Computational Cardiology Lecture 12

Local convolutional segmentation and geometric error

Zan Ahmad

Fall 2024

The output of segmentation

$$I \longmapsto \widehat{M} \in \{0,1\}^{H \times W}.$$

The output of segmentation

$$I \longmapsto \widehat{M} \in \{0,1\}^{H \times W}.$$

A spatial mask contains boundary and connectivity information absent from one image-level label.

A local shared predictor

$$z_i = w^{\mathsf{T}} x_i + b, \qquad p_i = \sigma(z_i).$$

A local shared predictor

$$z_i = w^{\mathsf{T}} x_i + b, \qquad p_i = \sigma(z_i).$$

The 3×3 patch supplies nine weights and one shared bias.

The sigmoid decision

$$\sigma(z) = \frac{1}{1 + e^{-z}}, \qquad \widehat{y}_i = \mathbf{1}_{\{p_i \geq .5\}}.$$

The sigmoid decision

$$\sigma(z) = \frac{1}{1+e^{-z}}, \qquad \widehat{y}_i = \mathbf{1}_{\{p_i \geq .5\}}.$$

With threshold $.5$, the class boundary is $z_i = 0$.

Receptive field and padding

The rule uses only a 3×3 neighborhood. Reflect padding supplies border context.

Receptive field and padding

The rule uses only a 3×3 neighborhood. Reflect padding supplies border context. A distant anatomical relationship lies outside this receptive field.

Binary cross-entropy from logits

$$J = \frac{1}{M} \sum_i [\log(1 + e^{z_i}) - y_i z_i] + \frac{\lambda}{2} \|w\|^2.$$

Binary cross-entropy from logits

$$J = \frac{1}{M} \sum_i [\log(1 + e^{z_i}) - y_i z_i] + \frac{\lambda}{2} \|w\|^2.$$

A stable log-add-exp evaluation avoids taking logarithms of rounded probabilities.

Differentiate the data term

$$\frac{\partial \ell_i}{\partial z_i} = p_i - y_i.$$

Differentiate the data term

$$\frac{\partial \ell_i}{\partial z_i} = p_i - y_i.$$

$$\nabla_w J = \frac{1}{M} X^{\mathsf{T}}(p - y) + \lambda w, \quad \partial_b J = \frac{1}{M} \mathbf{1}^{\mathsf{T}}(p - y).$$

Curvature of the local model

$$\nabla^2 \ell = \frac{1}{M} \widetilde{X}^{\mathsf{T}} \operatorname{diag}(p_i(1-p_i)) \widetilde{X} \succeq 0.$$

Curvature of the local model

$$\nabla^2 \ell = \frac{1}{M} \widetilde{X}^{\mathsf{T}} \operatorname{diag}(p_i(1-p_i)) \widetilde{X} \succeq 0.$$

The single-layer objective is convex; weight sharing does not make it a deep network.

A gradient update

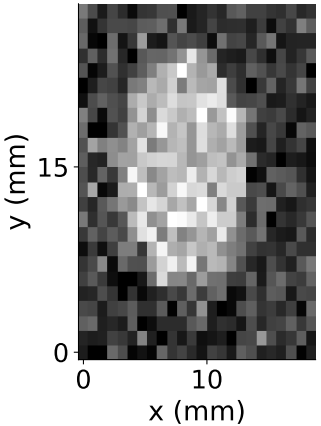
$$w^{k+1} = w^k - \eta \nabla_w J, \qquad b^{k+1} = b^k - \eta \partial_b J.$$

A gradient update

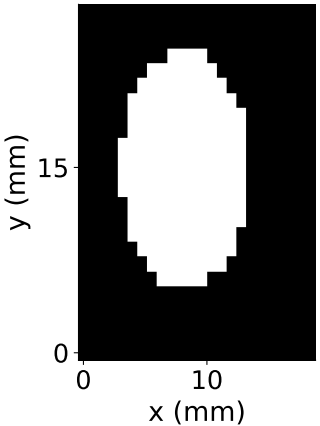
$$w^{k+1} = w^k - \eta \nabla_w J, \qquad b^{k+1} = b^k - \eta \partial_b J.$$

Validation chooses among saved iterates without using the test set.

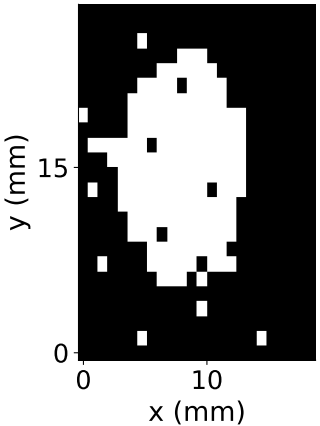
Synthetic test image



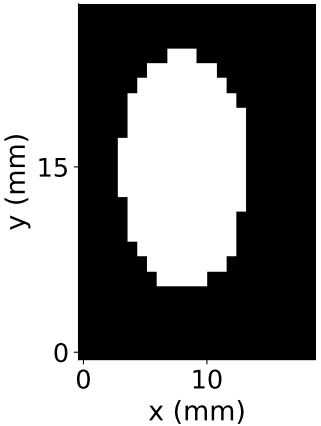
Generator reference



Fixed threshold 0.5



Learned 3×3 + sigmoid



The unit of separation

Complete generated images remain in one split.

The unit of separation

Complete generated images remain in one split. Neighboring patches from one image share both pixels and geometry; splitting pixels would answer an easier prediction problem.

Overlap as a set quantity

$$\mathrm{Dice}(P, R) = \frac{2|P \cap R|}{|P| + |R|}.$$

Overlap as a set quantity

$$\mathrm{Dice}(P, R) = \frac{2|P \cap R|}{|P| + |R|}.$$

Equal overlap need not imply equal physical boundary error.

Physical boundary distance

Compute nearest-boundary distances in both directions using physical coordinates.

Physical boundary distance

Compute nearest-boundary distances in both directions using physical coordinates. Average distance describes typical displacement; Hausdorff distance records the largest nearest-boundary discrepancy.

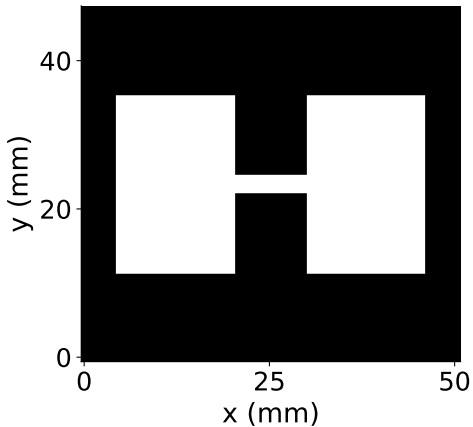
Empty-set conventions

Two empty masks have Dice one in this implementation. One empty mask has Dice zero.

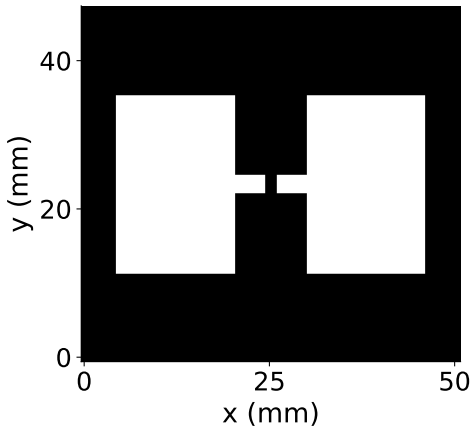
Empty-set conventions

Two empty masks have Dice one in this implementation. One empty mask has Dice zero. Boundary distance is unavailable when only one boundary is empty.

Reference: 1 component



Four missing pixels: 2 components



A nearly perfect overlap can sever a bridge

$$|R| = 824, \quad |P| = 820, \quad \text{Dice} = \frac{1640}{1644} \approx .99757.$$

A nearly perfect overlap can sever a bridge

$$|R| = 824, \quad |P| = 820, \quad \text{Dice} = \frac{1640}{1644} \approx .99757.$$

Four deleted bridge pixels change one connected component into two.

Area bias and overlap are distinct

$$e_A = \frac{A_P - A_R}{A_R}.$$

Area bias and overlap are distinct

$$e_A = \frac{A_P - A_R}{A_R}.$$

A signed mean can conceal cancellation of over- and under-segmentation across cases.

The downstream geometric requirement

A transport model depends on the paths and boundaries of its domain.

The downstream geometric requirement

A transport model depends on the paths and boundaries of its domain. An overlap score cannot by itself certify preservation of those paths.