Boundary loss for highly unbalanced segmentation

Hoel Kervadec, Jihene Bouchtiba, Christian Desrosiers
Eric Granger, Jose Dolz, Ismail Ben Ayed
ÉTS Montréal

Contributions

We propose a *boundary* loss that takes the form of a distance metric on the space of contours (or shapes), not regions. We argue that a boundary loss can mitigate the issues related to regional losses in highly unbalanced segmentation problems. Rather than using unbalanced integrals over the regions, a boundary loss uses integrals over the boundary (interface) between the regions. Furthermore, a boundary loss provides information that is complimentary to regional losses.

Formulation

Let $I : \Omega \subset \mathbb{R}^{2,3} \to \mathbb{R}$ denotes a training image with spatial domain Ω, and $g : \Omega \to \{0, 1\}$ a binary ground-truth segmentation of the image: $g(p) = 1$ if pixel/voxel p belongs to the target region $G \subset \Omega$ (foreground region) and 0 otherwise, i.e., $p \in \Omega \setminus G$ (background region). Let $s_\theta : \Omega \to [0, 1]$ denotes the softmax probability output of a deep segmentation network, and $S_\theta \subset \Omega$ the corresponding segmentation region: $S_\theta = \{p \in \Omega \,|\, s_\theta(p) \geq \delta\}$ for some threshold δ.

Our purpose is to build a boundary loss $\mathrm{Dist}(\partial G, \partial S_\theta)$, which takes the form of a distance metric on the space of contours (or region boundaries) in Ω, with ∂G denoting a representation of the boundary of ground-truth region G (e.g., the set of points of G, which have a spatial neighbor in background $\Omega \setminus G$) and ∂S_θ denoting the boundary of the segmentation region defined by the network output. However, it is not clear how to represent boundary points on ∂S_θ as a differentiable function of regional network outputs s_θ.

Our boundary loss is inspired from discrete (graph-based) optimization techniques for computing gradient flows of curve evolution [1]. Curve evolution methods require a measure for evaluating boundary changes (or variations). Consider the following non-symmetric L_2 distance on the space of shapes, which evaluates the change between two nearby boundaries ∂S and ∂G [1]:

$$\mathrm{Dist}(\partial G, \partial S) = \int_{\partial G} \|q_{\partial S}(p) - p\|^2 dp \tag{1}$$

where $p \in \Omega$ is a point on boundary ∂G and $q_{\partial S}(p)$ denotes the corresponding point on boundary ∂S, along the direction normal to ∂G, i.e., $q_{\partial S}(p)$ is the intersection of ∂S and the line that is normal to ∂G at p (See Fig. 1.a for an illustration). Similarly to any contour distance invoking directly points on contour ∂S, expression (1) cannot be used directly as a loss for $\partial S = \partial S_\theta$. However, it is easy to show that the differential boundary variation in (1) can be expressed using an *integral* approach [1]:

$$\mathrm{Dist}(\partial G, \partial S) = 2 \int_{\Delta S} \phi_G(p) dp \tag{2}$$

where ΔS denotes the region between the two contours and $\phi_G : \Omega \to \mathbb{R}$ is a *level set* representation of boundary ∂G: $\phi_G(p)$ evaluates a signed distance between point $p \in \Omega$ and the nearest point $z_{\partial G}(p)$ on contour ∂G: $\phi_G(p) = -\|p - z_{\partial G}(p)\|$ if $p \in G$ and $\phi_G(p) = \|p - z_{\partial G}(p)\|$ otherwise. Fig. 1.b illustrates this integral framework for evaluating the boundary distance in (1). To show that (2) holds, it suffices to notice that integrating the distance function $2\phi_G(p)$ over the normal segment connecting p and $q_{\partial S}(p)$ yields $\|q_{\partial S}(p) - p\|^2$. Thus, the non-symmetric L_2 distance between contours in Eq. (1) can be expressed as a sum of regional integrals:

$$\int_S \phi_G(p)dp - \int_G \phi_G(p)dp = \int_\Omega \phi_G(p)s(p)dp - \int_\Omega \phi_G(p)g(p)dp \tag{3}$$

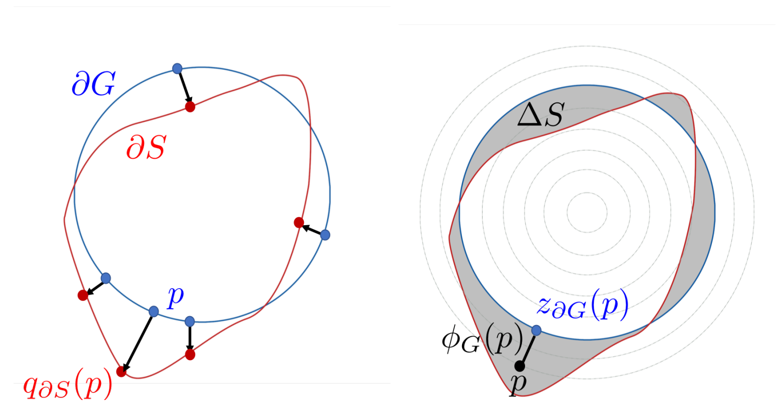

(a) Differential (b) Integral

Figure 1: The relationship between *differential* and *integral* approaches for evaluating boundary change (variation).

where $s : \Omega \rightarrow \{0, 1\}$ is binary indicator function of region S: $s(p) = 1$ if $p \in S$ belongs to the target and 0 otherwise. Now, for $S = S_\theta$, i.e., replacing binary variables $s(p)$ in Eq. (3) by the softmax probability outputs of the network $s_\theta(p)$, we obtain the following boundary loss which, up to a constant independent of θ, approximates boundary distance $\text{Dist}(\partial G, \partial S_\theta)$:

$$\mathcal{L}_B(\theta) = \int_\Omega \phi_G(p) s_\theta(p) dp \tag{4}$$

The level set function ϕ_G is pre-computed directly from the ground-truth region G. It can be easily combined with standard regional losses and implemented with any existing deep network architecture for N-D segmentation. In the experiments, we will use our boundary loss in conjunction with the regional generalized Dice loss:

$$\alpha \mathcal{L}_{GD}(\theta) + (1 - \alpha) \mathcal{L}_B(\theta) \tag{5}$$

Architecture and training

We employed UNet [2] as deep learning architecture in our experiments. To train our model, we employed Adam optimizer, with a learning rate of 0.001 and a batch size equal to 8. The learning rate is halved if the validation performances do not improve during 20 epochs.

To compute the level set function ϕ_G in Eq. (4), we used standard SciPy functions [1]. Note that, for slices containing only the background region, we used a zero-distance map, assuming that the GDL is sufficient in those cases. Furthermore, during training, the value of α in Eq. (5) was initially set to 1, and decreased by 0.01 after each epoch, following a simple scheduling strategy, until it reached the value of 0.01. In this way, we give more importance to the regional loss term at the beginning while gradually increasing the impact of the boundary loss term.

References

[1] Yuri Boykov, Vladimir Kolmogorov, Daniel Cremers, and Andrew Delong. An integral solution to surface evolution PDEs via geo-cuts. In *European Conference on Computer Vision*, pages 409–422. Springer, 2006.

[2] Olaf Ronneberger, Philipp Fischer, and Thomas Brox. U-Net: Convolutional networks for biomedical image segmentation. In *International Conference on Medical image computing and computer-assisted intervention*, pages 234–241. Springer, 2015.

[1] `https://docs.scipy.org/doc/scipy-0.14.0/reference/generated/scipy.ndimage.morphology.distance_transform_edt.html`