and passing them into a rectified linear unit (ReLU) $Re(X) = \max(0, X)$ according to

$$\alpha_r = \sigma(Re(C_{1 \times 1 \times 1}(e_{r,in}) + C_{1 \times 1 \times 1}(m_r))). \quad (1)$$

The obtained attention map α_r is then pixel-wise multiplied by $e_{r,in}$ and fed into a residual layer with kernel w_r . Therefore, the output of each resolution in EG-CNN, $e_{r,out}$, can be represented as

$$e_{r,out} = e_{r,in} \odot \alpha_r + e_{r,in}. \quad (2)$$

The computed attention map highlights the edge semantics that are embedded in the main stream feature maps. In general, there will be as many edge-gated layers as the number of different resolutions in the main encoder-decoder CNN architecture.

2.3. Loss Functions

The total loss of the EG-CNN is as follows:

$$L_{\text{Tot}} = L_{\text{Semantic}} + L_{\text{Consistency}} + L_{\text{Edge}}, \quad (3)$$

where L_{Semantic} represent standard loss functions used for supervising the main stream in a semantic segmentation network, L_{Edge} represent tailored losses for learning the edge representations, and $L_{\text{Consistency}}$ is a dual-task loss for the joint learning of edge and texture and enforces the class consistency of predictions.

Semantic Loss: Without loss of generality, we use the Dice loss^[15] for learning the semantic representations of texture according to

$$L_{\text{Dice}} = 1 - \frac{2 \sum y_{\text{true}} y_{\text{pred}}}{\sum y_{\text{true}}^2 + \sum y_{\text{pred}}^2 + \epsilon}, \quad (4)$$

where summation is carried over the total number of pixels, y_{pred} and y_{true} denote the pixel-wise semantic predictions of the main stream, and ϵ is a small constant to prevent division by zero.

Edge Loss: The edge loss used in EG-CNN comprises of Dice loss^[15] and balanced cross entropy^[21], as follows:

$$L_{\text{Edge}} = \lambda_1 L_{\text{Dice}} + \lambda_2 L_{\text{BCE}}, \quad (5)$$

where λ_1 and λ_2 are hyper-parameters. Let $e_{\text{pred},j}$ and $e_{\text{true},j}$ denote the edge prediction outputs of the EG-CNN and its corresponding groundtruth at voxel j , respectively. Then the balanced cross entropy L_{BCE} used in (5) can be defined as

$$L_{\text{BCE}} = -\beta \sum_{j \in e_+} \log P(e_{\text{pred},j} = 1|x; \theta) - (1 - \beta) \sum_{j \in e_-} \log P(e_{\text{pred},j} = 0|x; \theta), \quad (6)$$

where x , θ , e_- , and e_+ denote the input image, CNN parameters, edge, and non-edge voxel sets, respectively, β is the ratio of non-edge voxels to all voxels, and $P(e_{\text{pred},j})$ is the probability of the predicated class at voxel j . The cross entropy loss follows (6) except for the fact that non-edge voxels are not weighted.

Consistency Loss: We exploit the duality of edge and texture predictions and simultaneously supervise the outputs of the edge and main stream by the consistency loss. Inspired by^[20], the semantic probability predictions of the main CNN architectures and the ground truth masks are first converted into edge predictions by taking the spatial derivative in a differentiable manner as described in Section 2.1. Subsequently, we penalize the mismatch between the boundary predictions of the semantic masks and the corresponding ground truth by utilizing an L_1 loss. Let $y_{\text{pred},j}$ denote the output of the main stream and c represent the segmentation class. We propose a consistency loss function

$$L_{\text{Consistency}} = \sum_{j \in e_+} (\|\nabla(\arg \max(P(y_{\text{pred},j} = 1|e; c)))\| - \|\nabla(y_{\text{true},j})\|). \quad (7)$$

Due to the non-differentiability of the arg max function, we leverage the Gumbel softmax trick^[12] to avoid blocking the error-gradient. Thus, the gradient of the arg max can be approximated according to

$$\frac{\partial \arg \max_t P(y^t)}{\partial \gamma_j} = \nabla_j \frac{e^{(\log P(y_t) + g_t)/\tau}}{\sum_i e^{(\log P(y_i) + g_i)/\tau}}, \quad (8)$$

where γ is a differentiation dummy variable, τ is the temperature, set as a hyper-parameter, and g_i denotes the Gumbel density function.

3. Experiments

3.1. Datasets

BraTS 2019: The multimodal Brain Tumor Segmentation Challenge (BraTS 2019) serves to evaluate state-of-the-art methods for the segmentation of brain tumors by providing a 3D MRI dataset with ground truth tumor segmentation labels annotated by physicians^[4;14;3;2]. The BraTS 2019 training dataset includes 335 cases, each with four 3D MRI modalities (T1, T1c, T2 and FLAIR) rigidly aligned, resampled to $1 \times 1 \times 1$ mm isotropic resolution and skull-stripped. The input image size is $240 \times 240 \times 155$. Annotations include 3 tumor subregions: the enhancing tumor, the peritumoral edema, and the necrotic and non-enhancing tumor core. The annotations were combined into 3 nested subregions: Whole Tumor (WT), Tumor Core (TC), and Enhancing Tumor (ET).

KiTS 2019: The Kidney Tumor Segmentation Challenge (KiTS 2019) provides data comprising multi-phase 3D CTs and voxel-wise groundtruth labels of kidneys and kidney tumors for 300 patients who underwent nephrectomy for kidney tumors between 2010 to 2018 at the University of Minnesota^[11]. The input image size is $380 \times 380 \times 250$.

3.2. Implementation Details

We implemented all the models in Pytorch.¹ For our experiments with the BraTS 2019 dataset, we normalized all input images to have zero mean and unit std (based on non-zero voxels only). For the KiTS 2019 dataset, we normalized the CT data to the $[-1, 1]$ range by dividing the intensity values by 1,000 and clipping the values that fall outside this range. For training, images were resampled to $1 \times 1 \times 1$ mm isotropic resolution and resampled back to their original resolution after the inference. For both the BraTS 2019 and KiTS 2019 datasets, all the segmentation models were trained on and tested using a 80/20 split. We used $\lambda_1 = 1$, $\lambda_2 = 0.5$ as hyperparameters in (3) and (5). All models were trained on eight NVIDIA Tesla V100 16GB GPUs (DGX-1 server). We used a batch size of 8 and the Adam optimization algorithm was used with an initial learning rate of $\alpha_0 = 0.0001$ decreasing according to^[18]

$$\alpha = \alpha_0 * (1 - e/N_e)^{0.9} \quad (9)$$

with epoch counter e and total number of epochs N_e .

1. <https://pytorch.org/>

EG-CNN

Architecture	Edge Stream	Average Dice	ET Dice	TC Dice	WT Dice
U-Net	None	0.8305±0.0035	0.7375±0.0021	0.8480±0.0056	0.9060±0.0021
V-Net	None	0.8281±0.0035	0.7255±0.0049	0.8570±0.0042	0.9020±0.0014
Seg-Net	None	0.8300±0.0033	0.7330±0.0042	0.8550±0.0049	0.9015±0.0007
U-Net	EG-CNN	0.8406±0.0028	0.7530±0.0113	0.8630±0.0014	0.9006±0.0042
V-Net	EG-CNN	0.8386±0.0051	0.7460±0.0056	0.8605±0.0035	0.9095±0.0063
Seg-Net	EG-CNN	0.8570±0.0007	0.7680±0.0113	0.8850±0.0070	0.9180±0.0028

Table 1: BraTS 2019 segmentation results in terms of overall and tumor sub-regions Dice scores. The Edge Stream column determines whether EG-CNN is utilized with the backbone architecture.

3.3. Evaluation Metrics

For the BraTS 2019 challenge, we used the Dice function as a standard metric for image segmentation to assess the quality of the segmentation in the vicinity of boundaries. For the KiTS 2019 challenge, we adopted the three evaluation metrics of Kidneys, Tumor, and Composite Dice, as outlined by the organizers. Kidneys Dice denotes the segmentation performance when considering both kidneys and tumors as the foreground, whereas Tumor Dice considers everything except the tumor as background. Composite Dice is simply the average of Kidneys Dice and Tumor Dice.

4. Results and Discussion

We evaluate the EG-CNN module when it is used to augment popular medical image segmentation models: U-Net^[19], V-Net^[15], and Seg-Net^[16]. We modified each architecture to adopt them to the given task and to be similar to the others for a more fair comparison. For both the U-net and V-net, we changed the normalization to Groupnorm, to better handle a small batch size, and adjusted the number of layers to a roughly equivalent number between the networks. For each dataset we trained the main CNN segmentation network with and without the EG-CNN in order to validate the contribution of our proposed module. We estimated the accuracy of each model in terms of Dice score for each class and of the overall average.

BraTS 2019: Table 1 reports the accuracy of the model on each of the classes: Whole Tumor (WT), Tumor Core (TC), and Enhancing Tumor (ET), as well as the overall average accuracy. According to our benchmarks, including the EG-CNN consistently increases the overall and sub-region Dice scores in all cases. In the case of brain tumor segmentation, the EG-CNN has effectively learned highly complex and irregular boundaries of certain sub-regions. Therefore, it improves the segmentation quality around the edges, which leads to overall better segmentation performance. Figure 2 illustrates how the addition of the EG-CNN to a standalone Seg-Net^[16] improves the quality of segmentation.

The quality of the predicted edges also validates the effectiveness of our proposed edge-aware loss function, since the boundaries are crisp and avoid the thickening effect around edges. Such a phenomenon usually occurs when a naive loss function such as binary cross

EG-CNN

Figure 2: In-plane visualization of the segmentation outputs for tumor sub-regions in BraTS 2019. Red, green and yellow labels denote TC, WT and ET sub-regions respectively.

entropy is utilized for the task of edge prediction without taking precautions. Moreover, our model results in more fine-grained boundaries and visually attractive edges because the learned predicted boundaries are eventually fused with the final prediction output of the main encoder-decoder architecture.

Since the addition of the EG-CNN module increases the number of free parameters of the overall model, we have also experimented with larger standalone models (by increasing their depth and/or width), but doing so did not result in the better validation accuracy. This indicates that our module improves the overall segmentation accuracy not due to the model capacity increase, but due to the extra emphasis of edge information.

KiTS 2019: The achieved accuracy of the model for kidneys and kidney tumor classes, as well as the overall accuracy are presented in Table 2. Similar to the results achieved on BraTS 2019 dataset, the addition of EG-CNN has consistently improved the segmentation performance. Visual comparisons of the output segmentation and boundary predictions are presented in Figure 3. As such, the predicted edges visually conform to the region outlines, demonstrating that the EG-CNN module and our proposed loss functions helped to captured the details of the edges. This has also been reflected in the final predictions of semantic masks.

EG-CNN

Figure 3: Visualization of the segmentation performance of standalone Seg-Net and Seg-Net+EG-CNN on the KiTS 2019 challenge. The green and red masks denote tumor and kidneys, respectively.

Architecture	Edge Stream	Kidneys Dice	Tumor Dice	Composite Dice
U-Net	None	0.9515 $\pm$ 0.0049	0.8245 $\pm$ 0.0091	0.8880 $\pm$ 0.0070
V-Net	None	0.9370 $\pm$ 0.0065	0.8072 $\pm$ 0.0072	0.8720 $\pm$ 0.0068
Seg-Net	None	0.9530 $\pm$ 0.0028	0.8235 $\pm$ 0.0049	0.8892 $\pm$ 0.0038
U-Net	EG-CNN	0.9620 $\pm$ 0.0056	0.8270 $\pm$ 0.0084	0.8945 $\pm$ 0.0070
V-Net	EG-CNN	0.9483 $\pm$ 0.0048	0.8275 $\pm$ 0.0087	0.8879 $\pm$ 0.0067
Seg-Net	EG-CNN	0.9647$\pm$0.0051	0.8353$\pm$0.0025	0.9000$\pm$0.0038

Table 2: KiTS 2019 segmentation results for kidneys, tumor, and composite Dice functions.

5. Conclusion

We proposed the EG-CNN, a plug-and-play module for boundary-aware CNN segmentation, which can be paired with an existing encoder-decoder architecture to improve the segmentation accuracy. Our EG-CNN does not require any additional annotation effort since edge information can be extracted from the ground truth segmentation masks. Supervised by edge-aware and consistency loss functions, the EG-CNN learns to emphasize the edge representations by leveraging the feature maps of intermediate resolutions in the encoder of the main stream and feeding them into a series of edge-gated layers. We evaluated the EG-CNN by utilizing three popular 3D segmentation architectures, U-Net, V-Net, and Seg-Net, in

EG-CNN

the tasks of brain and kidney tumor segmentation on the BraTS19 and KiTS19 datasets. Our results indicate that the addition of the proposed EG-CNN consistently improves the segmentation accuracy in all the benchmarks.

References

- [1] Spyridon Bakas, Hamed Akbari, Aristeidis Sotiras, Michel Bilello, Martin Rozycki, Justin Kirby, John Freymann, Keyvan Farahani, and Christos Davatzikos. Segmentation labels and radiomic features for the pre-operative scans of the tcga-lgg collection. *The Cancer Imaging Archive*, 286, 2017.
- [2] Spyridon Bakas, Hamed Akbari, Aristeidis Sotiras, Michel Bilello, Martin Rozycki, Justin Kirby, Keyvan Farahani John Freymann, and Christos Davatzikos. Segmentation labels and radiomic features for the pre-operative scans of the tcga-lgg collection. *The Cancer Imaging Archive*, 2017. URL <https://doi.org/10.7937/K9/TCIA.2017.GJQ7ROEF>.
- [3] Spyridon Bakas, Hamed Akbari, Aristeidis Sotiras, Michel Bilello, Martin Rozycki, Justin S. Kirby, John B. Freymann, Keyvan Farahani, and Christos Davatzikos. Advancing the cancer genome atlas glioma mri collections with expert segmentation labels and radiomic features. *Scientific data*, 4, 9 2017. ISSN 2052-4463.
- [4] Spyridon Bakas, Mauricio Reyes, et Int, and Bjoern Menze. Identifying the best machine learning algorithms for brain tumor segmentation, progression assessment, and overall survival prediction in the BRATS challenge. In *arXiv:1811.02629*, 2018.
- [5] Jose Dolz, Christian Desrosiers, and Ismail Ben Ayed. 3d fully convolutional networks for subcortical segmentation in mri: A large-scale study. *NeuroImage*, 170:456–470, 2018.
- [6] Robert Geirhos, Patricia Rubisch, Claudio Michaelis, Matthias Bethge, Felix A. Wichmann, and Wieland Brendel. Imagenet-trained CNNs are biased towards texture; increasing shape bias improves accuracy and robustness. In *International Conference on Learning Representations (ICLR)*, 2019.
- [7] Eli Gibson, Francesco Giganti, Yipeng Hu, Ester Bonmati, Steve Bandula, Kurinchi Gurusamy, Brian Davidson, Stephen P Pereira, Matthew J Clarkson, and Dean C Barratt. Automatic multi-organ segmentation on abdominal ct with dense v-networks. *IEEE transactions on medical imaging*, 37(8):1822–1834, 2018.
- [8] Ali Hatamizadeh, Assaf Hoogi, Debleena Sengupta, Wuyue Lu, Brian Wilcox, Daniel Rubin, and Demetri Terzopoulos. Deep active lesion segmentation. *arXiv preprint arXiv:1908.06933*, 2019.
- [9] Ali Hatamizadeh, Hamid Hosseini, Zhengyuan Liu, Steven D Schwartz, and Demetri Terzopoulos. Deep dilated convolutional nets for the automatic segmentation of retinal vessels. *arXiv preprint arXiv:1905.12120*, 2019.

- [10] Ali Hatamizadeh, Debleena Sengupta, and Demetri Terzopoulos. End-to-end deep convolutional active contours for image segmentation. *arXiv preprint arXiv:1909.13359*, 2019.
- [11] Nicholas Heller, Niranjana Sathianathan, Arveen Kalapara, Edward Walczak, Keenan Moore, Heather Kaluzniak, Joel Rosenberg, Paul Blake, Zachary Rengel, Makinna Oestreich, et al. The kits19 challenge data: 300 kidney tumor cases with clinical context, ct semantic segmentations, and surgical outcomes. *arXiv preprint arXiv:1904.00445*, 2019.
- [12] Eric Jang, Shixiang Gu, and Ben Poole. Categorical reparameterization with gumbel-softmax. *arXiv preprint arXiv:1611.01144*, 2016.
- [13] Konstantinos Kamnitsas, Christian Ledig, Virginia FJ Newcombe, Joanna P Simpson, Andrew D Kane, David K Menon, Daniel Rueckert, and Ben Glocker. Efficient multi-scale 3d cnn with fully connected crf for accurate brain lesion segmentation. *Medical image analysis*, 36:61–78, 2017.
- [14] Bjoern H. Menze et al. The multimodal brain tumor image segmentation benchmark (brats). *IEEE Trans. Med. Imaging*, 34(10):1993–2024, 2015.
- [15] Fausto Milletari, Nassir Navab, and Seyed-Ahmad Ahmadi. V-net: Fully convolutional neural networks for volumetric medical image segmentation. In *Fourth International Conference on 3D Vision (3DV)*, 2016.
- [16] Andriy Myronenko. 3D MRI brain tumor segmentation using autoencoder regularization. In *BrainLes, Medical Image Computing and Computer Assisted Intervention (MICCAI)*, LNCS, pages 311–320. Springer, 2018. URL <https://arxiv.org/abs/1810.11654>.
- [17] Andriy Myronenko and Ali Hatamizadeh. 3d kidneys and kidney tumor semantic segmentation using boundary-aware networks. *arXiv preprint arXiv:1909.06684*, 2019.
- [18] Andriy Myronenko and Ali Hatamizadeh. Robust semantic segmentation of brain tumor regions from 3d mris. *arXiv preprint arXiv:2001.02040*, 2020.
- [19] O. Ronneberger, P. Fischer, and T. Brox. U-net: Convolutional networks for biomedical image segmentation. In *Medical Image Computing and Computer-Assisted Intervention (MICCAI)*, volume 9351 of *LNCS*, pages 234–241. Springer, 2015.
- [20] Towaki Takikawa, David Acuna, Varun Jampani, and Sanja Fidler. Gated-SCNN: Gated shape cnns for semantic segmentation. *ICCV*, 2019.
- [21] Zhiding Yu, Chen Feng, Ming-Yu Liu, and Srikumar Ramalingam. Casenet: Deep category-aware semantic edge detection. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 5964–5973, 2017.
- [22] Wentao Zhu, Yufang Huang, Liang Zeng, Xuming Chen, Yong Liu, Zhen Qian, Nan Du, Wei Fan, and Xiaohui Xie. Anatomynet: Deep learning for fast and fully automated whole-volume segmentation of head and neck anatomy. *Medical physics*, 46(2):576–589, 2019.