Lyon, Avril 17– 21

DLMI2023

CREATIS

UNIVERSITÉ DE LYON

ÉTS

Le génie pour l'industrie

Weakly Supervised Segmentation

Jose Dolz

ÉTS, Montreal

Motivation

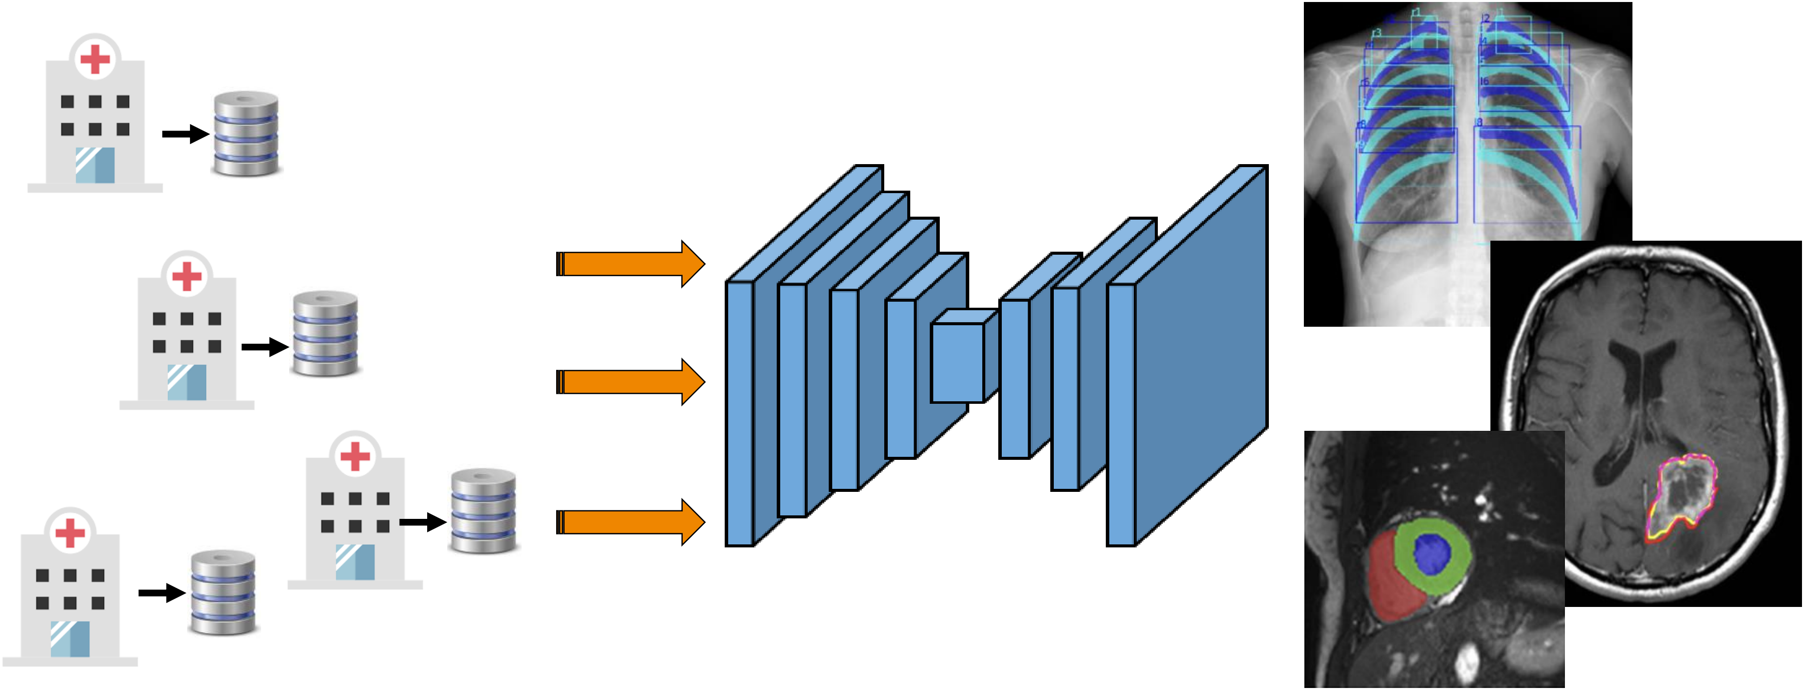

Large labeled (pixel-wise) datasets

Very good performance in many tasks

Motivation

Pixel-wise annotation is a time-consuming task…

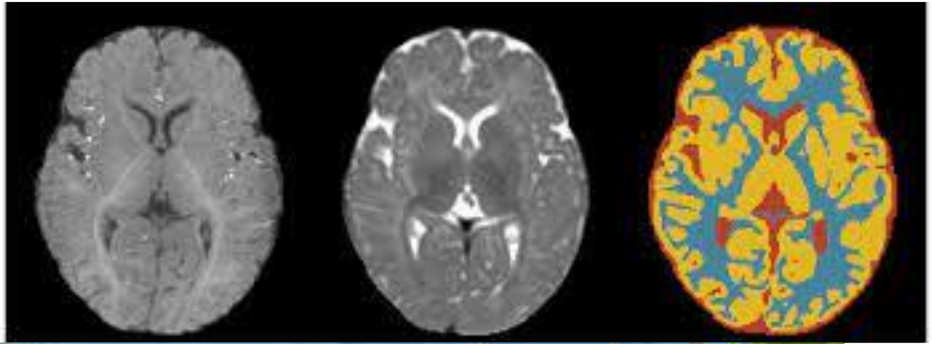

Motivation

… which further requires domain expertise

Which images contain the class ‘*esophagus*’?

Motivation

… which further requires domain expertise

Which images contain the class '*esophagus*'?

Motivation

… which further requires domain expertise

Which images contain the class ‘*esophagus*’?

And where on these images the esophagus is?

Motivation

… which further requires domain expertise

Which images contain the class ‘*esophagus*’?

And where on these images the esophagus is?

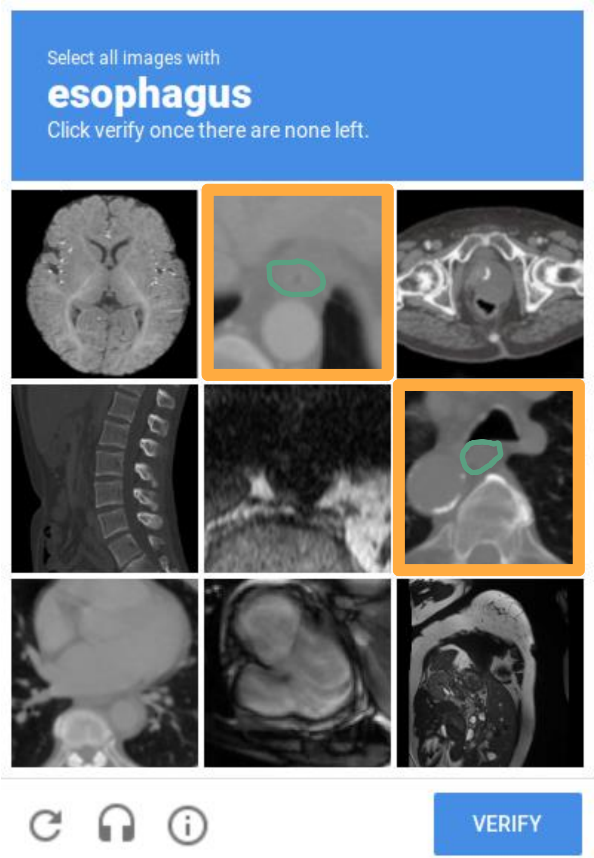

Data-driven priors (cues)

Image tags

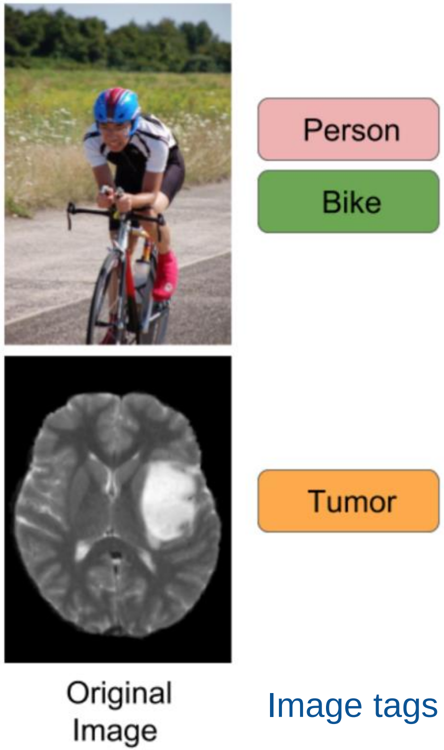

- Pathak et al., Constrained convolutional neural networks for weakly supervised segmentation, ICCV 2015
- Kervadec et al., Constrained-CNN losses for weakly supervised segmentation, MedIA 2019.

Data-driven priors (cues)

Image tags

Bounding boxes

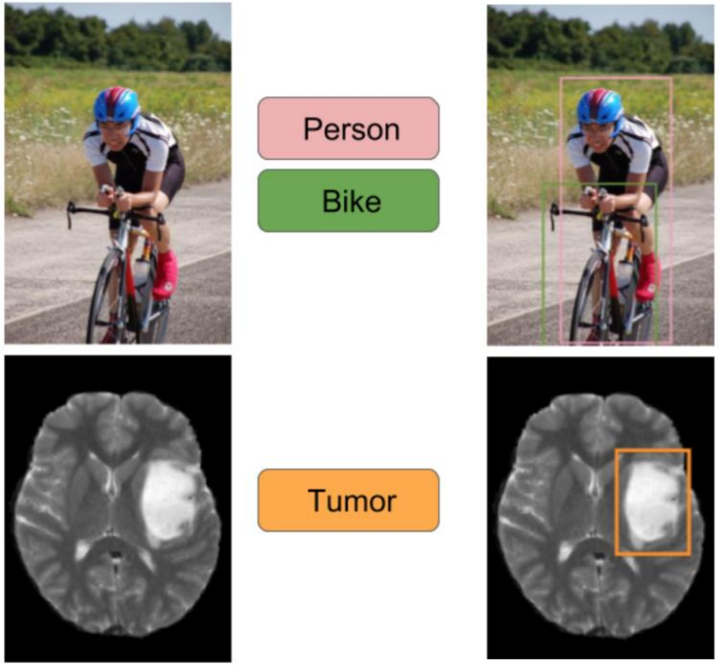

Original Image | Image tags | Bounding boxes

- Pathak et al., Constrained convolutional neural networks for weakly supervised segmentation, ICCV 2015
- Kervadec et al., Constrained-CNN losses for weakly supervised segmentation, MedIA 2019.

Data-driven priors (cues)

Image tags

Bounding boxes

Scribbles

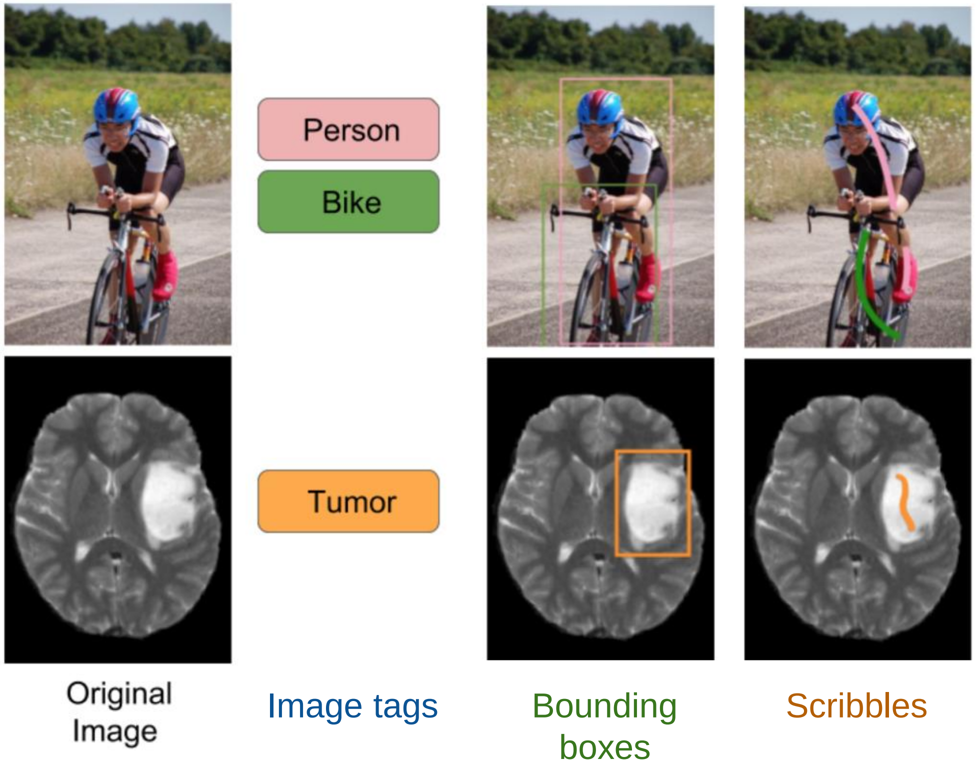

- Pathak et al., Constrained convolutional neural networks for weakly supervised segmentation, ICCV 2015
- Kervadec et al., Constrained-CNN losses for weakly supervised segmentation, MedIA 2019.

Data-driven priors (cues)

Image tags

Bounding boxes

Scribbles

Points

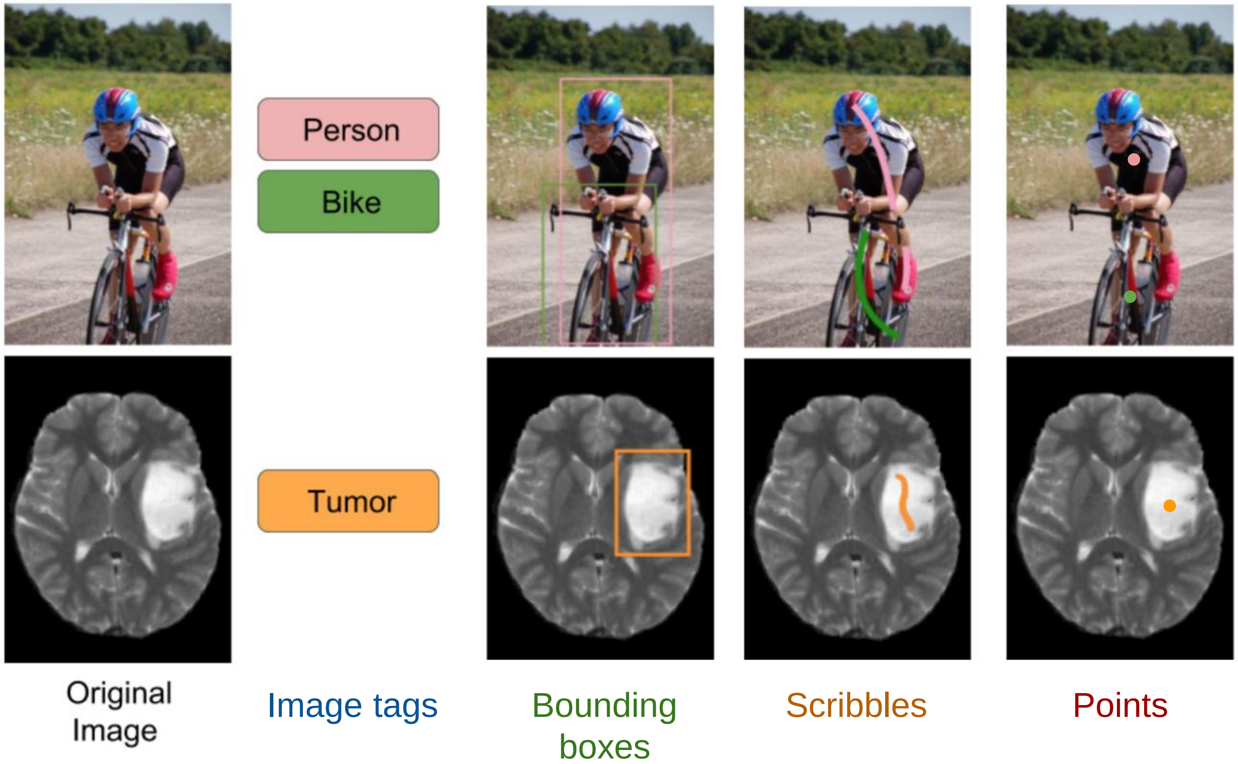

- Pathak et al., Constrained convolutional neural networks for weakly supervised segmentation, ICCV 2015
- Kervadec et al., Constrained-CNN losses for weakly supervised segmentation, MedIA 2019.

Data-driven priors (cues)

Another data-driven priors

Image captions

A boy jumping on a skateboard

Extreme points

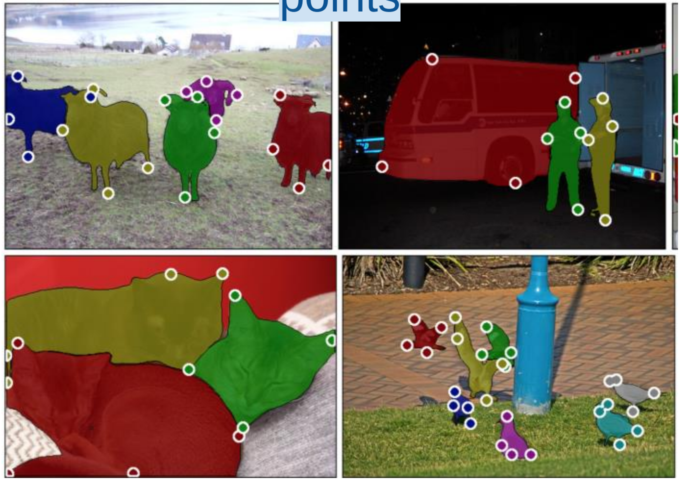

Image from Maninis et al, CVPR'18

- Maninis et al. Deep extreme cut: From extreme points to object segmentation. CVPR 2018

Knowledge-driven priors

Common priors in natural images

Target Size

Incorrect sizes

Correct sizes

- Pathak et al., Constrained convolutional neural networks for weakly supervised segmentation, ICCV 2015
- Xu et al., Learning to Segment Under Various Forms of Weak Supervision, CVPR 2015
- Zhang et al., Curriculum Domain Adaptation for Semantic Segmentation of Urban Scenes. ICCV'17

Knowledge-driven priors

Common priors in natural images

Incorrect location

Target Location

Correct location

- Remez et al. Learning to segment via cut-and-paste. ECCV 2018
- Georgakis et al Synthesizing training data for object detection in indoor scenes. RSS 2017

Knowledge-driven priors

Common priors in natural images

Contrast
Foreground/Background

Saliency

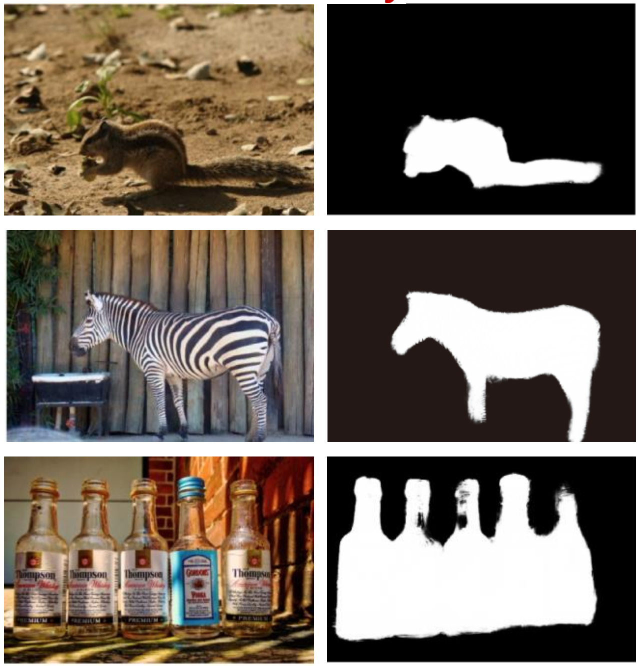

Images from Hou et al, CVPR'17

- Hou et al. Deeply supervised salient object detection with short connections. CVPR 2017
- Li et al. Instance-level salient object segmentation. CVPR 2017

Knowledge-driven priors

Common priors in natural images

Motion

Images from the DAVIS Challenge Dataset

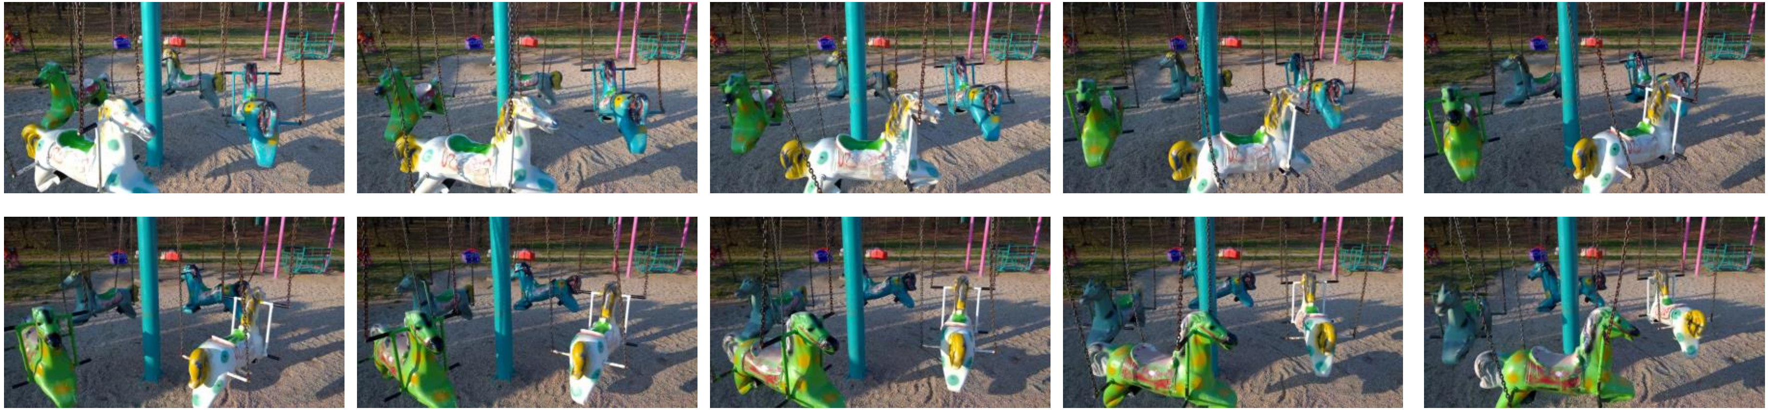

- Tokmakov et al. Weakly-supervised semantic segmentation using motion cues. ECCV 2016
- Pathak et al. Learning features by watching objects move. CVPR 2017

Knowledge-driven priors

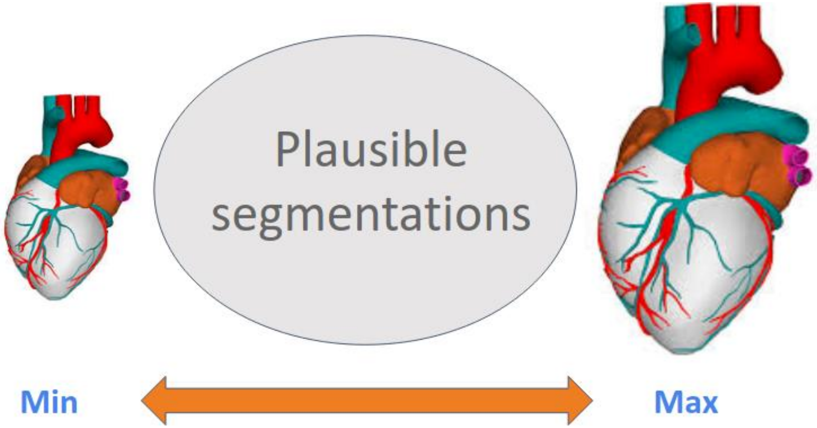

Anatomical priors

What about priors in the medical domain?

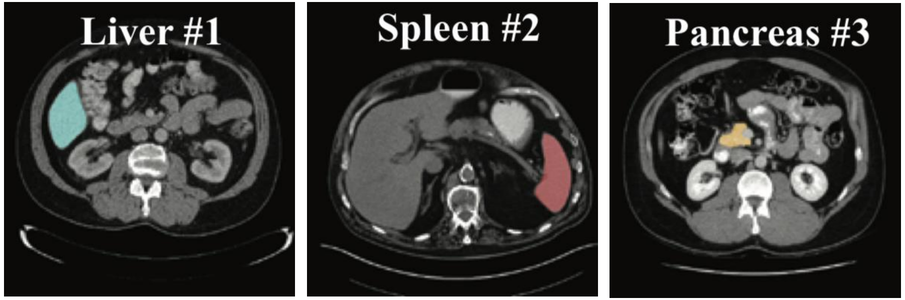

Partial labeled data
(exploit target relationships)

Recall of what learning means (from a gradient descent standpoint)…

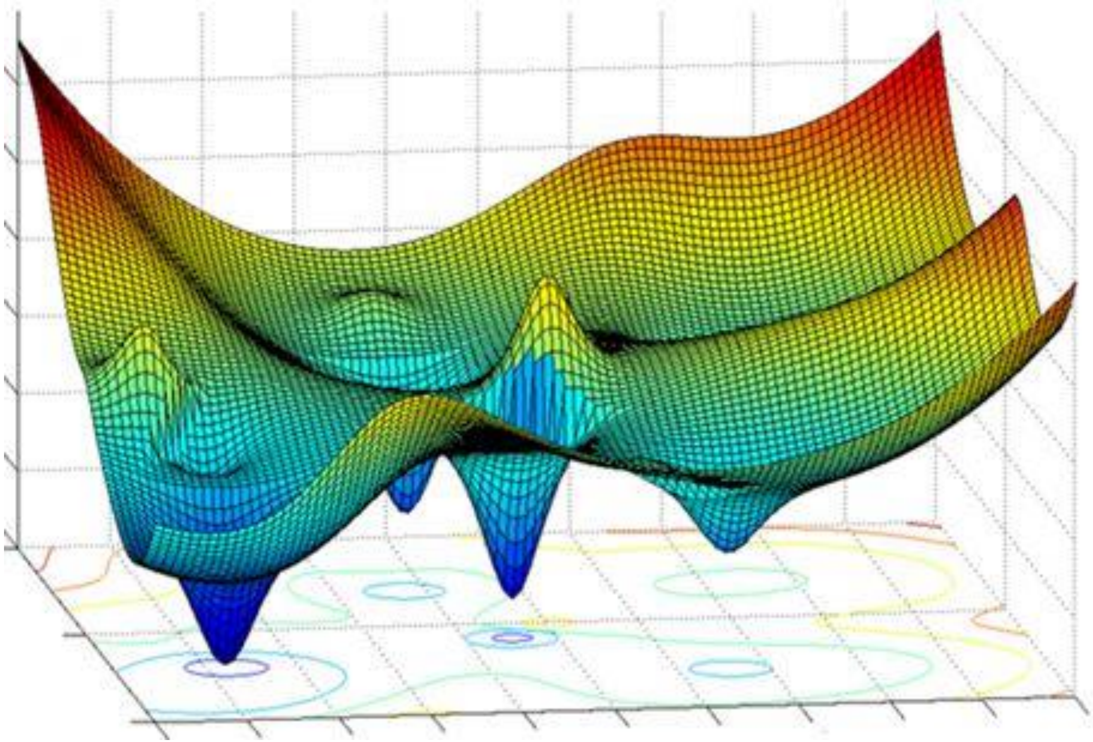

Knowledge vs data driven priors

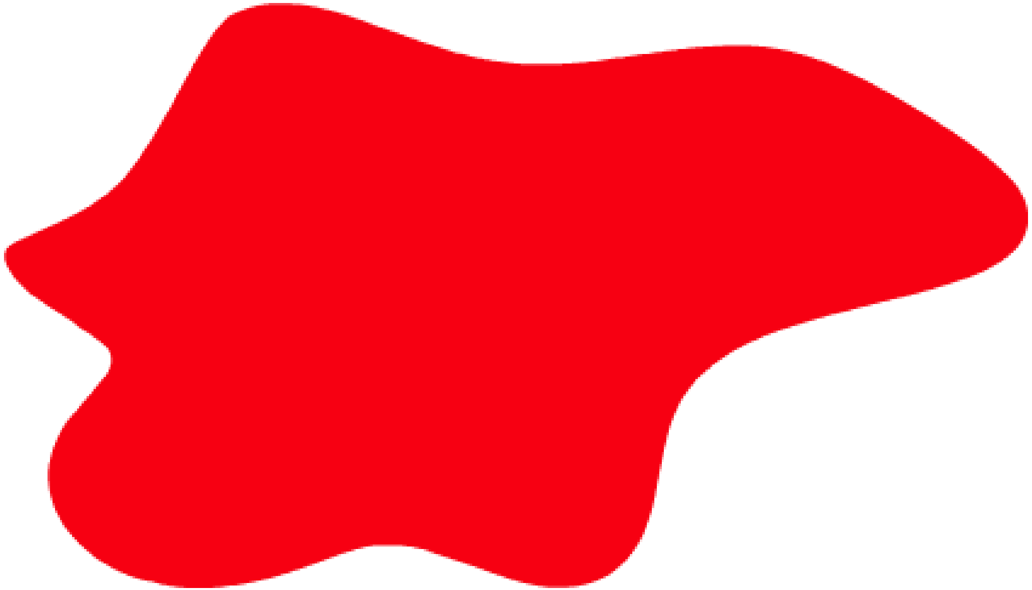

Knowledge vs data driven priors

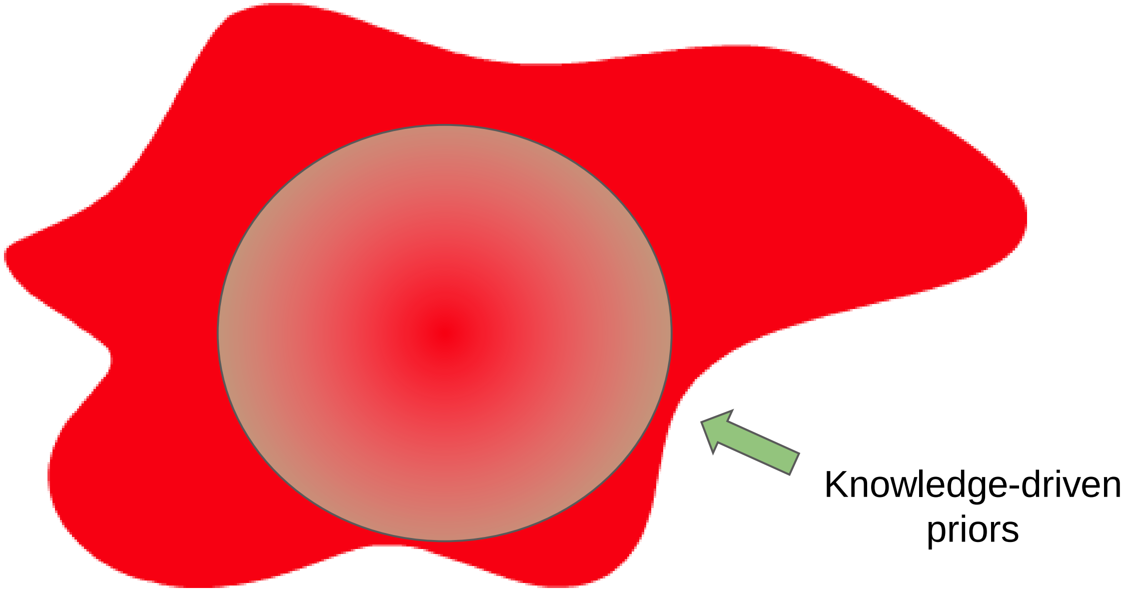

Knowledge vs data driven priors

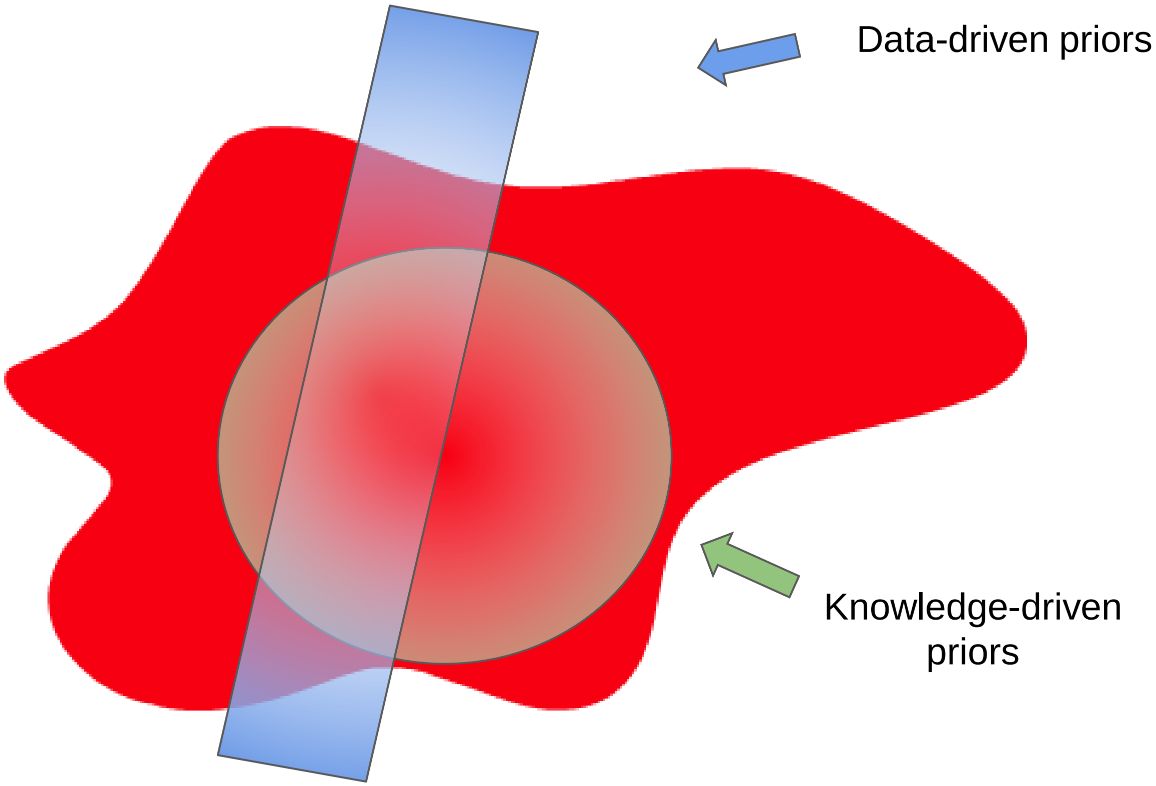

Knowledge vs data driven priors

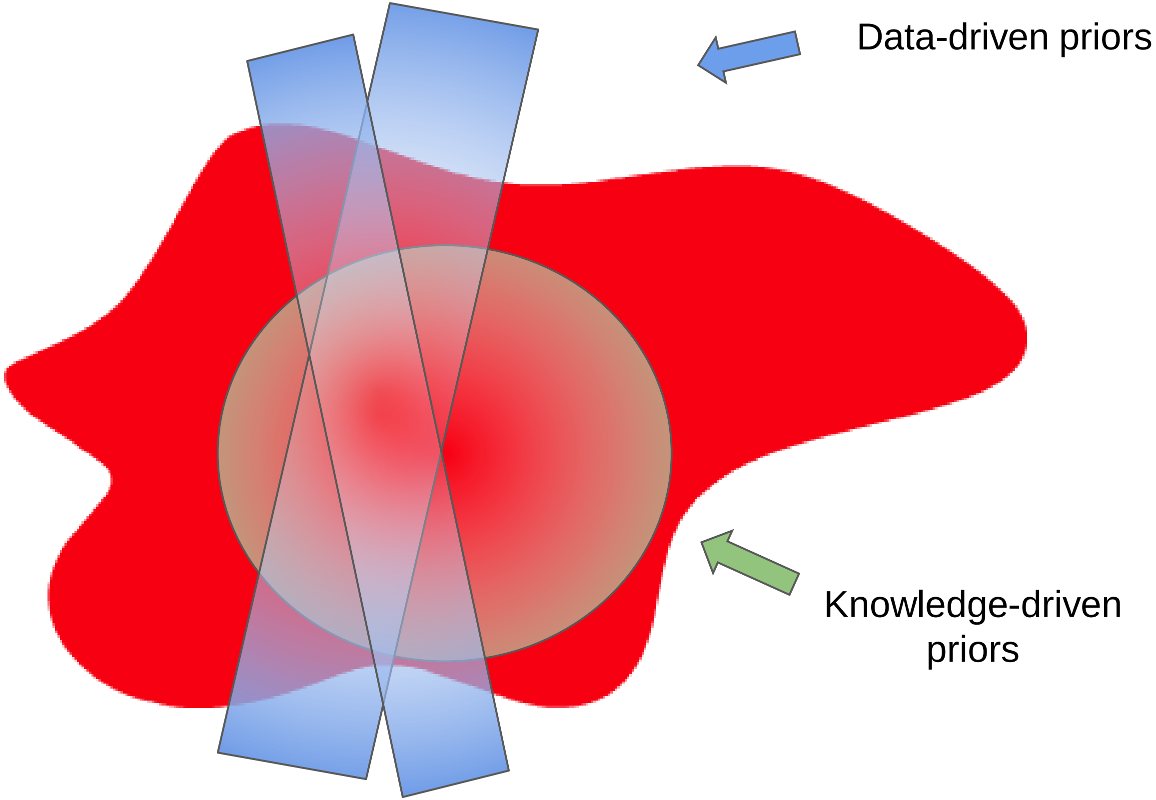

Knowledge vs data driven priors

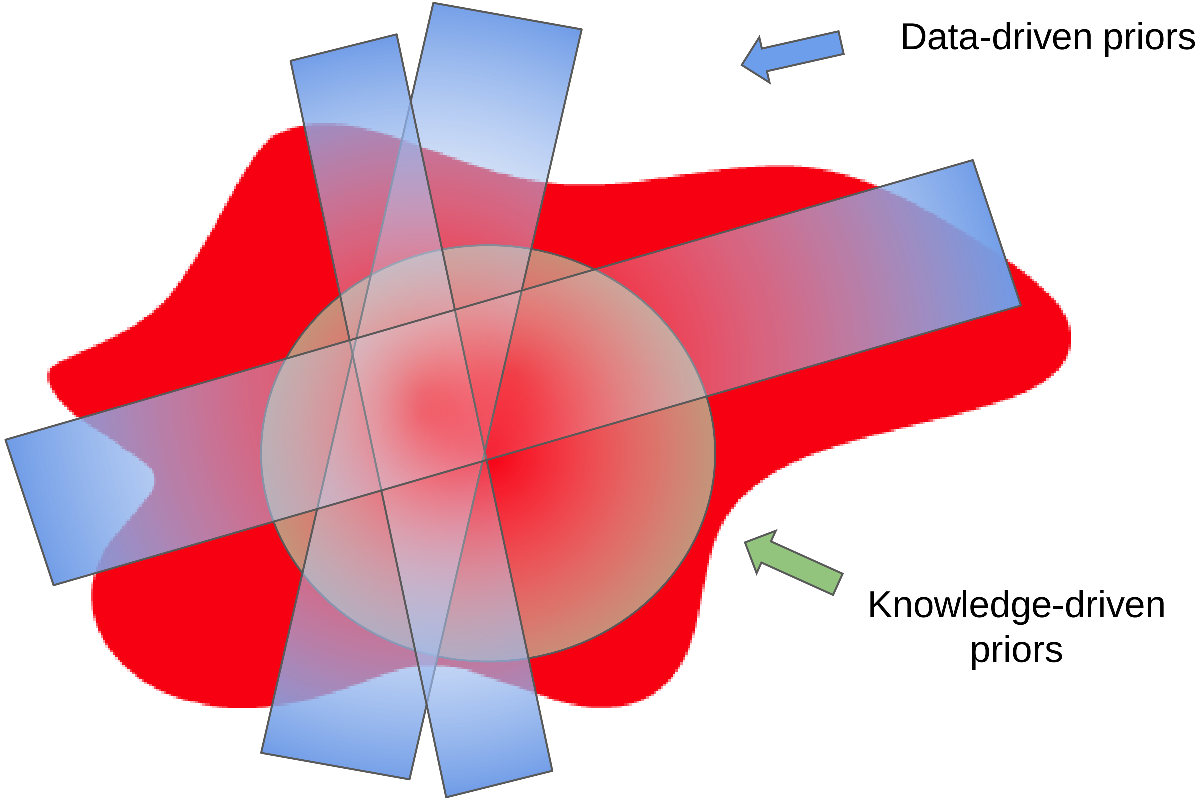

Knowledge vs data driven priors

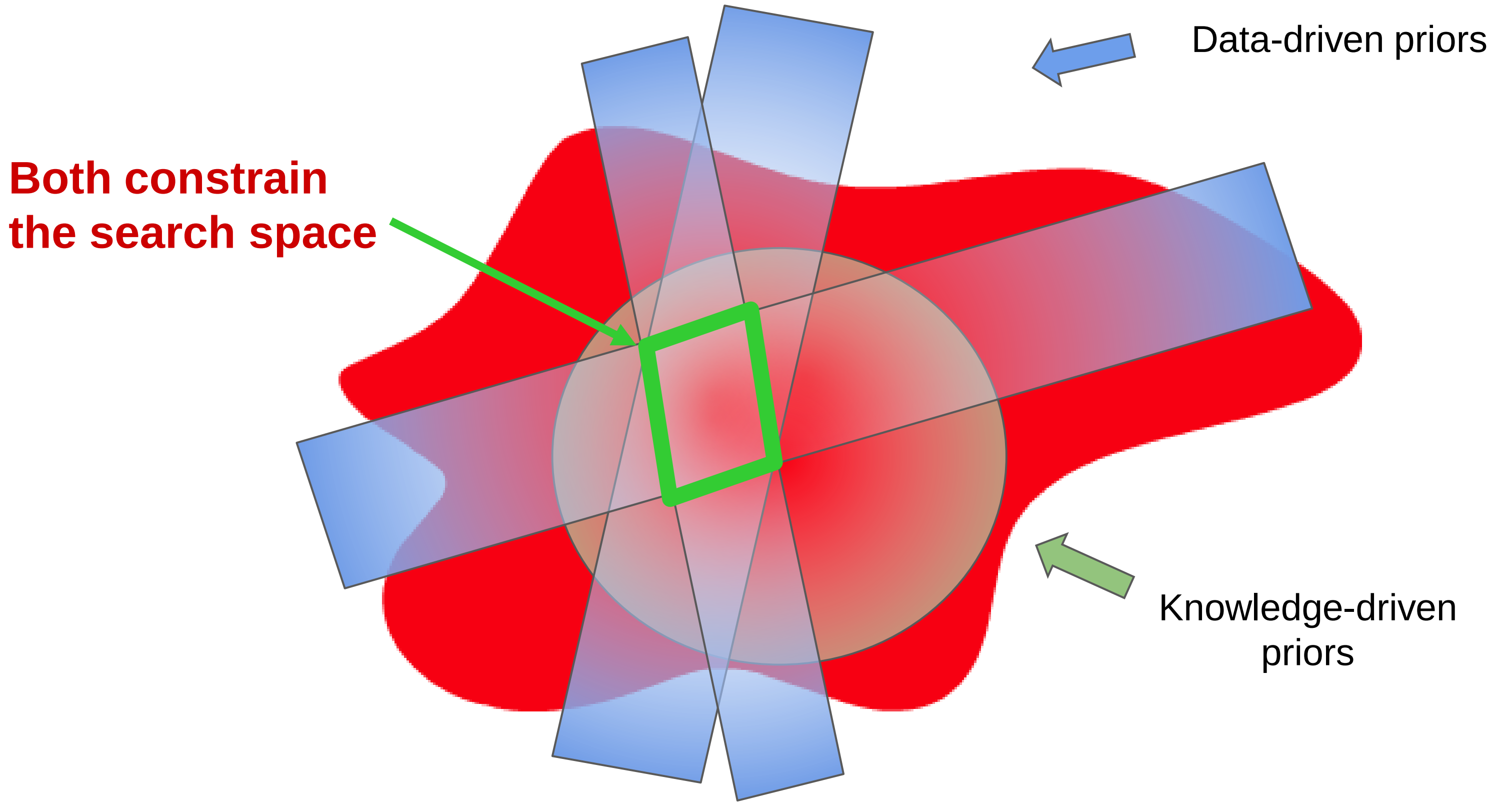

From global cues to pixel labels

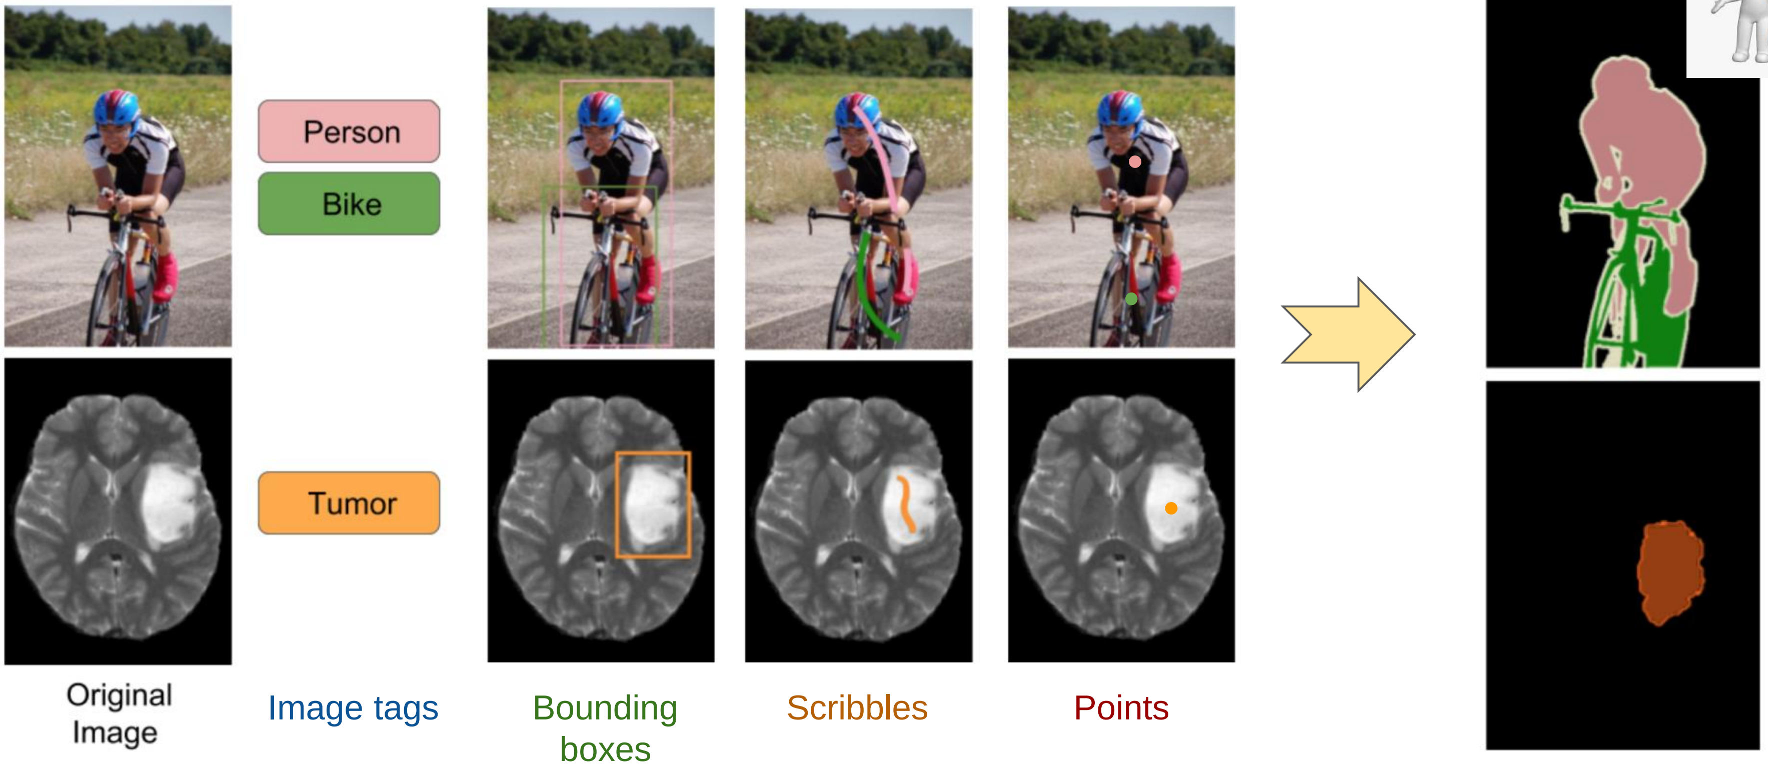

From global cues to pixel labels

Step 1: Get a classification CNN

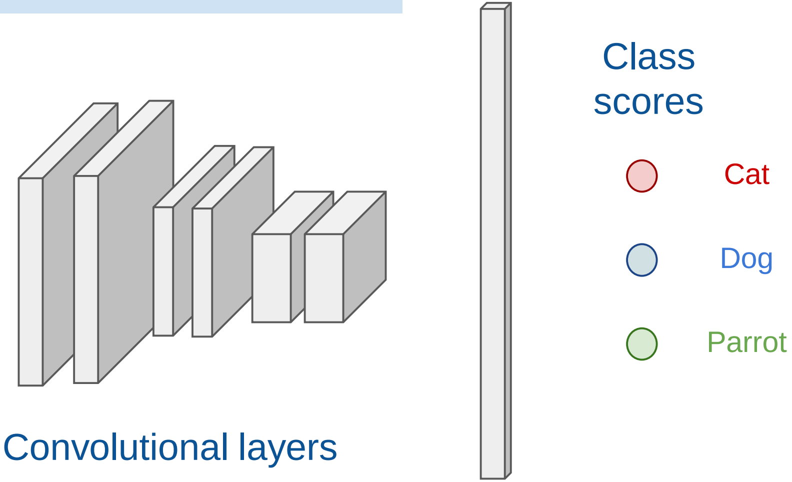

From global cues to pixel labels

Step 2: Modify the last layers

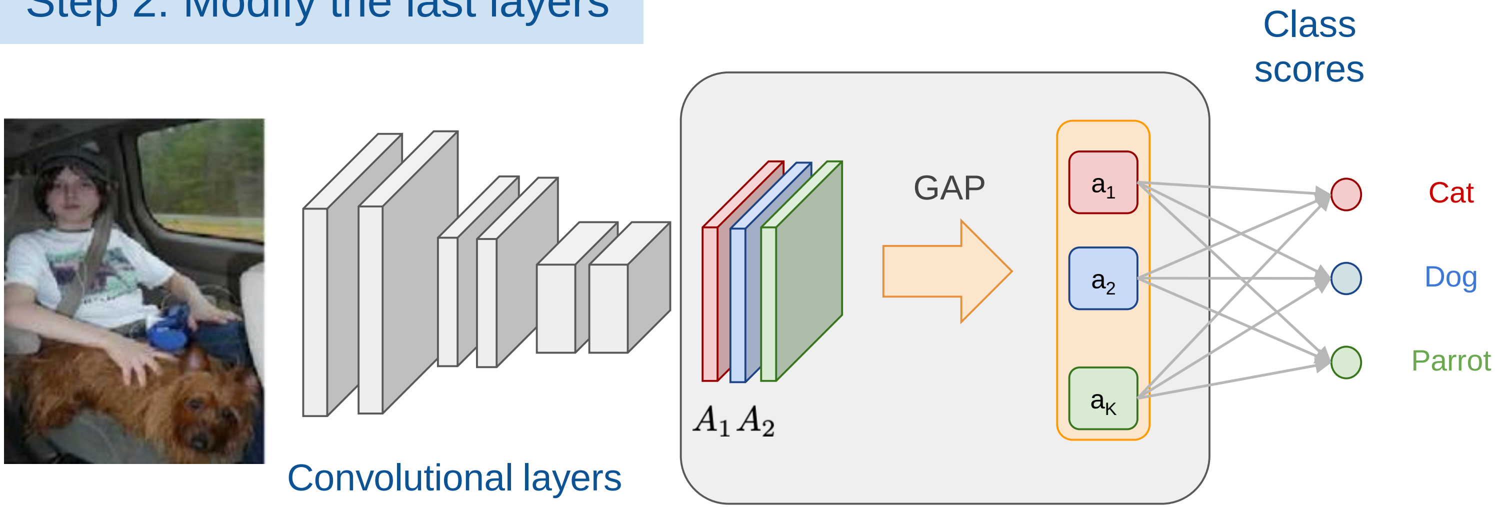

$$GAP(A_k) = a_k = \frac{1}{|N|}\sum_{x,y} A_k(x,y)$$

From global cues to pixel labels

Step 2: Modify the last layers

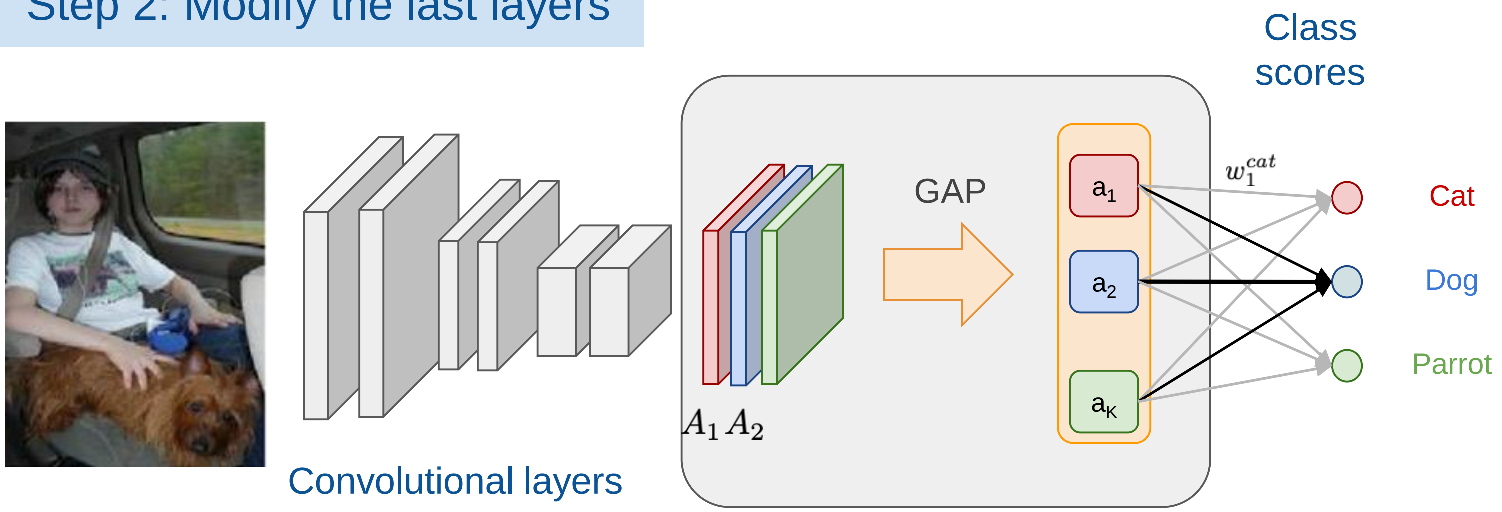

$$GAP(A_k) = a_k = \frac{1}{|N|}\sum_{x,y} A_k(x,y)$$

Class score
(logits)

$$S_c = \sum_k w_k^c a_k = \frac{1}{N}\sum_k w_k^c \sum_{x,y} a_k(x,y)$$

From global cues to pixel labels

Step 3: Get the CAMs

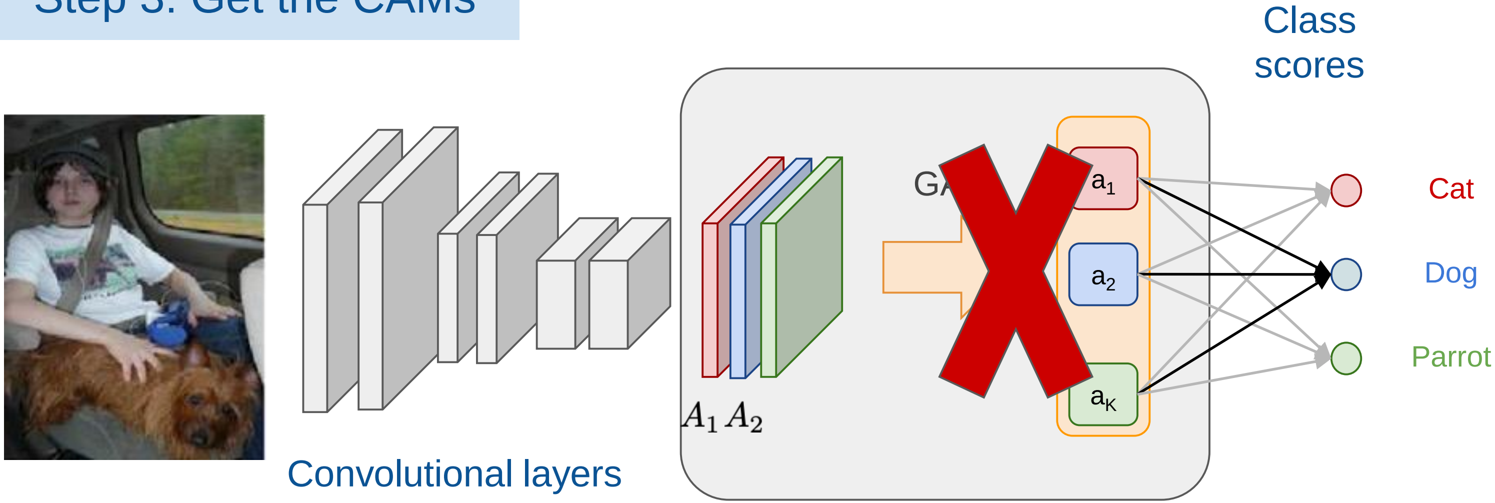

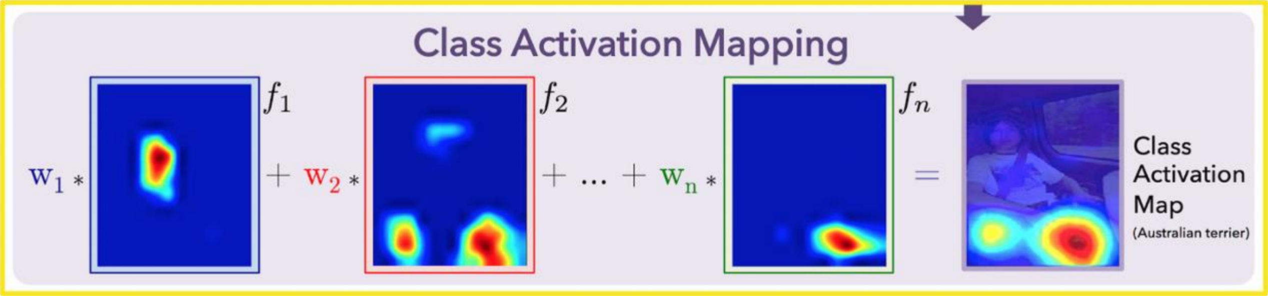

From global cues to pixel labels

Step 3: Get the CAMs

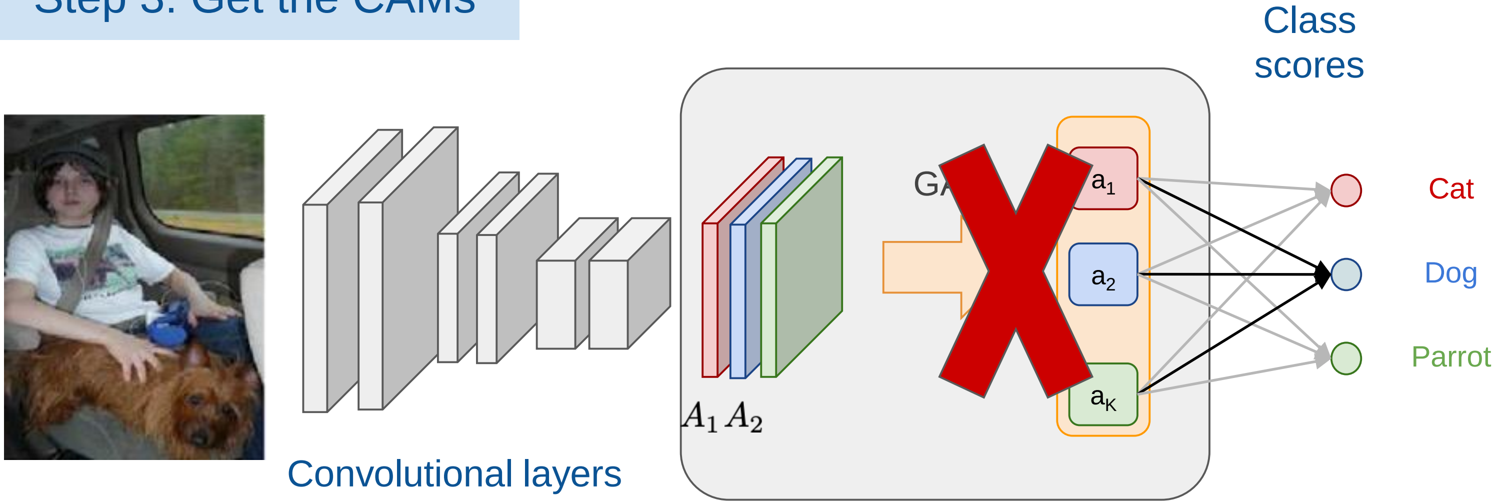

$$CAM_{Dog}(x,y) = \sum_k w_k^{Dog} A_k(x,y) =$$

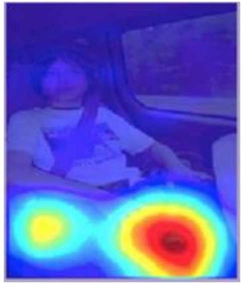

From global cues to pixel labels

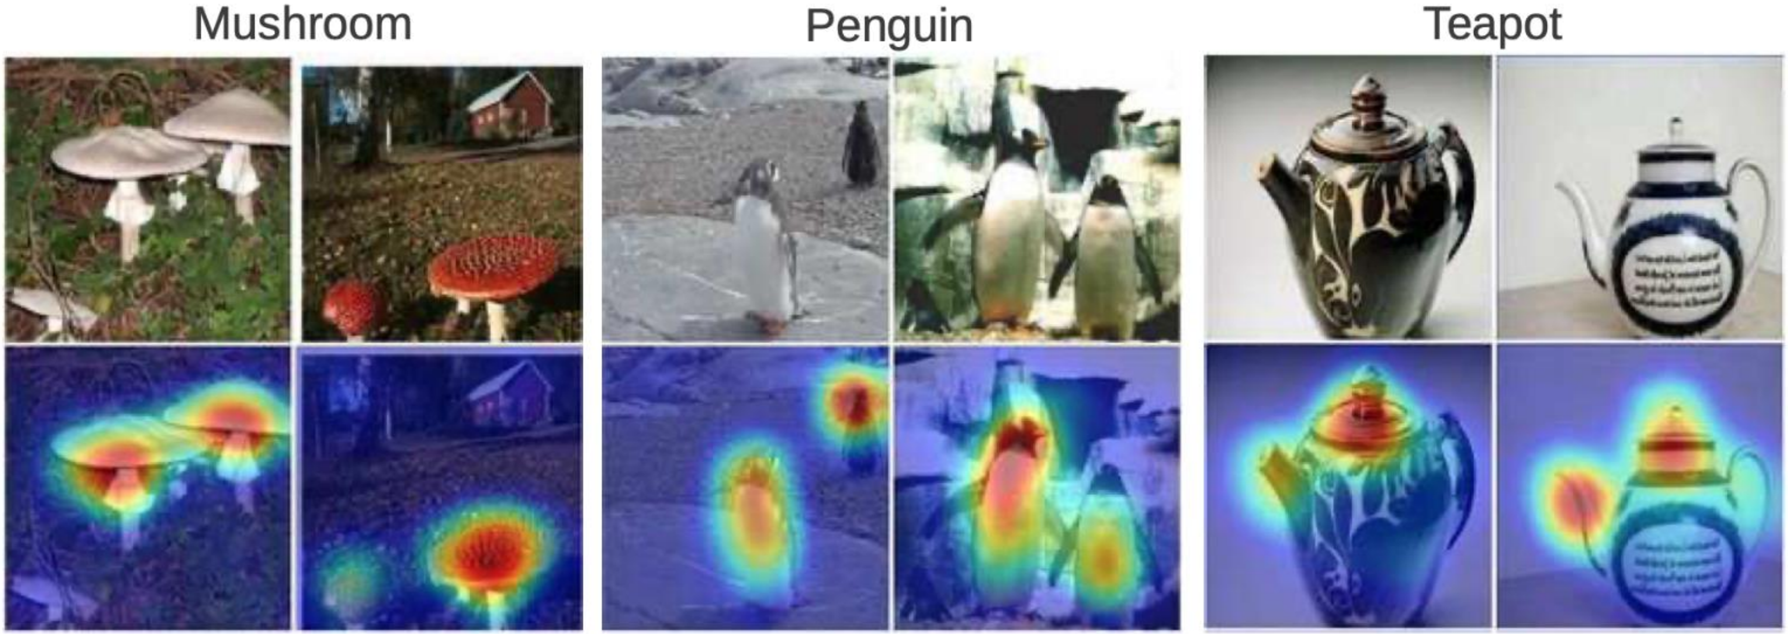

- Zhou et al., Learning deep features for discriminative localization. CVPR 2016

From global cues to pixel labels

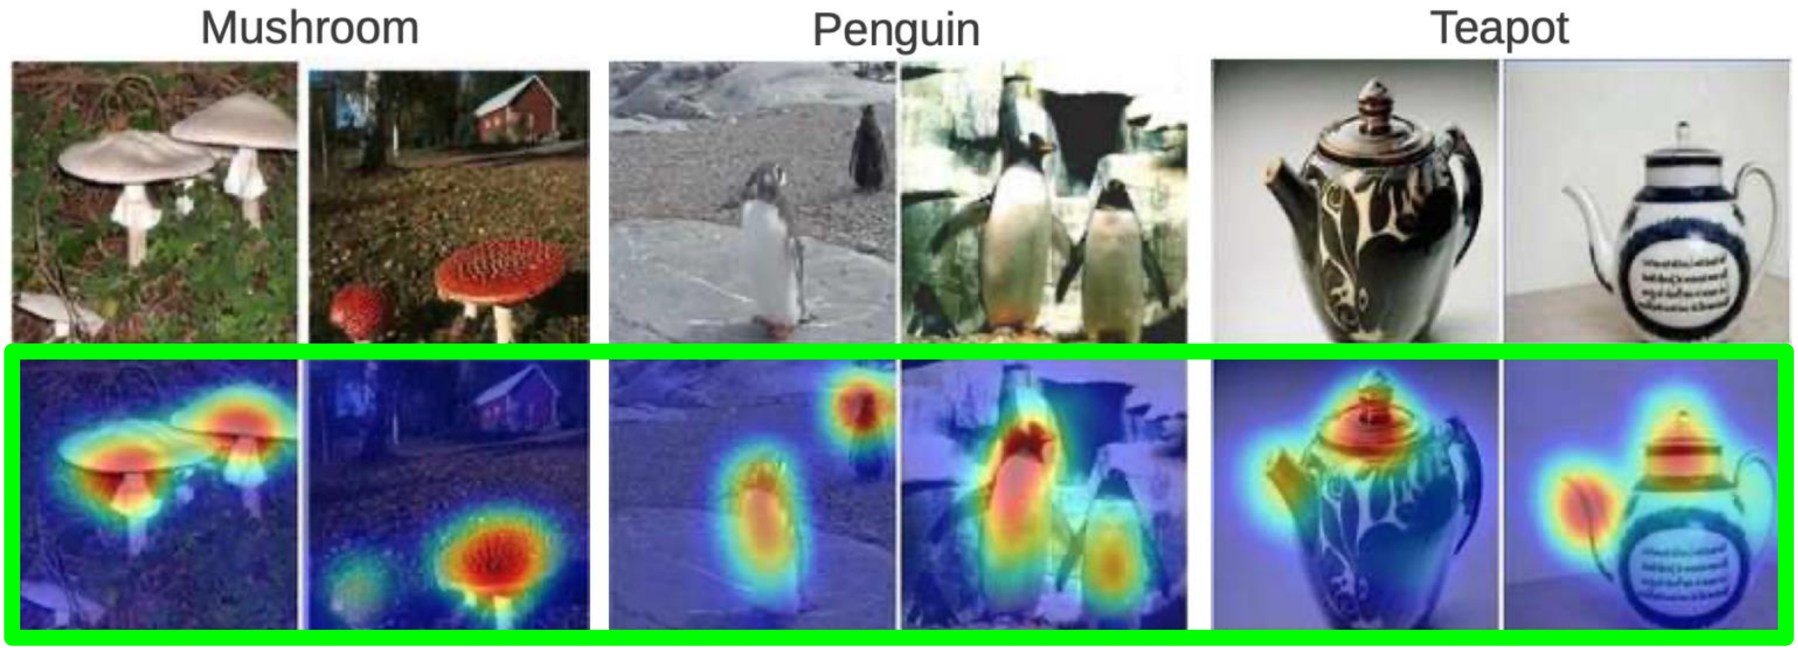

These activations maps can be used as **pseudo-masks**

$$-\sum_{i \in \Omega_i} \hat{y}_i \log(s_i)$$

- Zhou et al., Learning deep features for discriminative localization. CVPR 2016

From global cues to pixel labels

Problem: they focus only on highly discriminative regions

From global cues to pixel labels

How to improve CAMs?

Incorporate saliency maps

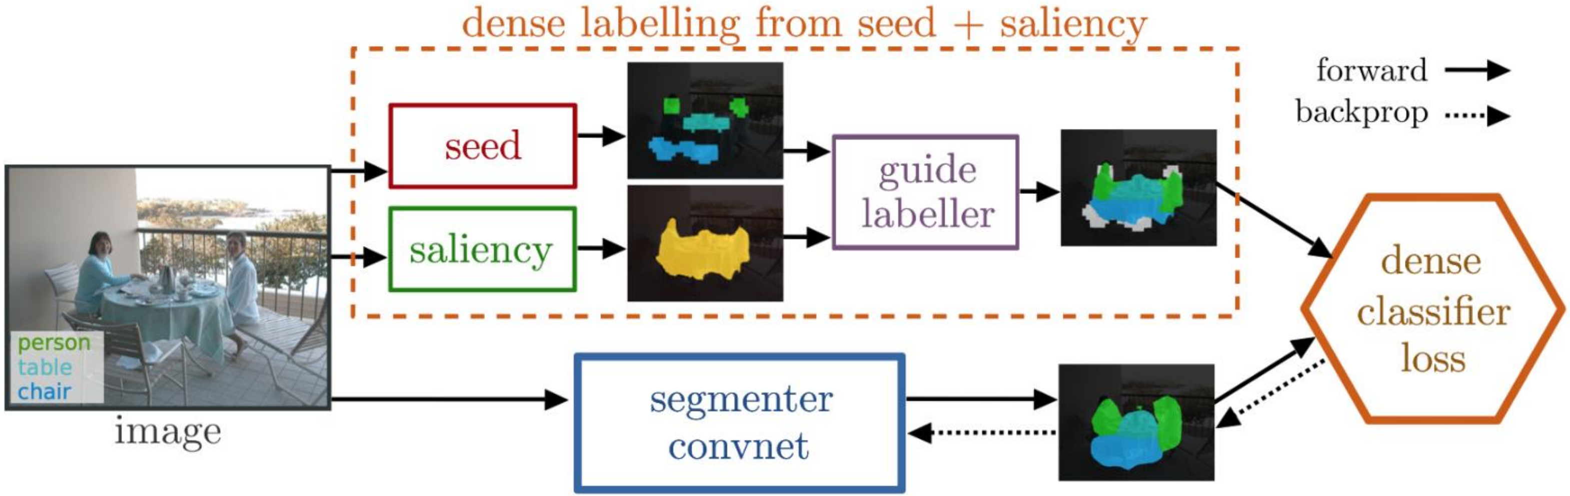

- Oh et al. Exploiting Saliency for Object Segmentation from Image Level Labels. CVPR 2017
- Fan et al. Learning Integral Objects With Intra-Class Discriminator for Weakly-Supervised Semantic Segmentation. CVPR 2020

From global cues to pixel labels

How to improve CAMs?

Region mining

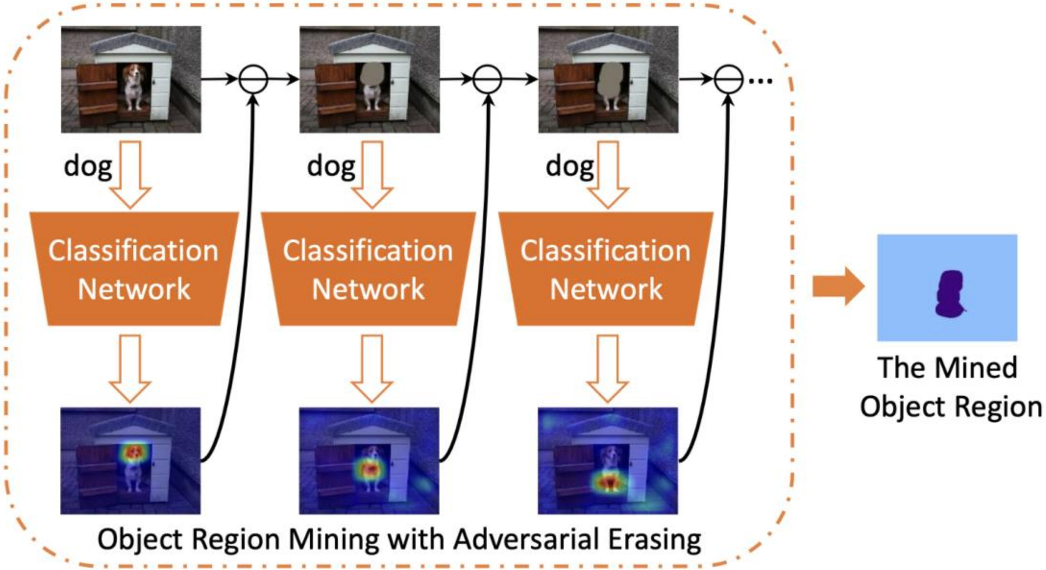

- Wei et al. Object Region Mining with Adversarial Erasing: A Simple Classification to Semantic Segmentation Approach. CVPR 2017
- Wang et al. Weakly-Supervised Semantic Segmentation by Iteratively Mining Common Object Features. CVPR 2018

From global cues to pixel labels

How to improve CAMs?

Region mining

1 - Obtain CAM
2 - Mask image with that region
3 - Repeat

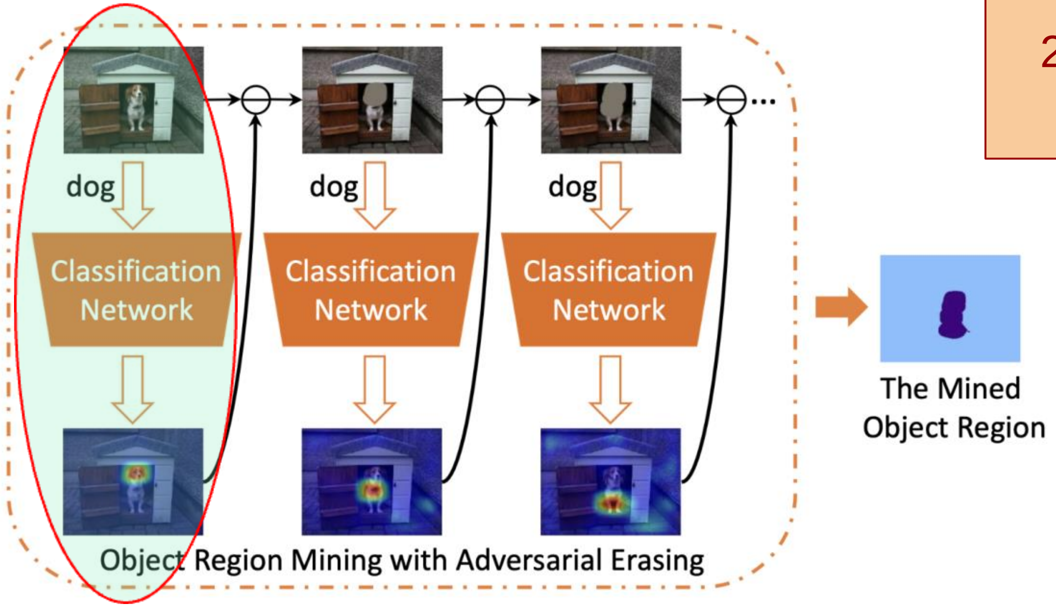

- Wei et al. Object Region Mining with Adversarial Erasing: A Simple Classification to Semantic Segmentation Approach. CVPR 2017
- Wang et al. Weakly-Supervised Semantic Segmentation by Iteratively Mining Common Object Features. CVPR 2018

From global cues to pixel labels

How to improve CAMs?

Region mining

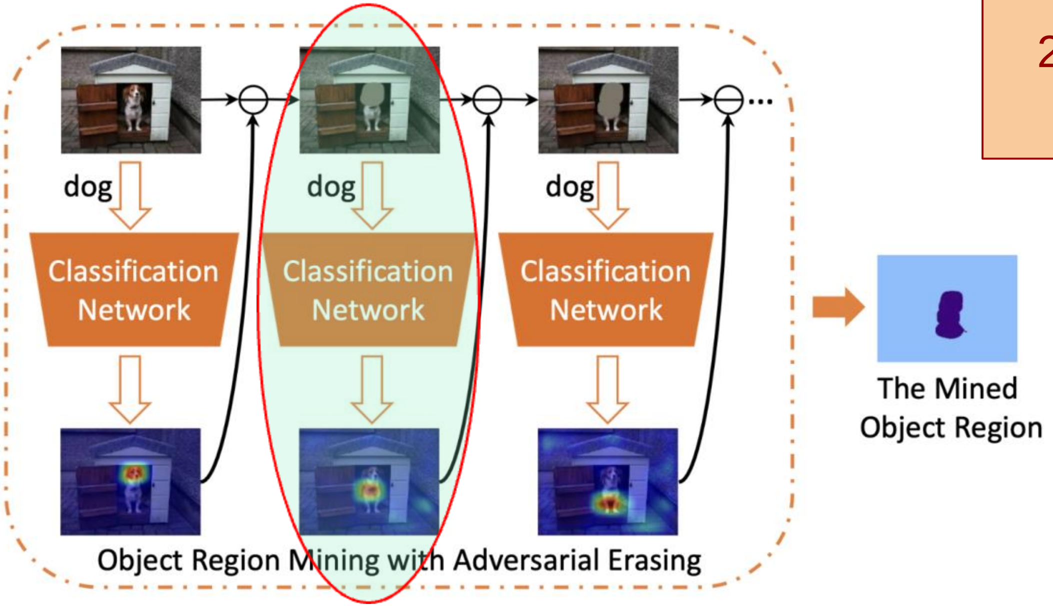

1 - Obtain CAM
2 - Mask image with that region
3 - Repeat

- Wei et al. Object Region Mining with Adversarial Erasing: A Simple Classification to Semantic Segmentation Approach. CVPR 2017
- Wang et al. Weakly-Supervised Semantic Segmentation by Iteratively Mining Common Object Features. CVPR 2018

From global cues to pixel labels

How to improve CAMs?

Region mining

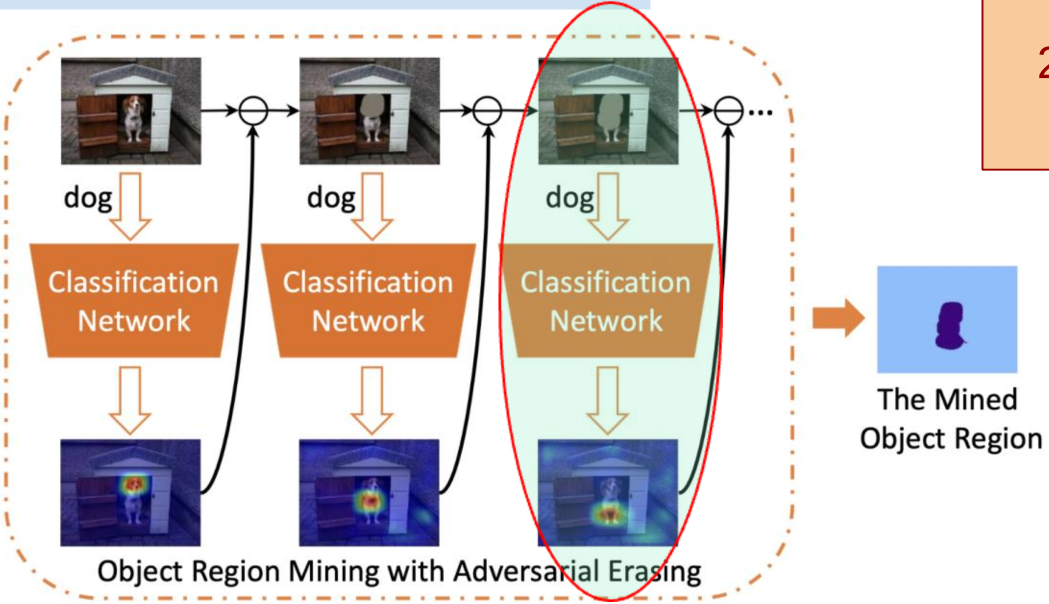

1 - Obtain CAM
2 - Mask image with that region
3 - Repeat

- Wei et al. Object Region Mining with Adversarial Erasing: A Simple Classification to Semantic Segmentation Approach. CVPR 2017
- Wang et al. Weakly-Supervised Semantic Segmentation by Iteratively Mining Common Object Features. CVPR 2018

From global cues to pixel labels

How to improve CAMs?

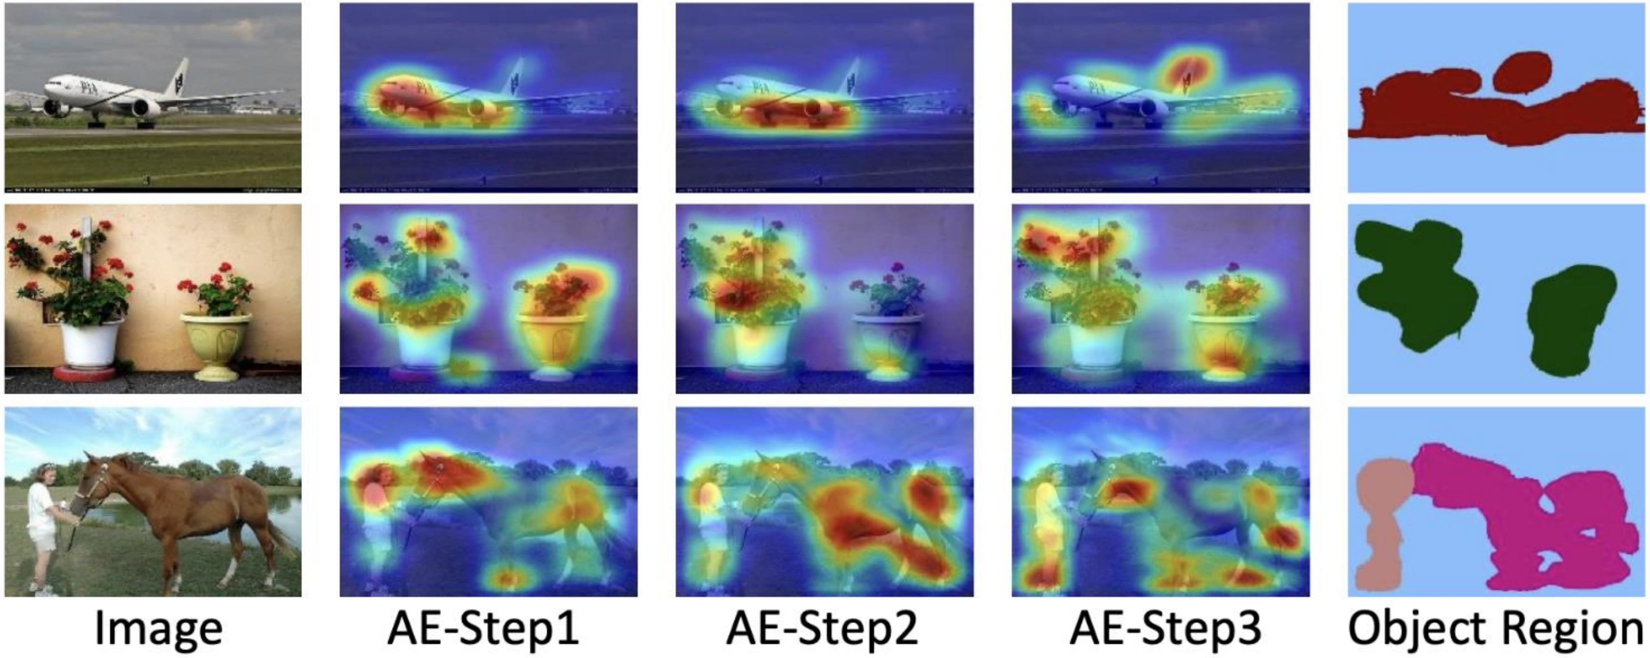

- Wei et al. Object Region Mining with Adversarial Erasing: A Simple Classification to Semantic Segmentation Approach. CVPR 2017
- Wang et al. Weakly-Supervised Semantic Segmentation by Iteratively Mining Common Object Features. CVPR 2018

From global cues to pixel labels

How to improve CAMs?

Equivariant constraints

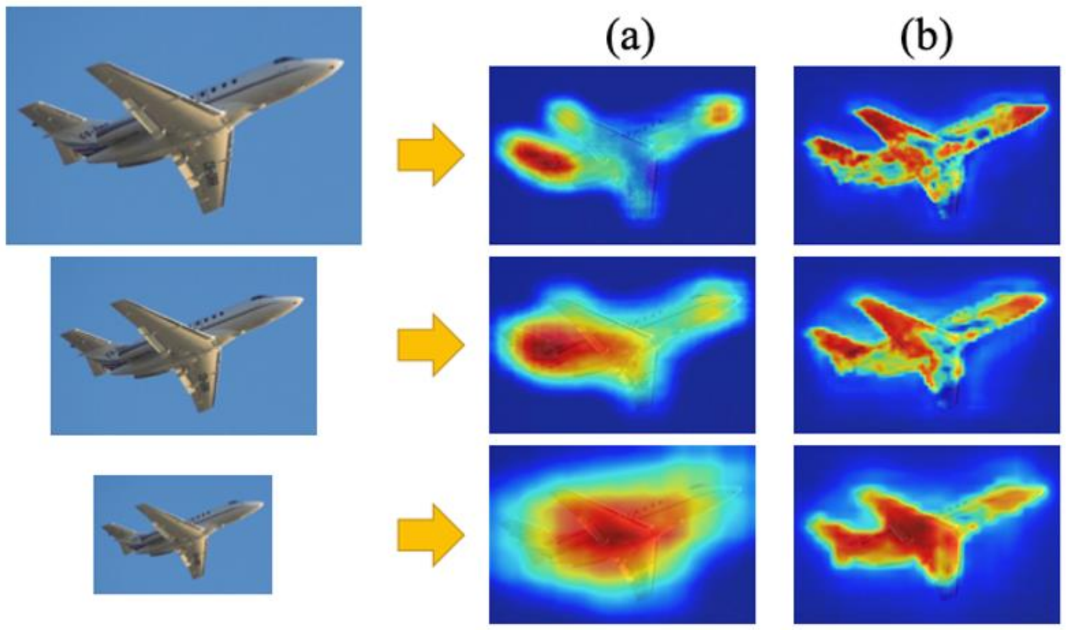

Wang et al. Self-supervised Equivariant Attention Mechanism for Weakly Supervised Semantic segmentation. CVPR’20

From global cues to pixel labels

Integrating language-vision models

How to improve CAMs?

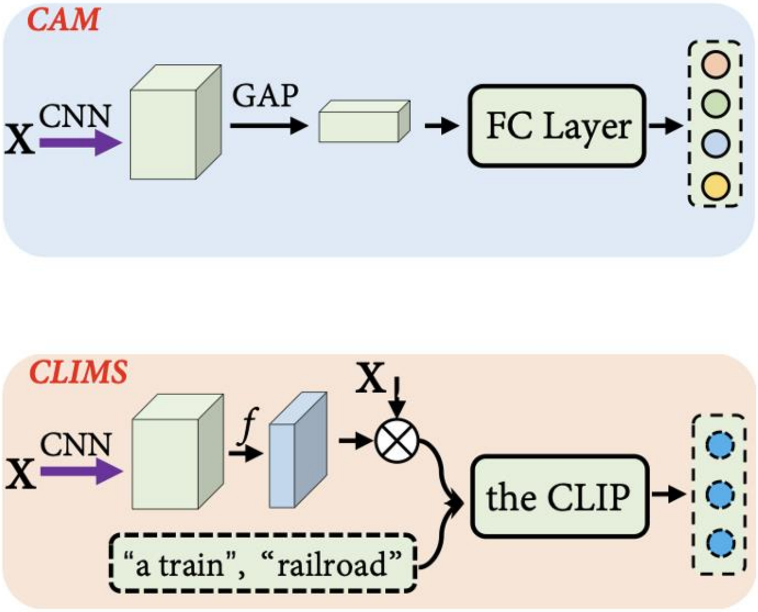

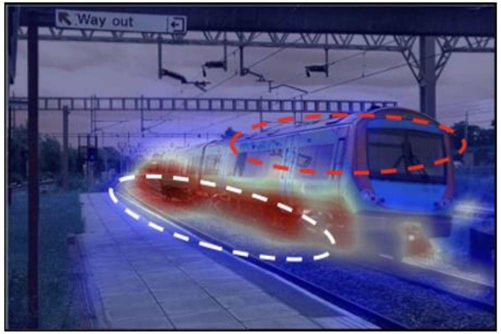

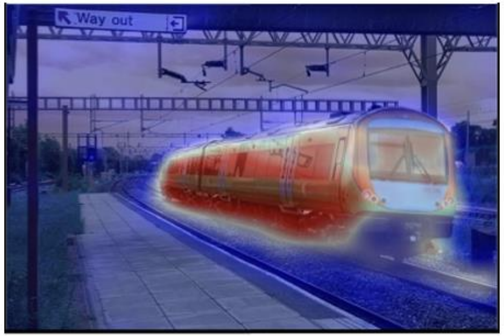

Xie et al, CLIMS: Cross Language Matching for weakly supervised semantic segmentation

From global cues to pixel labels

Integrating language-vision models

How to improve CAMs?

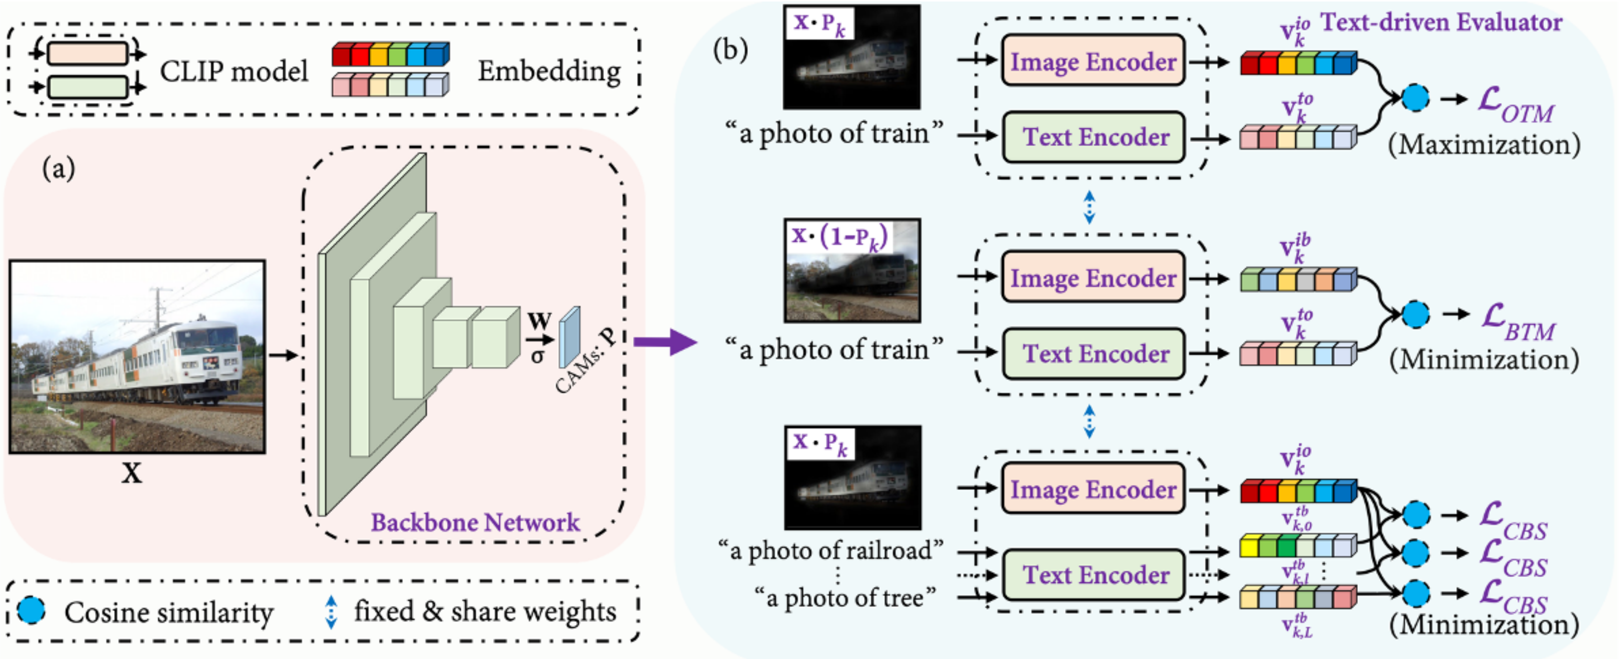

Xie et al, CLIMS: Cross Language Matching for weakly supervised semantic segmentation

CAMs in the medical domain

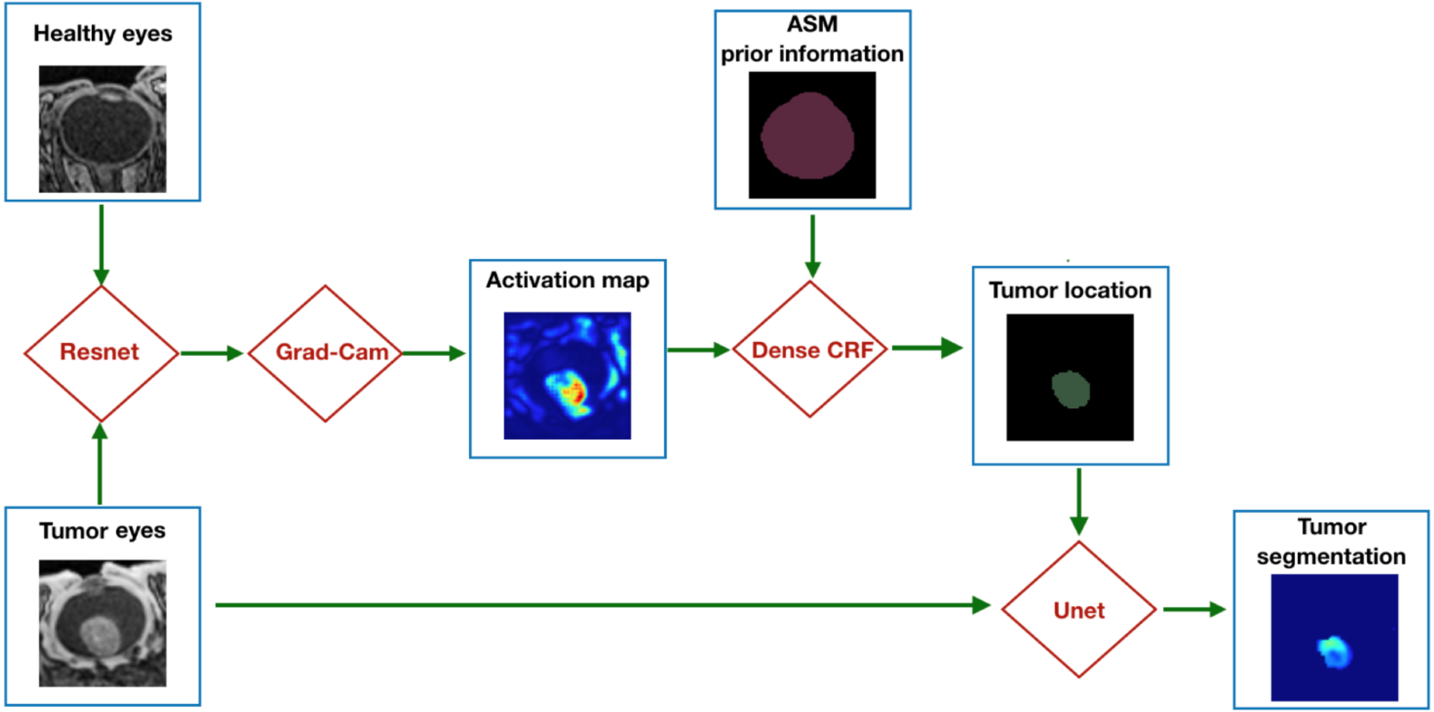

Nguyen et al. A novel segmentation framework for uveal melanoma based on magnetic resonance imaging and class activation maps. MIDL 2019.

CAMs in the medical domain

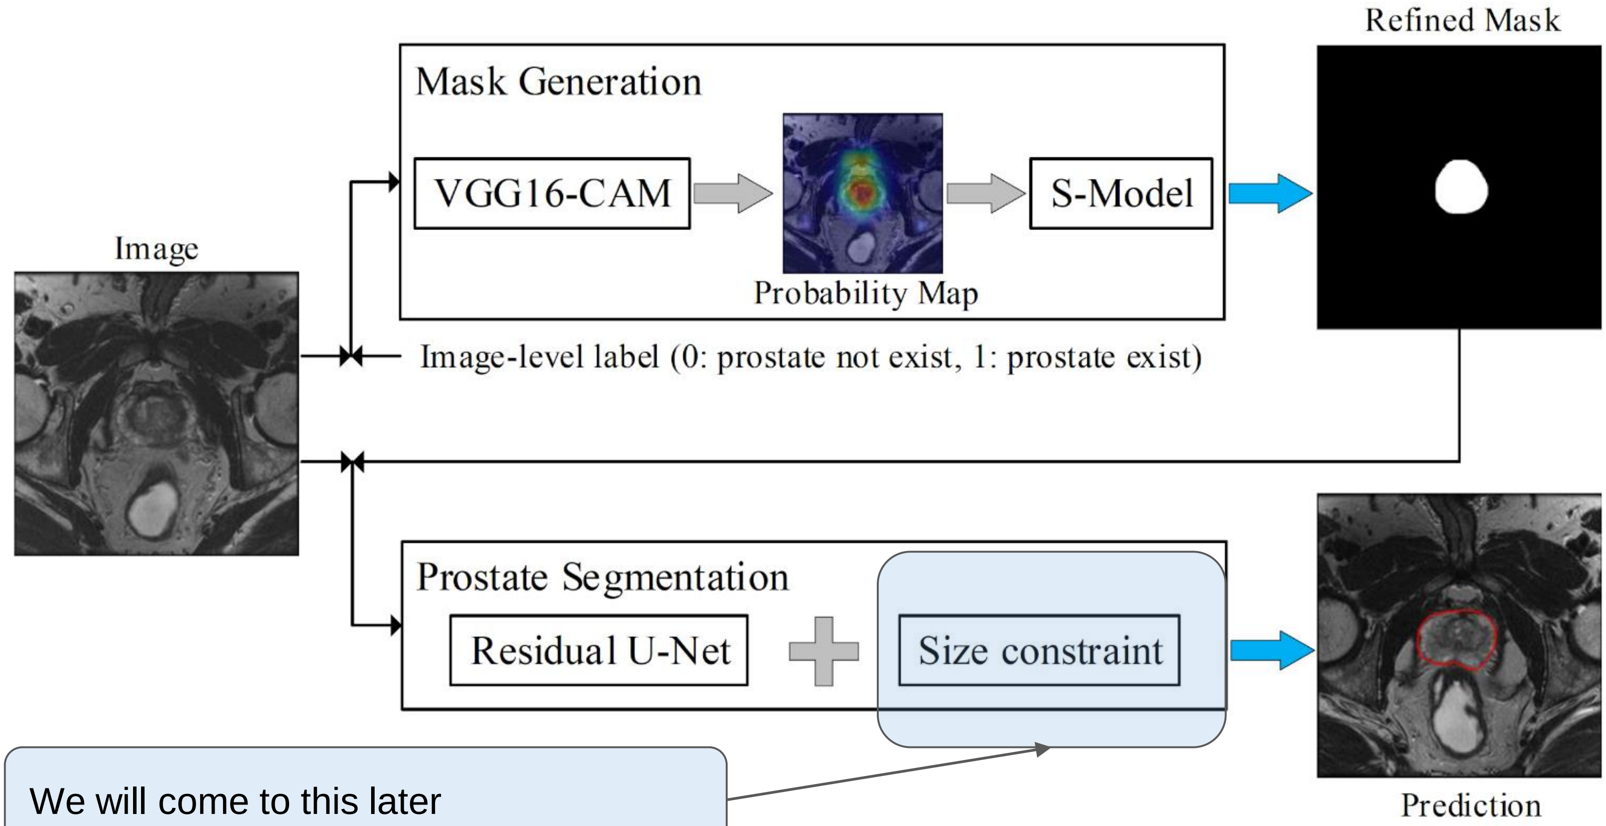

We will come to this later

Chen et al. Exploiting confident information for weakly supervised prostate segmentation based on image-level labels. SPIE Medical Imaging 2020

CAMs in the medical domain

Equivariant constraints

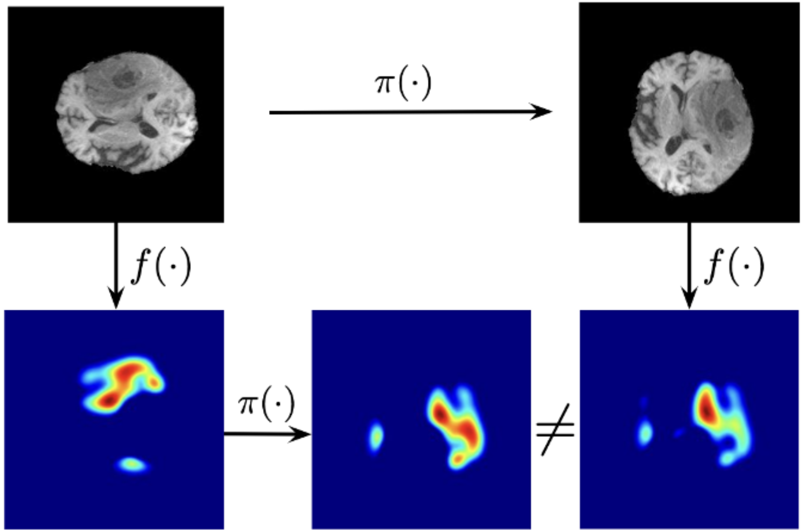

CAMs not equivariant to spatial transformations

Patel and Dolz, Weakly supervised segmentation with equivariant constraints. MedIA'22

CAMs in the medical domain

Equivariant constraints

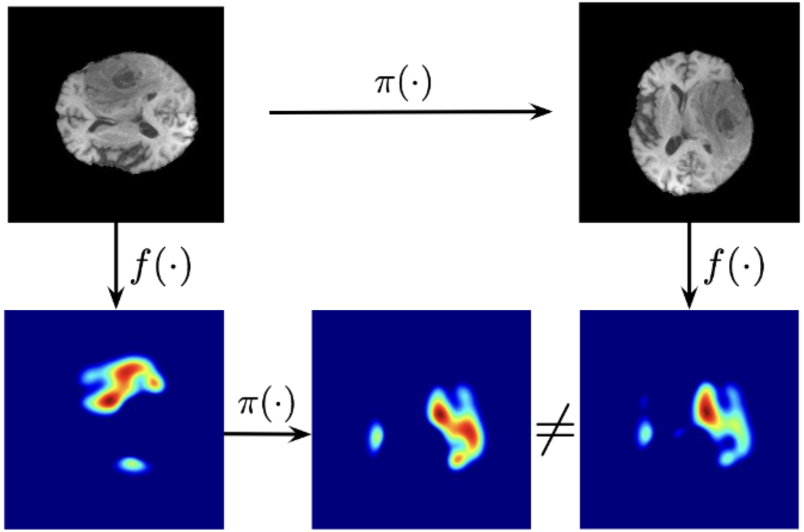

CAMs not equivariant to spatial transformations

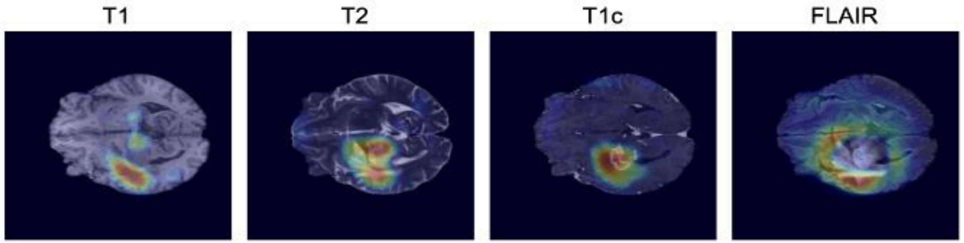

CAMs not consistent across modalities

Patel and Dolz, Weakly supervised segmentation with equivariant constraints. MedIA'22

CAMs in the medical domain

Equivariant constraints

Same-modality equivariant constraints

$$\mathcal{L}_{reg}(\tau(f(M_1)), f(\tau(M_1)))$$

transformations

Cross-modality equivariant constraints

$$\mathcal{L}_{reg}(\tau(f(M_1)), f(\tau(M_2)))$$

transformations

Modality 1

Modality 2

Patel and Dolz, Weakly supervised segmentation with equivariant constraints. MedIA'22

CAMs in the medical domain

Equivariant constraints

Same-modality equivariant constraints

$$\mathcal{L}_{reg}(\tau(f(M_1)), f(\tau(M_1)))$$

Cross-modality equivariant constraints

$$\mathcal{L}_{reg}(\tau(f(M_1)), f(\tau(M_2)))$$

Classification loss

$$\mathcal{L}_{class} + \mathcal{L}_{reg}(\tau(f(M_1)), f(\tau(M_1))) + \mathcal{L}_{reg}(\tau(f(M_1)), f(\tau(M_2)))$$

Intra-modal regularization

Inter-modal regularization

Patel and Dolz, Weakly supervised segmentation with equivariant constraints. MedIA'22

CAMs in the medical domain

Equivariant constraints

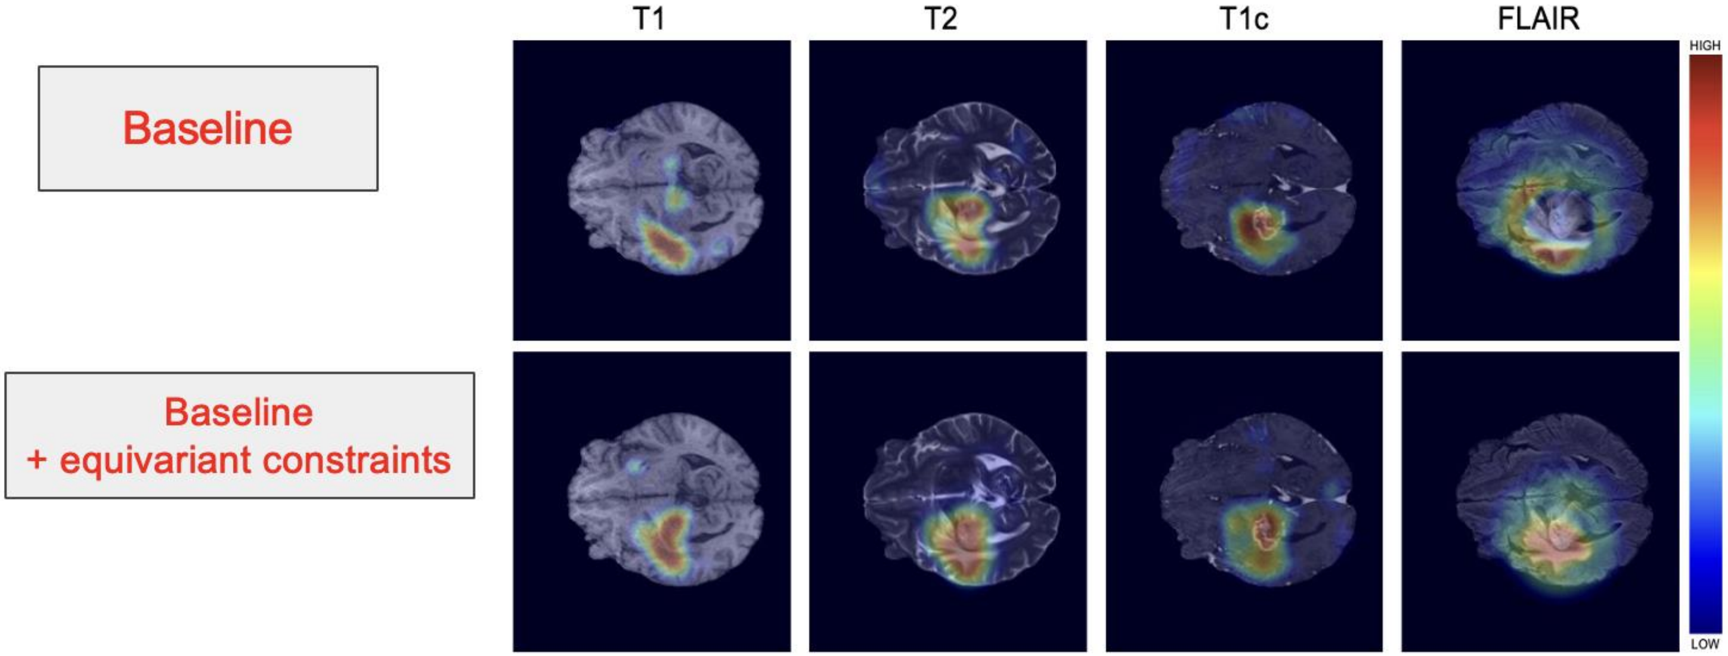

Patel and Dolz, Weakly supervised segmentation with equivariant constraints. MedIA'22

Constrained optimization (CNN)

Optimize (A) such that (B)

Task

Set of constraints

Constrained optimization (CNN)

Optimize (A)

Task

such that (B)

Set of constraints

How we can go from point A to B?

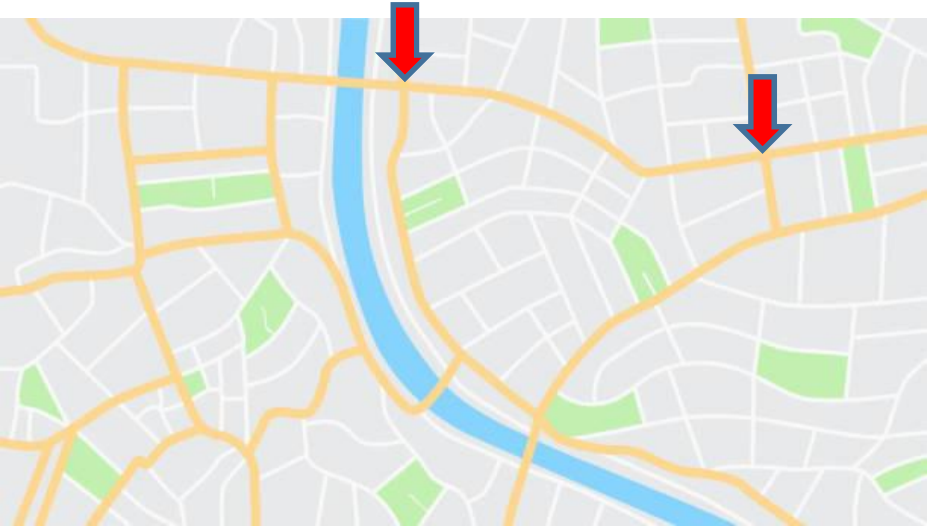

Constrained optimization (CNN)

Optimize (A) such that (B)

Task

Set of constraints

How we can go from point A to B?

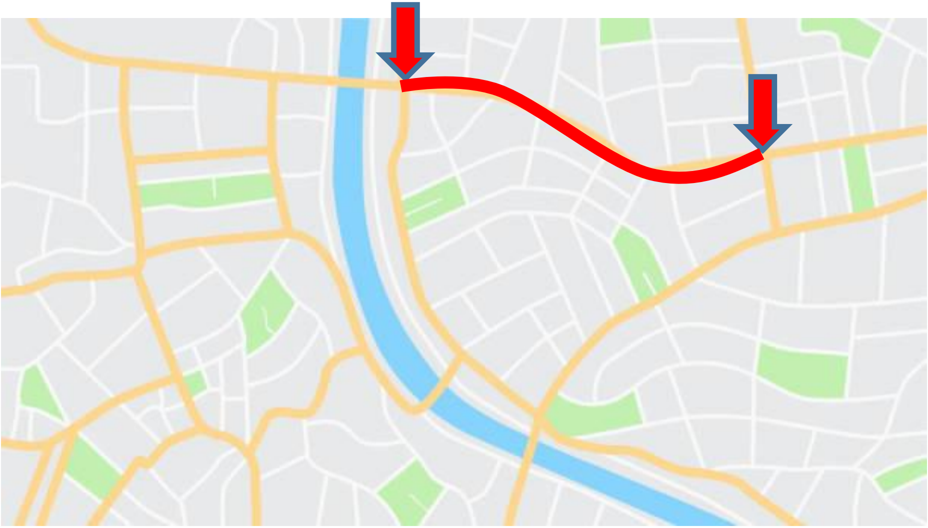

Constrained optimization (CNN)

Optimize (A)

Task

such that (B)

Set of constraints

How we can go from point A to B?

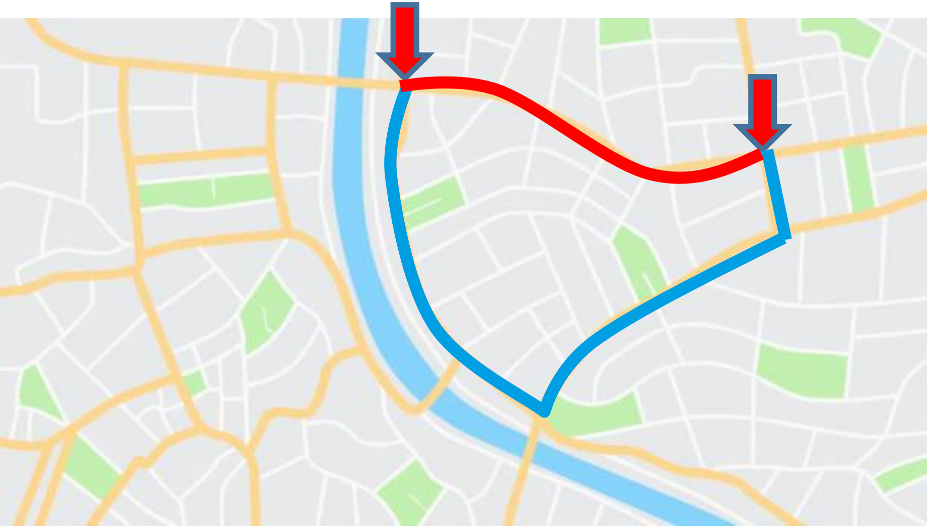

Constrained optimization (CNN)

Optimize (A) such that (B)

Task

Set of constraints

How we can go from point A to B?

Which is the best route?

Constrained optimization (CNN)

Optimize (A)

Task

such that (B)

Set of constraints

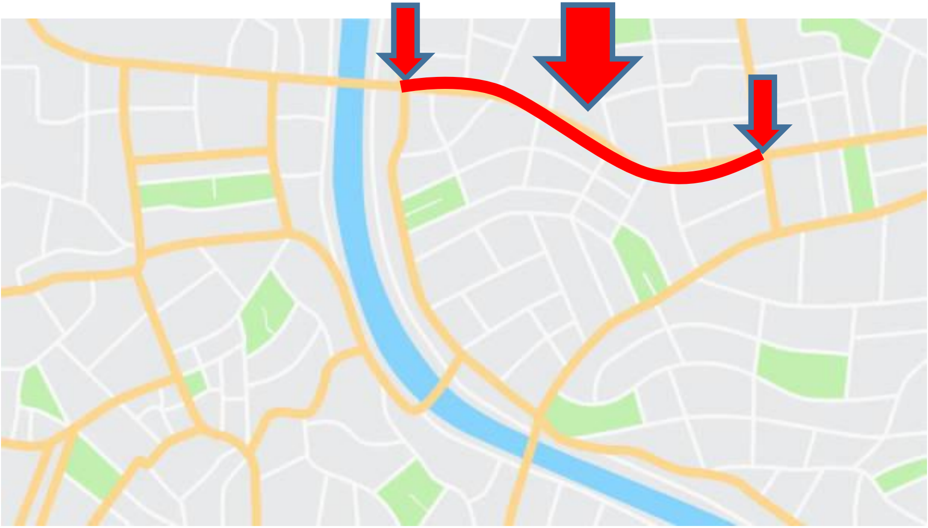

How we can go from point A to B?

Which is the best route?

Constraint: shortest

Constrained optimization (CNN)

Optimize (A)

Task

such that (B)

Set of constraints

How we can go from point A to B?

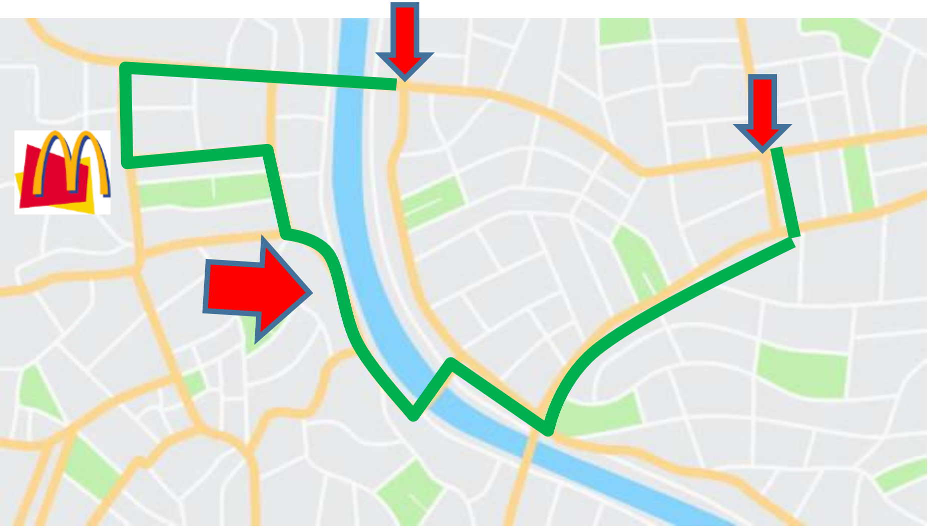

Which is the best route?

Constraint: shortest but going through a McDonnalds

Constrained optimization (CNN)

Optimize (A) such that (B)

$$\min_{\theta} \mathcal{H}(S, Y) \quad s.t. \quad \sum_{n=0}^{N} s_n = A$$

Constrained problem

Constrained optimization (CNN)

Optimize (A) such that (B)

$$\min_{\theta} \mathcal{H}(S, Y) \quad s.t. \quad \sum_{n=0}^{N} s_n = A$$

Constrained problem

$$\min_{\theta} \mathcal{H}(S, Y) + \lambda(\sum_{n=0}^{N} s_n - A)$$

Unconstrained problem
(penalties to the rescue!)

Equality constraints

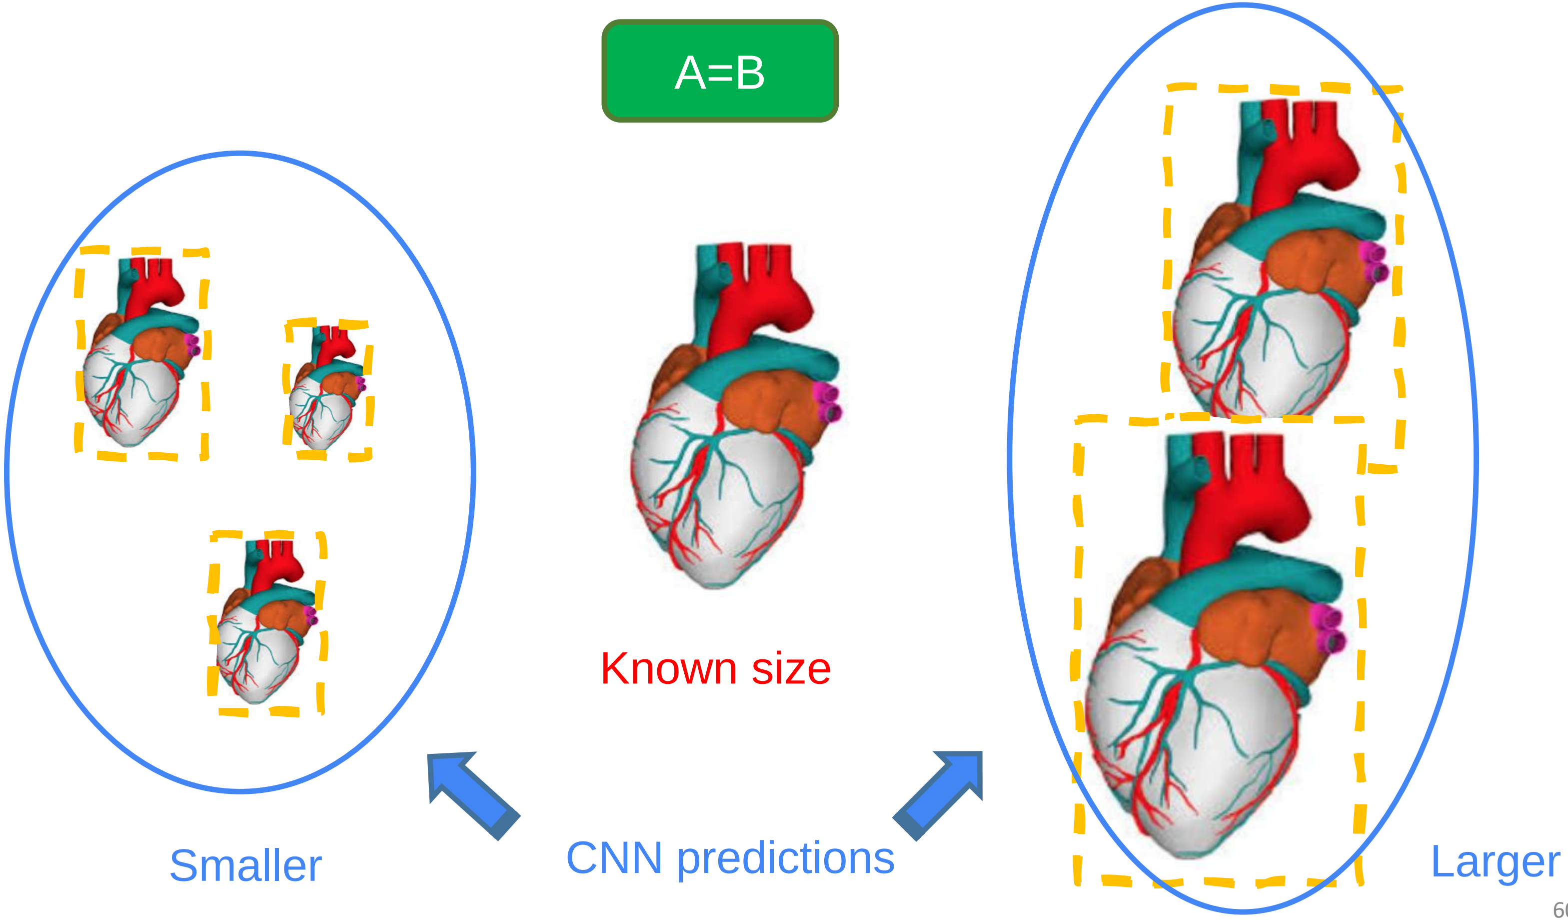

Equality constraints

A=B

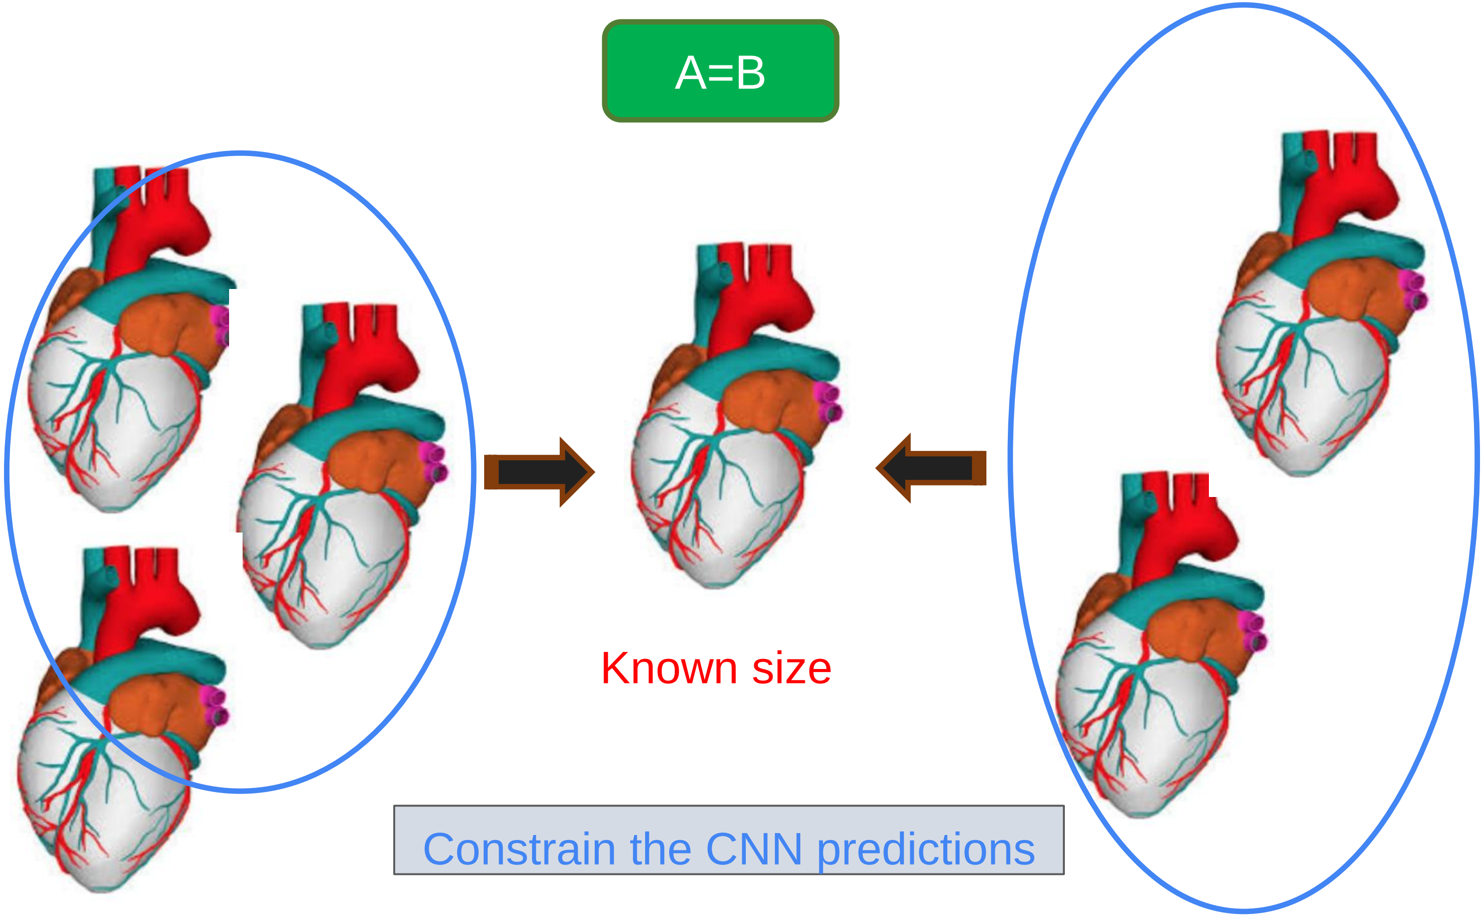

Known size

Constrain the CNN predictions

Equality constraints

A=B

General definition

$$\min_{\theta} \mathcal{H}(S) \quad s.t. \quad g(\mathbf{s}) = C$$

Constraint

Equality constraints

A=B

General definition

$$\min_{\theta} \mathcal{H}(S) \quad s.t. \quad g(\mathbf{s}) = C$$

Constraint

$$\min_{\theta} \mathcal{H}(S) + \lambda(g(\mathbf{s}) - C)$$

Penalty

Equality constraints

A=B

General definition

$$\min_{\theta} \mathcal{H}(S) \quad s.t. \quad g(\mathbf{s}) = C$$

Constraint

$$\min_{\theta} \mathcal{H}(S) + \lambda(g(\mathbf{s}) - C)$$

Penalty

This can be modeled with linear/quadratic penalties, KL divergence, etc

Equality constraints

L2 Penalty

Input
(Histology image)

Any CNN architecture

Output
(Pixel-wise prediction)

Jia et al. Constrained deep weak supervision for histopathology image segmentation. IEEE TMI 2017

Equality constraints

L2 Penalty

Input
(Histology image)

Any CNN architecture

Output
(Pixel-wise prediction)

Additional term

$$\boxed{l_{ac} = \boldsymbol{I}(Y_i = 1)\sum_i (v_i - a_i)^2}$$

Predicted relative size (%)

$$v_i = \frac{1}{N}\sum_{p\in\Omega} s_\theta^{p,1}$$

Predicted size given by experts

Jia et al. Constrained deep weak supervision for histopathology image segmentation. IEEE TMI 2017

Equality constraints

Kullback-Leibler (KL) Divergence

Partial annotations

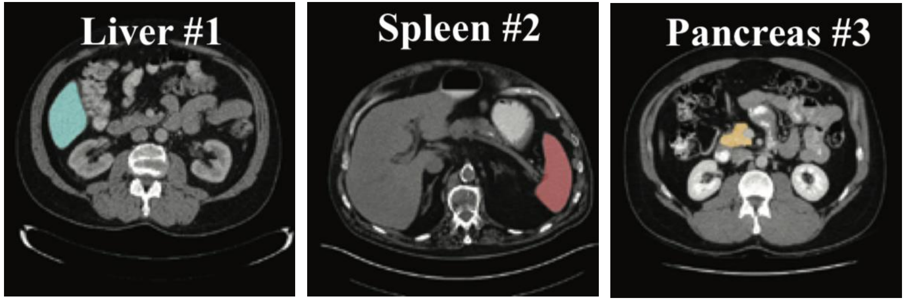

[Zhou et al., Prior-aware Neural Network for Partially-Supervised Multi-Organ Segmentation, ICCV'19]

Equality constraints

Kullback-Leibler (KL) Divergence

Partial annotations

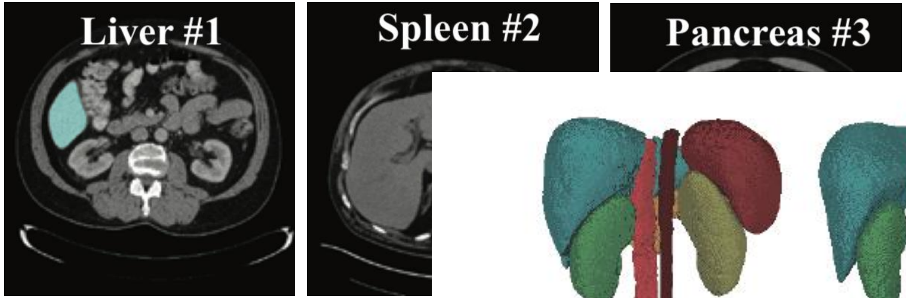

Prior on the proportion

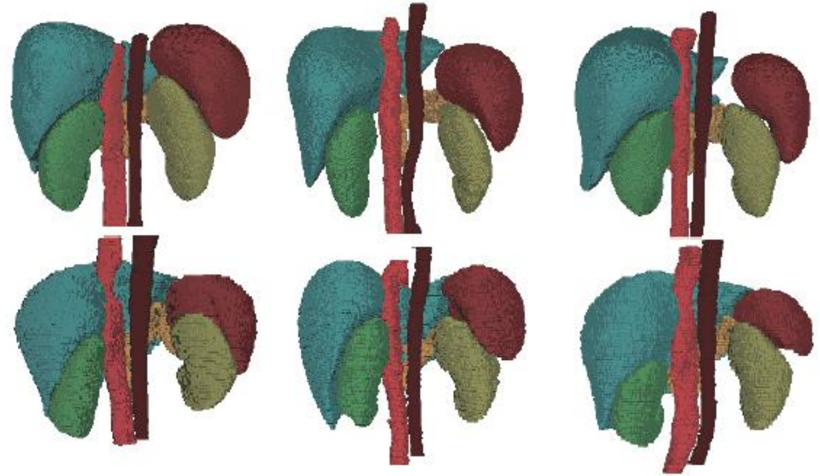

Figure 1. 3D Visualization of several abdominal organs (liver, spleen, left kidney, right kidney, aorta, inferior vena cava) to show the similarity of patient-wise abdominal organ size distributions.

[Zhou et al., Prior-aware Neural Network for Partially-Supervised Multi-Organ Segmentation, ICCV'19]

Equality constraints

Kullback-Leibler (KL) Divergence

Partial annotations

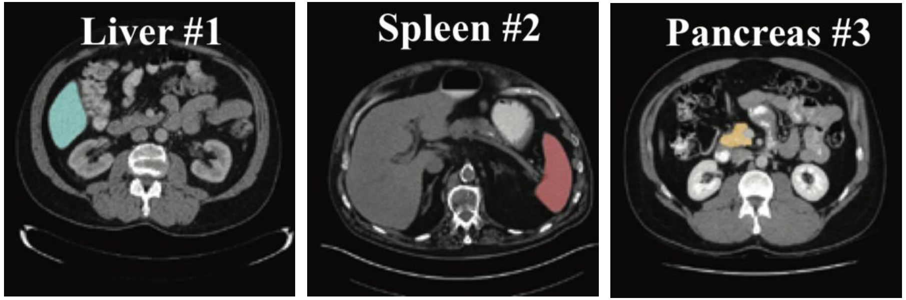

Main objective:

$$\underbrace{min \frac{1}{|\mathcal{L}|} \sum_{p \in \mathcal{L}} l(\mathbf{y}^p, \mathbf{s}_\theta^p) + \lambda_1 \frac{1}{|\mathcal{P}|} \sum_{q \in \mathcal{P}} l(\mathbf{y}^q, \mathbf{s}_\theta^q)} + \lambda_2 \mathcal{J}(\theta)$$

Prior-aware loss

Fully labeled images

Partially labeled images

[Zhou et al., Prior-aware Neural Network for Partially-Supervised Multi-Organ Segmentation, ICCV'19]

Equality constraints

Kullback-Leibler (KL) Divergence

Partial annotations

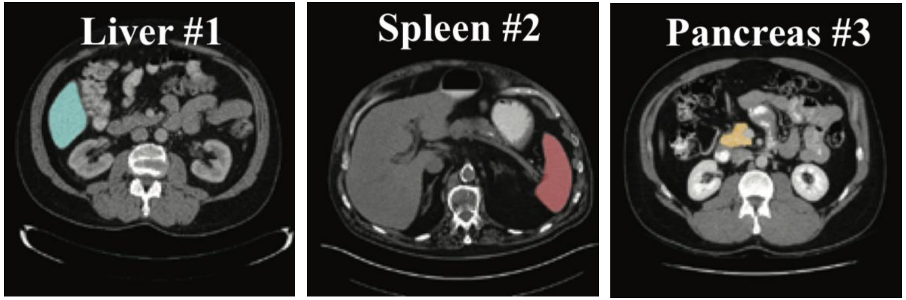

Prior-aware loss

Averaged predicted distribution

$$\hat{\mathbf{p}} = \frac{1}{N} \sum_{p \in \mathcal{P}} \mathbf{s}_\theta^p$$

$$[s_\theta^{p,0}, s_\theta^{p,1}, \ldots, s_\theta^{p,|K|}]$$

On partially labeled images

[Zhou et al., Prior-aware Neural Network for Partially-Supervised Multi-Organ Segmentation, ICCV'19]

Equality constraints

Kullback-Leibler (KL) Divergence

Partial annotations

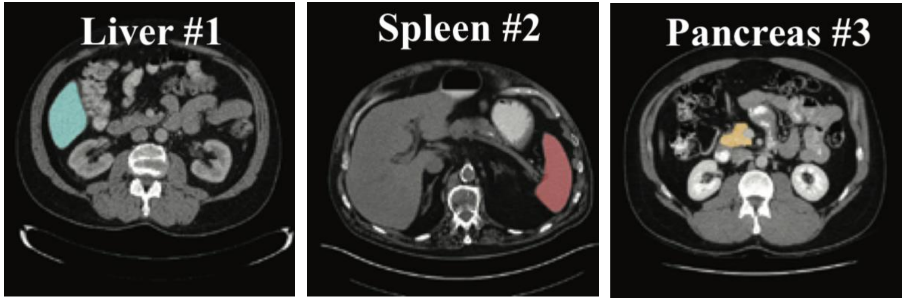

Prior-aware loss

Embed prior knowledge

$$KL(\mathbf{q}|\hat{\mathbf{p}})$$

Real label distribution Average predicted distribution

[Zhou et al., Prior-aware Neural Network for Partially-Supervised Multi-Organ Segmentation, ICCV'19]

Equality constraints

At pixel level

Imposing Consistency across image modalities

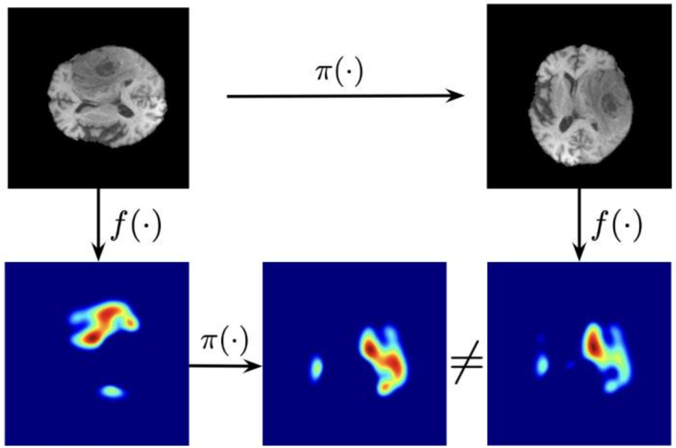

CAMs not equivariant to spatial transformations

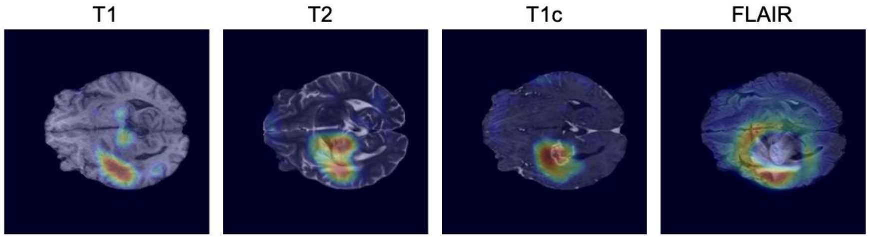

CAMs not consistent across modalities

Patel and Dolz, Weakly supervised segmentation with equivariant constraints. MedIA'22

Equality constraints

At pixel level

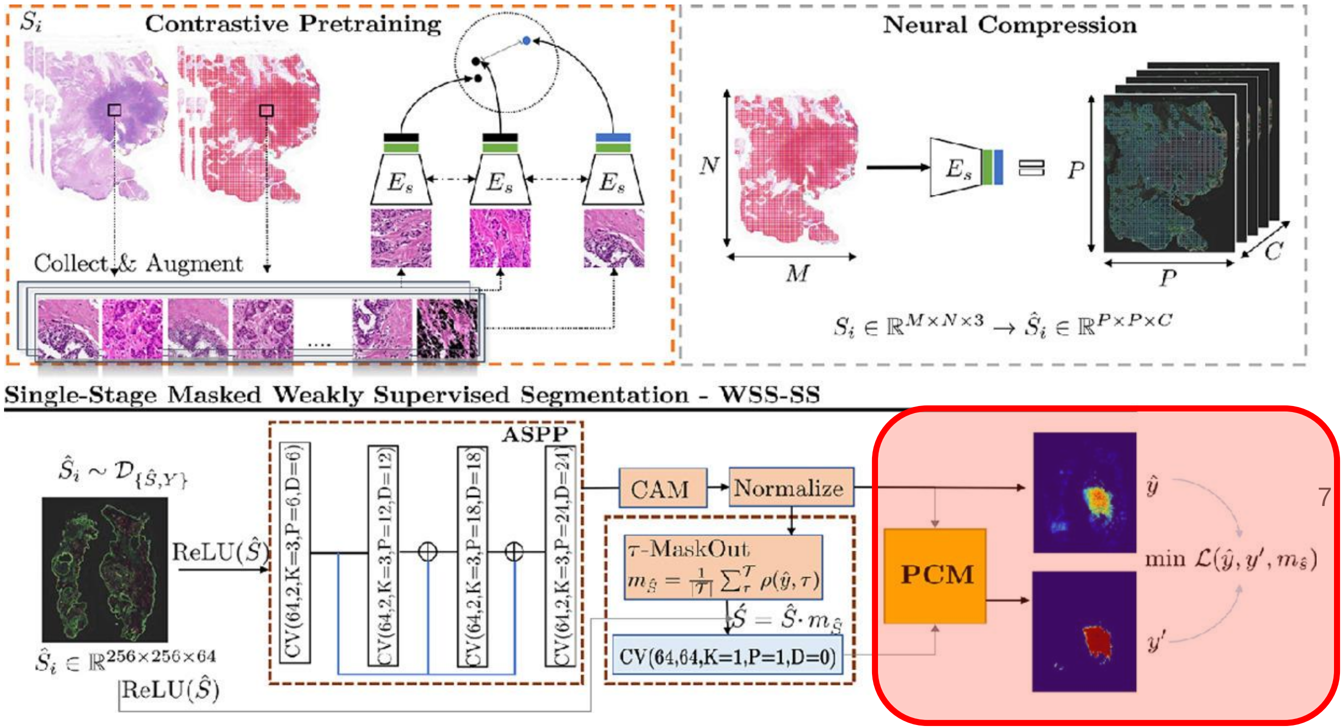

Equivariant regularization

Chikontwe et al. Weakly Supervised Segmentation on Neural Compressed Histopathology with Self-Equivariant Regularization. MedIA’22

Inequality constraints

A<B,
A>B,
B1 < A < B2

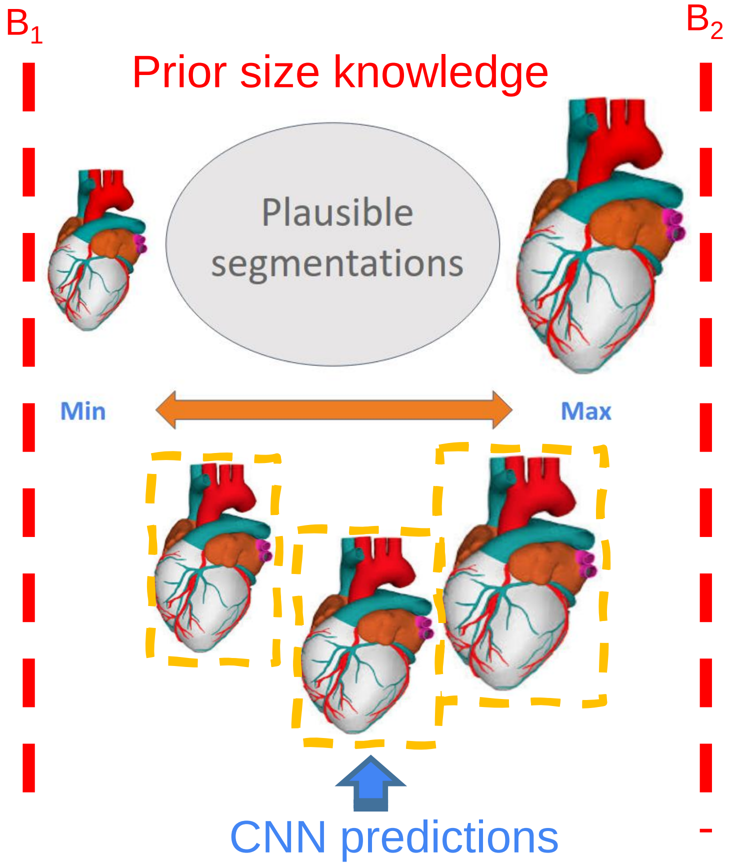

Inequality constraints

$A<B$,
$A>B$,
$B1 < A < B2$

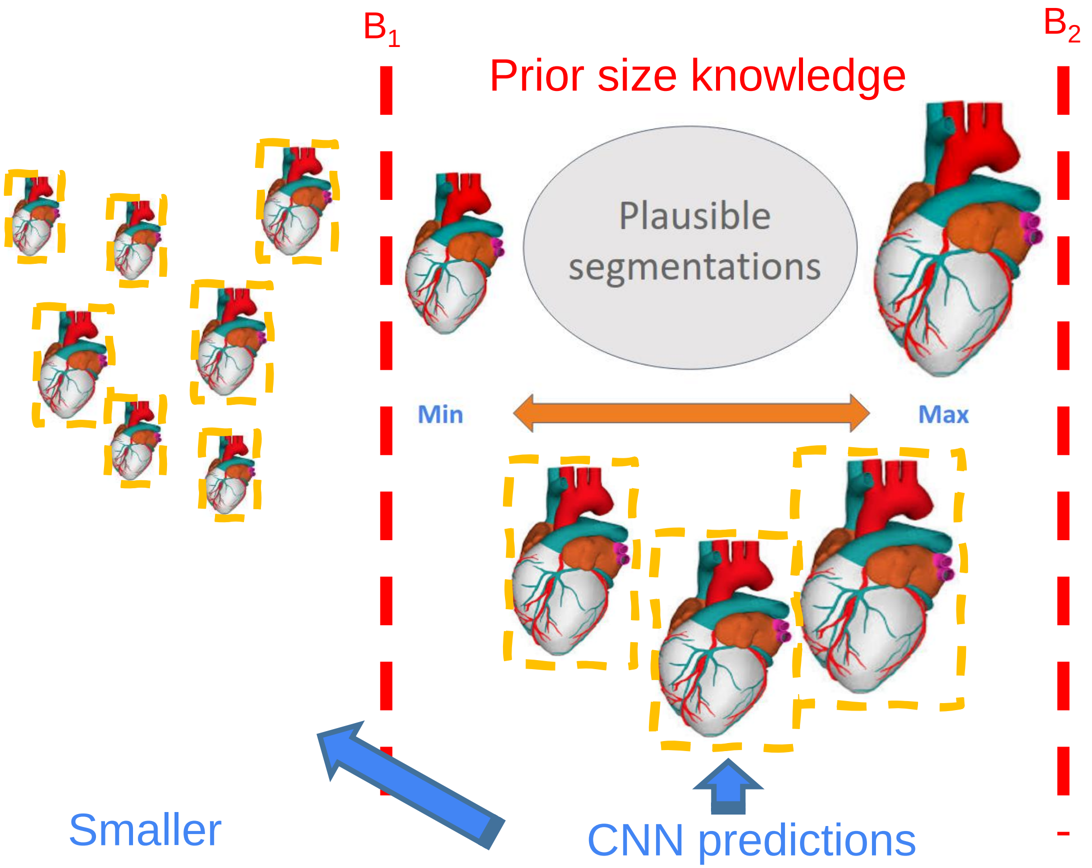

Inequality constraints

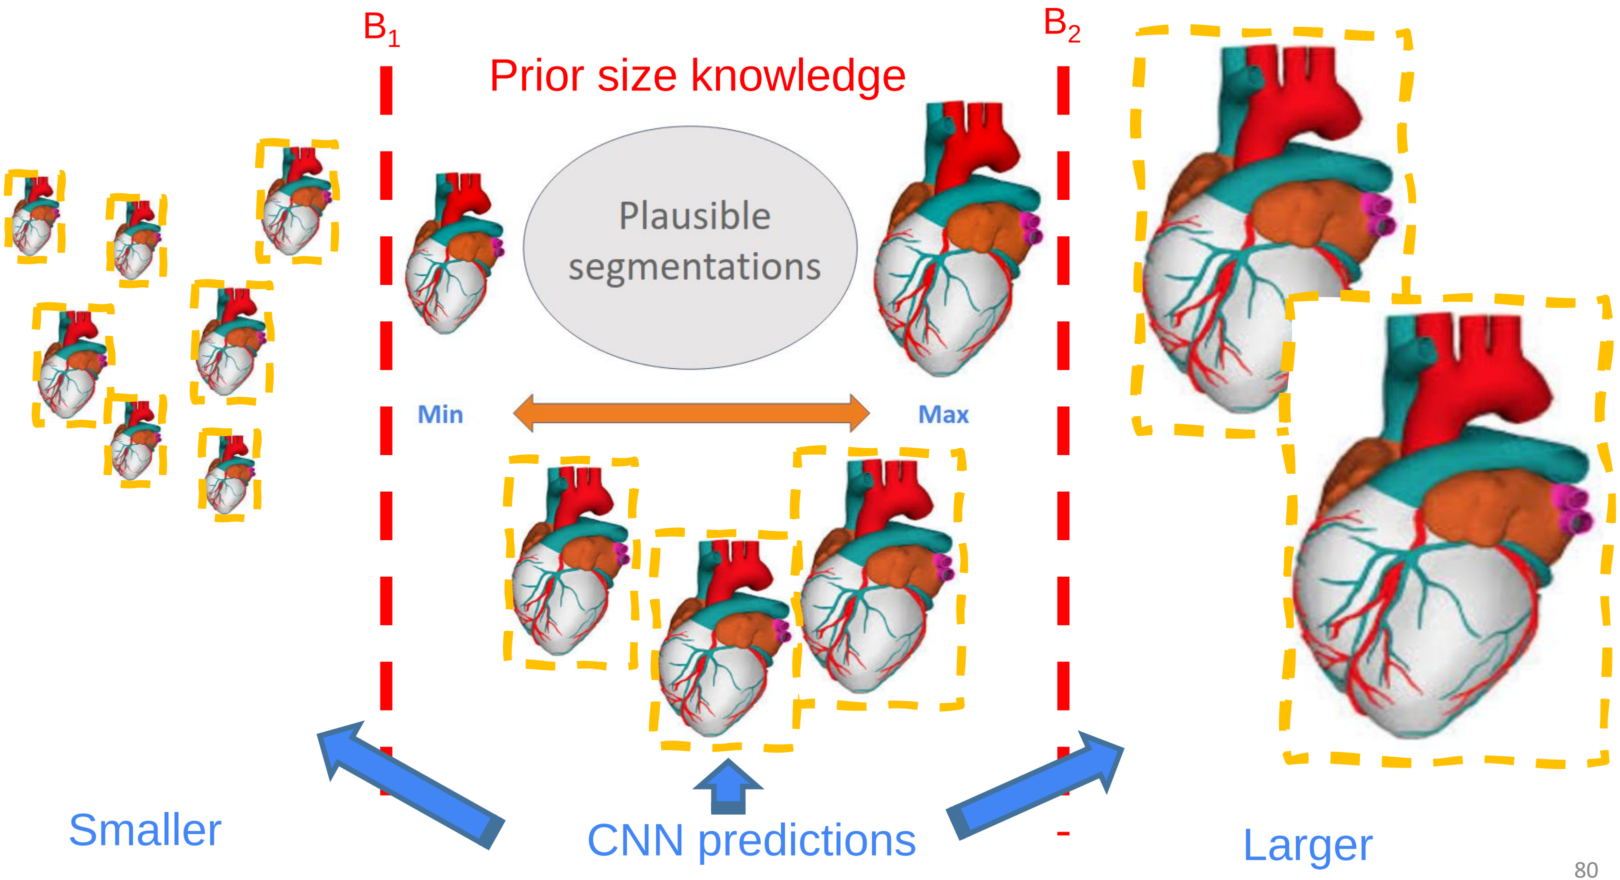

Inequality constraints

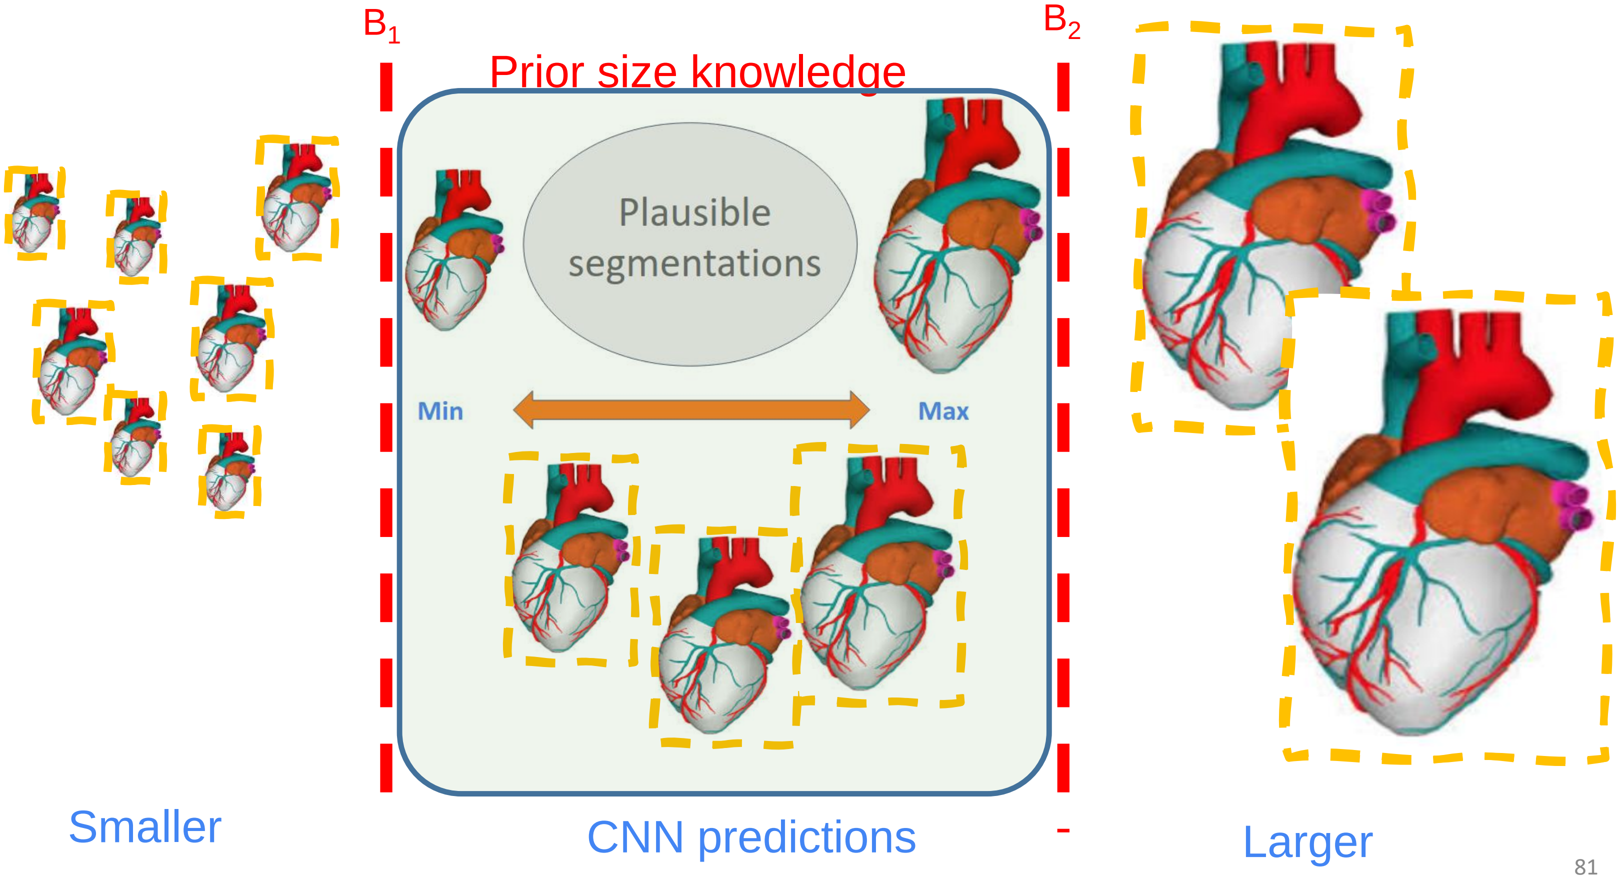

Inequality constraints

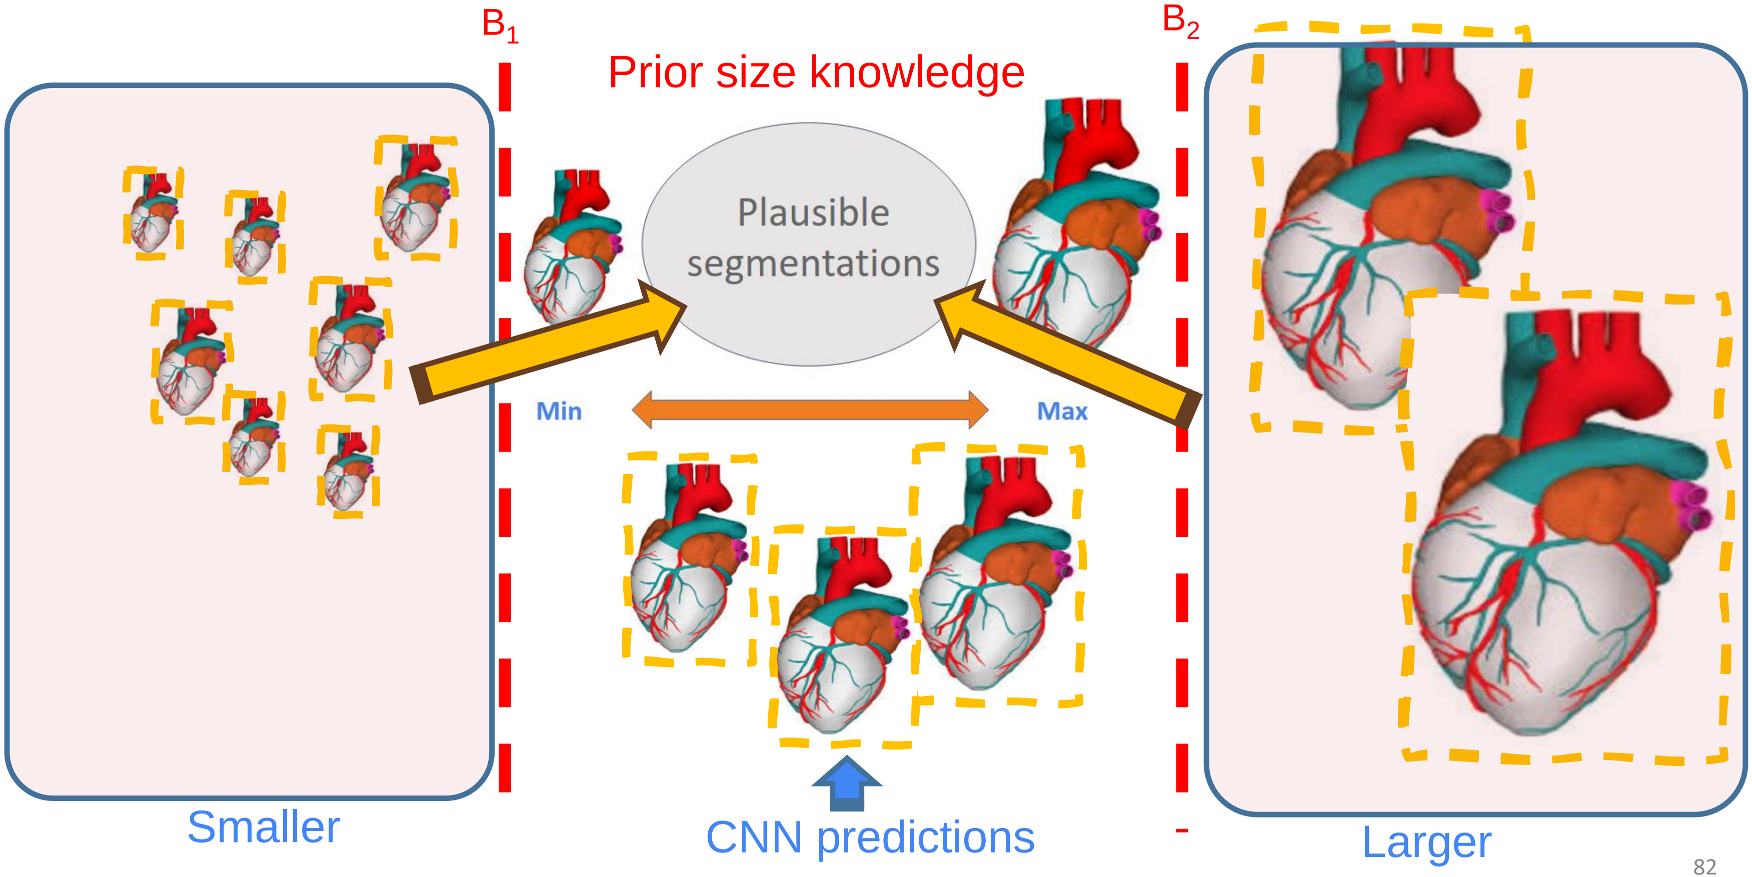

Inequality constraints

Information is given in the form of image-tags

Suppression $$\sum_{p \in \Omega} s_\theta^{p,c} \leq 0 \quad \forall c \notin C$$

"Person"

[Pathak et al., Constrained Convolutional Neural Networks for Weakly Supervised Segmentation, ICCV'15]

Inequality constraints

Information is given in the form of image-tags

Inclusion (or existence)

$$\sum_{p \in \Omega} s_\theta^{p,c} \geq 1 \quad \forall c \in C$$

“Car”

[Pathak et al., Constrained Convolutional Neural Networks for Weakly Supervised Segmentation, ICCV’15]

Inequality constraints

Information is given in the form of image-tags

Target Size
a > 1

$$\sum_{p \in \Omega} s_\theta^{p,c} \geq a \quad \forall c \in C$$

"Car"

[Pathak et al., Constrained Convolutional Neural Networks for Weakly Supervised Segmentation, ICCV'15]

Inequality constraints

Use case:
Size constraint

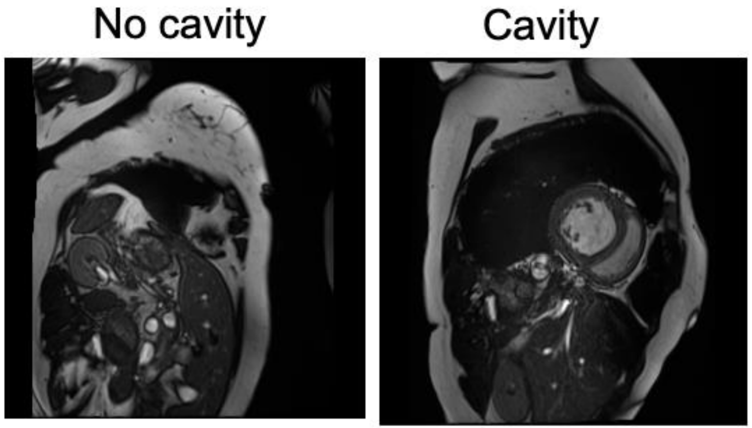

Image-tag information

$$\sum_{p \in \Omega} s_{\theta}^{p,c} \leq 0$$

For negative image tags

Inequality constraints

Use case:
Size constraint

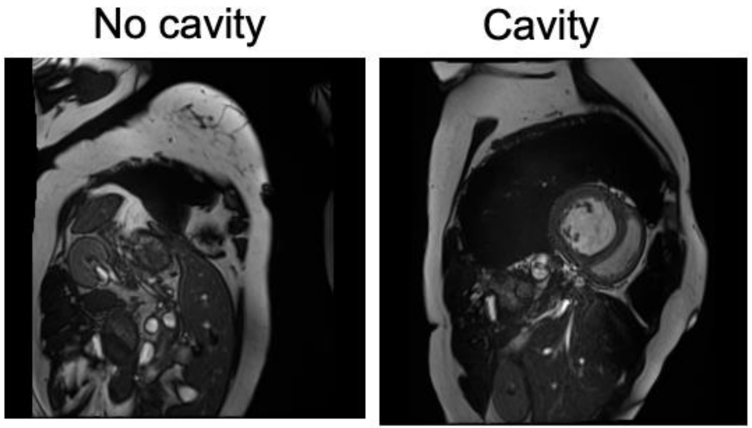

Image-tag information

$$\sum_{p\in\Omega} s_\theta^{p,c} \leq 0$$

For negative image tags

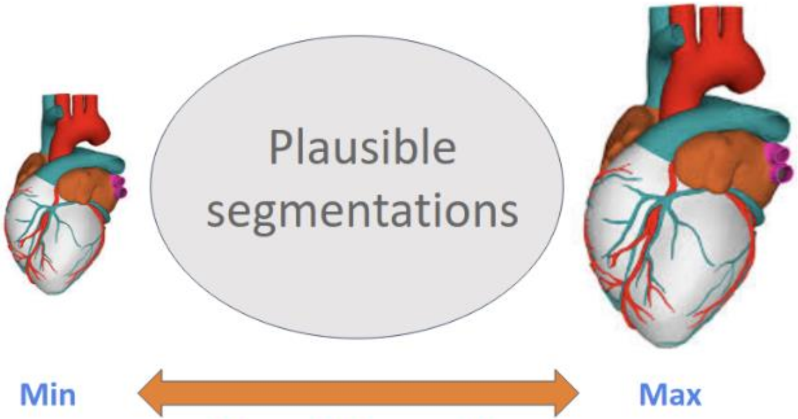

Size information

$$min \leq \sum_{p\in\Omega} s_\theta^{p,c} \leq max$$

For positive image tags

Inequality constraints

Formal definition

Use case:
Size constraint

$$\min_{\boldsymbol{\theta}} \ \mathcal{H}(S) \quad \text{s.t} \quad a \leqslant \sum_{p \in \Omega} S_p \leqslant b$$

Inequality constraint

[Kervadec et al., Constrained-CNN losses for weakly supervised segmentation. MedIA 2019]

Inequality constraints

Formal definition

Use case:
Size constraint

$$\min_{\boldsymbol{\theta}} \ \mathcal{H}(S) \quad \text{s.t} \quad a \leqslant \sum_{p \in \Omega} S_p \leqslant b$$

$$\mathcal{H}(S) + \lambda \, \mathcal{C}\left(V_S\right)$$

$$\boxed{V_S = \sum_{p \in \Omega} \mathbf{s}_{\theta}^{p,c}}$$

[Kervadec et al., Constrained-CNN losses for weakly supervised segmentation. MedIA 2019]

Inequality constraints

Formal definition

Use case:
Size constraint

$$\min_{\boldsymbol{\theta}} \ \mathcal{H}(S) \quad \text{s.t} \quad a \leqslant \sum_{p \in \Omega} S_p \leqslant b$$

$$\mathcal{H}(S) + \lambda \, \mathcal{C}\left(V_S\right)$$

$$\boxed{V_S = \sum_{p \in \Omega} \mathbf{s}_{\theta}^{p,c}}$$

$$\mathcal{H}(S) = -\sum_{p \in \mathcal{L}} log(s_{\theta}^{p})$$

CE on the labeled pixels (if any)

[Kervadec et al., Constrained-CNN losses for weakly supervised segmentation. MedIA 2019]

Inequality constraints

Formal definition

Use case: Size constraint

$$\min_{\boldsymbol{\theta}} \mathcal{H}(S) \quad \text{s.t} \quad a \leqslant \sum_{p \in \Omega} S_p \leqslant b$$

$$\mathcal{H}(S) + \lambda \mathcal{C}(V_S)$$

$$\boxed{V_S = \sum_{p \in \Omega} \mathbf{s}_{\theta}^{p,c}}$$

$$\mathcal{H}(S) = -\sum_{p \in \mathcal{L}} log(s_{\theta}^{p})$$

CE on the labeled pixels (if any)

$$\mathcal{C}(V_S) = \begin{cases} (V_S - a)^2, & \text{if } V_S < a \\ (V_S - b)^2, & \text{if } V_S > b \\ 0, & \text{otherwise} \end{cases}$$

[Kervadec et al., Constrained-CNN losses for weakly supervised segmentation. MedIA 2019]

Inequality constraints

Visual intuition

Use case:
Size constraint

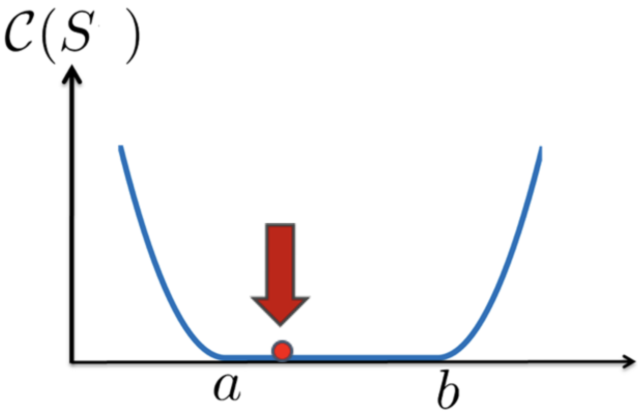

Constraint A satisfied

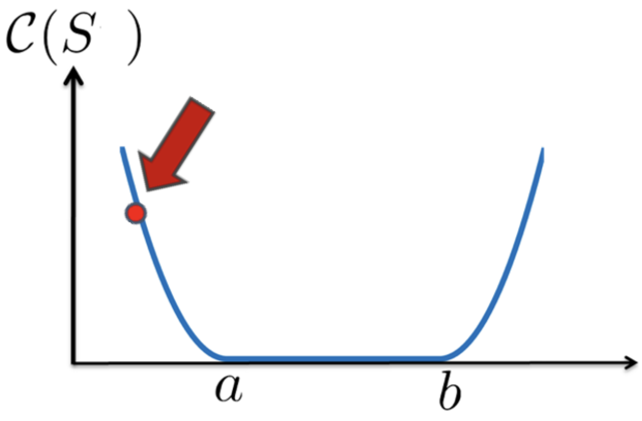

Constraint B violated

[Kervadec et al., Constrained-CNN losses for weakly supervised segmentation. MedIA 2019]

Inequality constraints

Information (proportion) is given as a prior

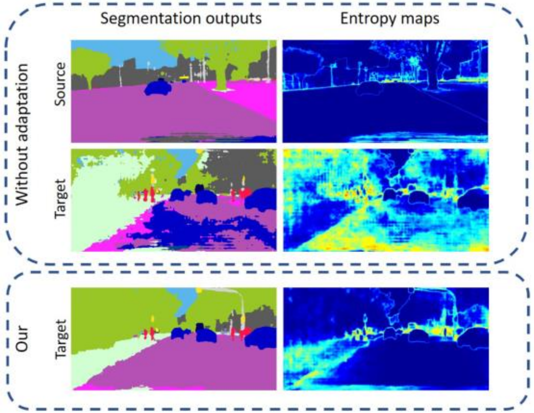

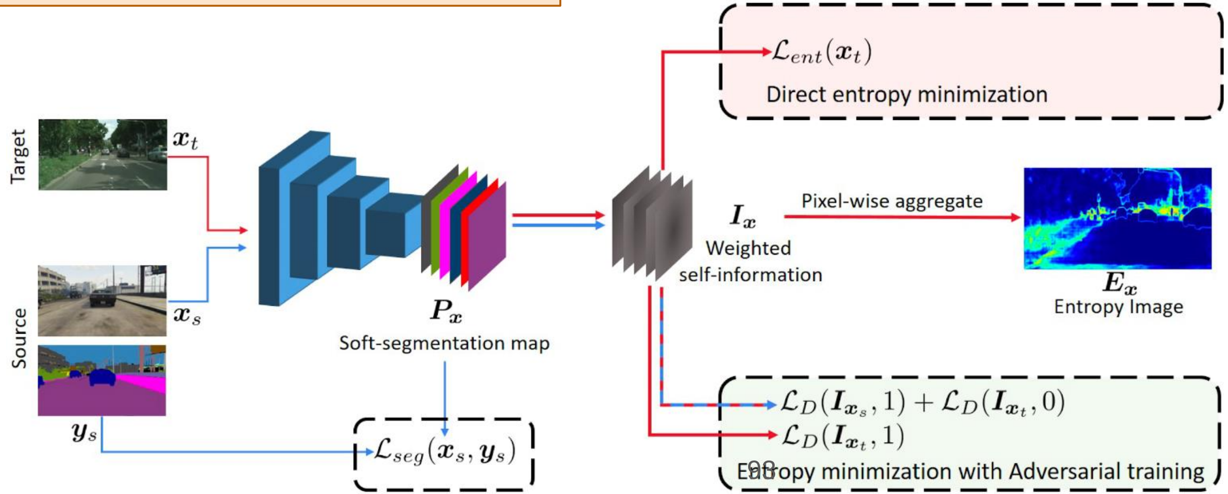

Images from [Vu et al., ADVENT: Adversarial Entropy Minimization for Domain Adaptation in Semantic Segmentation. CVPR'19]

Inequality constraints

Information (proportion) is given as a prior

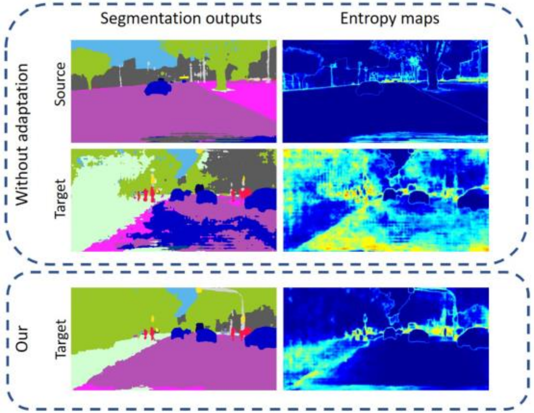

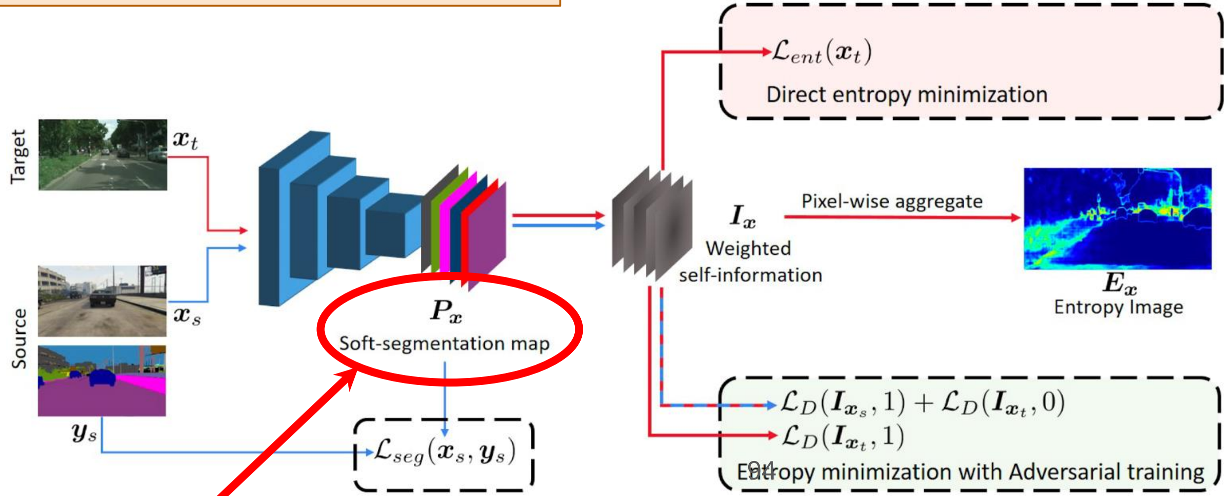

We focus on this now

Images from [Vu et al., ADVENT: Adversarial Entropy Minimization for Domain Adaptation in Semantic Segmentation. CVPR'19]

Inequality constraints

Information (proportion) is given as a prior

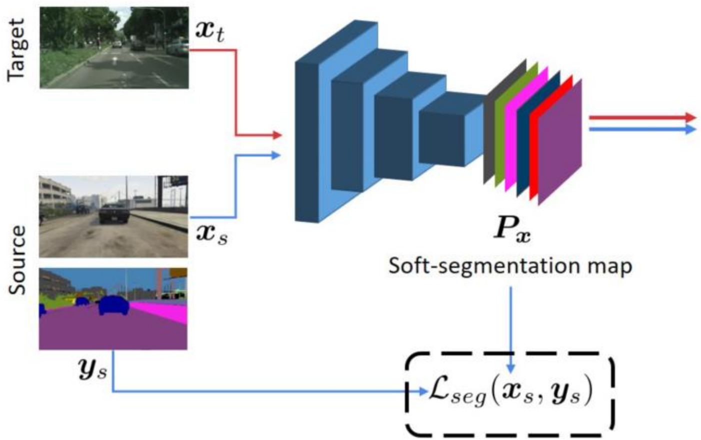

Class-ratio priors

$$\mathcal{L}_{cp}(\boldsymbol{x}_t) = \sum_{c=1}^{C} \max\left(0, \mu \boldsymbol{p}_s^{(c)} - \mathbb{E}_c(\boldsymbol{P}_{\boldsymbol{x}_t}^{(c)})\right)$$

It relaxes the class prior constraint

$\ell 1$-normalized histogram (source)

Estimated size on the prediction

Vu et al., ADVENT: Adversarial Entropy Minimization for Domain Adaptation in Semantic Segmentation. CVPR'19

Inequality constraints

Information (proportion) is given as a prior

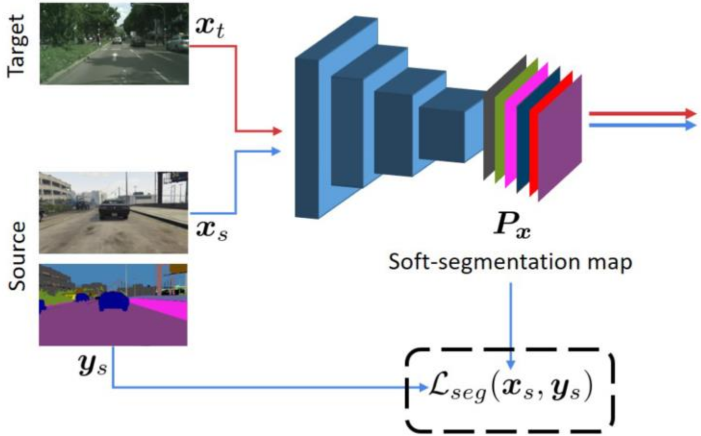

Class-ratio priors

$$\mathcal{L}_{cp}(\boldsymbol{x}_t) = \sum_{c=1}^{C} \max\left(0, \mu \boldsymbol{p}_s^{(c)} - \mathbb{E}_c(\boldsymbol{P}_{\boldsymbol{x}_t}^{(c)})\right)$$

It relaxes the class prior constraint

ℓ1-normalized histogram (source)

Estimated size on the prediction

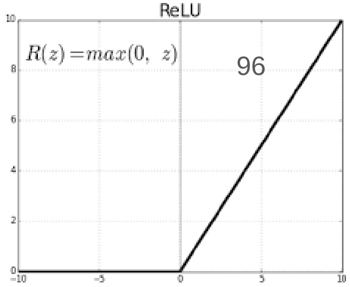

Vu et al., ADVENT: Adversarial Entropy Minimization for Domain Adaptation in Semantic Segmentation. CVPR'19

Inequality constraints

Source-free Domain Adaptation

KL divergence

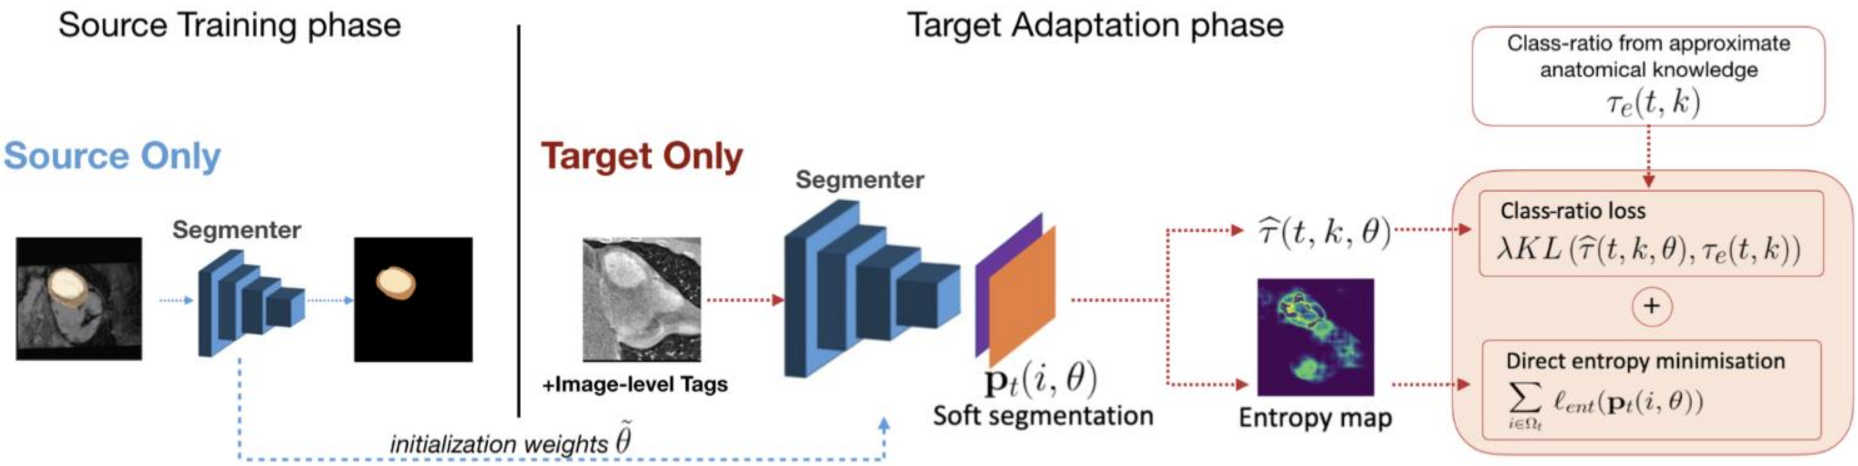

Bateson et al. Source-relaxed domain adaptation for image segmentation. MICCAI'20

Inequality constraints

KL divergence

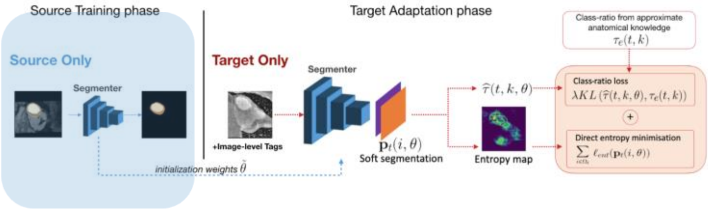

Source-free: no access to source data when adapting

1-Train the network on the source domain

$$\min_{\theta} \sum_{p \in \mathcal{L}} l(\mathbf{y}^p, \mathbf{s}_\theta^p)$$

Set of labeled SOURCE pixels

98

Bateson et al., Source-Relaxed Domain Adaptation for Image Segmentation. MICCAI'19

Inequality constraints

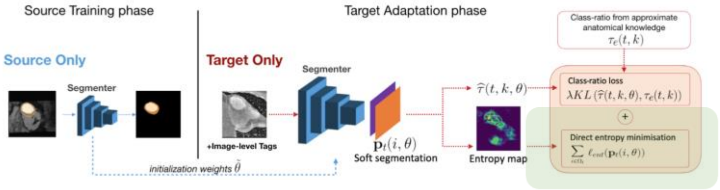

KL divergence

Source-free: no access to source data when adapting

2-Adapt the model without accessing the source data

$$\mathcal{L}_{\mathcal{H}} = -\sum_{p \in \mathcal{T}} \sum_{k} \mathbf{s}_{\theta}^{p,k} \log \mathbf{s}_{\theta}^{p,k} + \mathcal{D}_{KL}(\hat{\tau}, \tau_e)$$

Minimize entropy on predicted TARGET pixels

Bateson et al., Source-Relaxed Domain Adaptation for Image Segmentation. MICCAI'19

Inequality constraints

KL divergence

Source-free: no access to source data when adapting

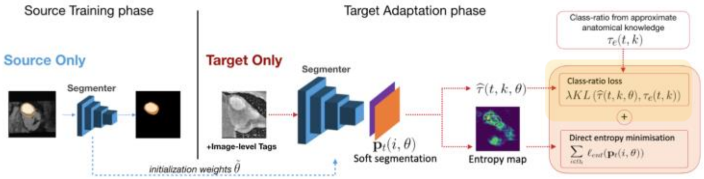

2-Adapt the model without accessing the source data

$\tau_e(t,k)$ Estimated size by an auxiliary network trained on the source

$$\mathcal{L}_{\mathcal{H}} = -\sum_{p\in\mathcal{T}}\sum_{k} \mathbf{s}_\theta^{p,k}\log \mathbf{s}_\theta^{p,k} + \mathcal{D}_{KL}(\hat{\tau}, \tau_e)$$

$$\hat{\tau}(t,k,\theta) = \frac{1}{|\Omega_t|}\sum_{i\in\Omega_t} \mathbf{s}_\theta^{i,k}$$

Minimize entropy on predicted TARGET pixels

Size regularizer

Computed size from the segmentation of the target image

Bateson et al., Source-Relaxed Domain Adaptation for Image Segmentation. MICCAI'19

Inequality constraints

L2 Penalty

But we can do more than simply the size

Shape moment $$\mu_{p,q}^{(k)}(s_{\boldsymbol{\theta}}) := \sum_{i \in \Omega} s_{\boldsymbol{\theta}}^{(i,k)} x_{(i)}^{p} y_{(i)}^{q},$$

Central moment $$\bar{\mu}_{p,q}^{(k)} := \sum_{i \in \Omega} s_{\boldsymbol{\theta}}^{(i,k)} \left(x_{(i)} - \frac{\mu_{1,0}^{(k)}}{\mu_{0,0}^{(k)}} \right)^{p} \left(y_{(i)} - \frac{\mu_{0,1}^{(k)}}{\mu_{0,0}^{(k)}} \right)^{q}.$$

[Kervadec et al., Beyond pixel-wise supervision for segmentation: A few global shape descriptors might be surprisingly good!. MIDL'21]

Inequality constraints

L2 Penalty

But we can do more than simply the size

From shape and central moment

Volume

$$\mathfrak{V}^{(k)}(s_{\boldsymbol{\theta}}) := \mu_{0,0}^{(k)}(s_{\boldsymbol{\theta}}).$$

Centroid

$$\mathfrak{C}^{(k)}(s_{\boldsymbol{\theta}}) := \left(\frac{\mu_{1,0}^{(k)}(s_{\boldsymbol{\theta}})}{\mu_{0,0}^{(k)}(s_{\boldsymbol{\theta}})}, \frac{\mu_{0,1}^{(k)}(s_{\boldsymbol{\theta}})}{\mu_{0,0}^{(k)}(s_{\boldsymbol{\theta}})}\right).$$

Length

Laplacian

$$\mathfrak{L}^{(k)}(s_{\boldsymbol{\theta}}) := \sum_{i,j \in \mathcal{G}_\Omega} |s_{\boldsymbol{\theta}}^{(i,k)} - s_{\boldsymbol{\theta}}^{(j,k)}| L_{\Omega,i,j}.$$

[Kervadec et al., Beyond pixel-wise supervision for segmentation: A few global shape descriptors might be surprisingly good!. MIDL'21]

Inequality constraints

L2 Penalty

But we can do more than simply the size

From shape and central moment

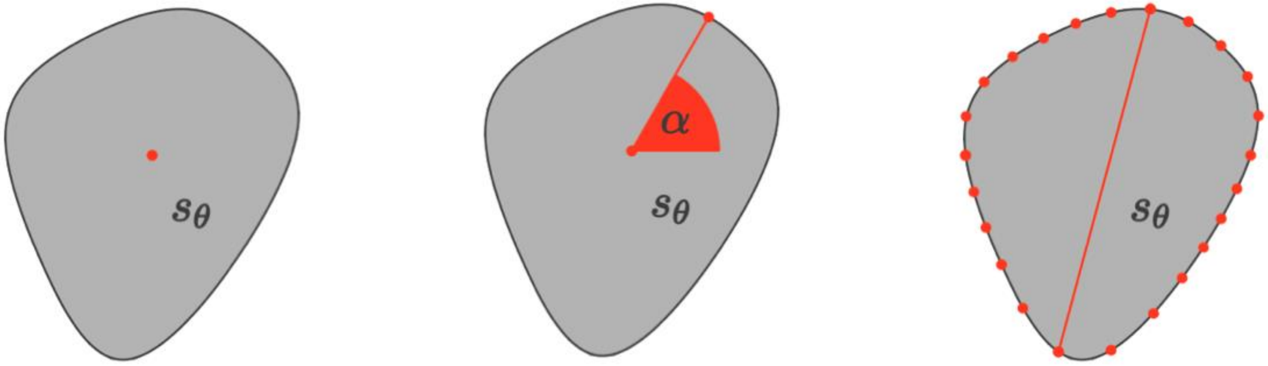

(a) Centroid $\mathfrak{C}(s_\theta)$ (b) Radius $\widehat{\mathfrak{R}}_\beta(s_\theta, \alpha)$ (c) Diameter $\mathfrak{D}(s_\theta)$

Camarasa et al. Differentiable Boundary Point Extraction for Weakly Supervised Star-shaped Object Segmentation. MIDL'22

Inequality constraints

L2 Penalty

Test-Time Adaptation (TTA)

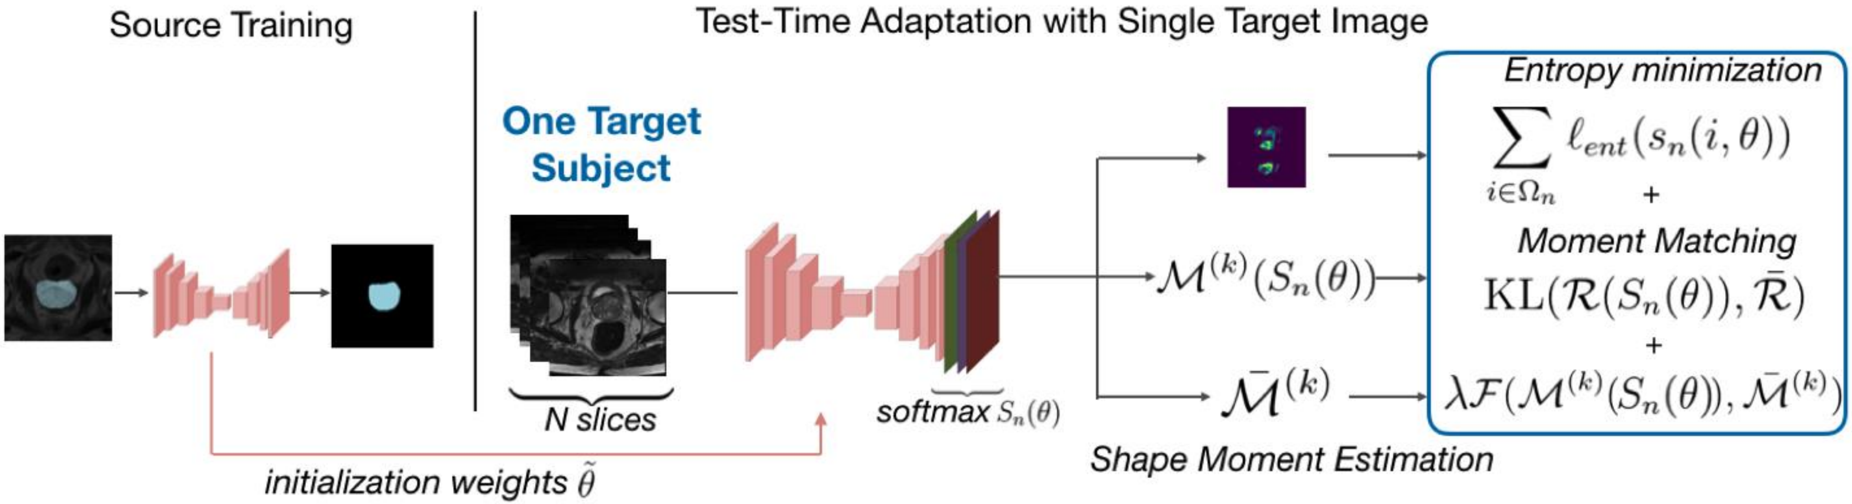

Bateson et al., Test-Time Adaptation with Shape Moments for Image Segmentation. MICCAI'22

Take-home message

- Imposing constraints helps weakly-supervised segmentation learning by restricting plausible segmentations on weakly labeled and unlabeled images

- Few constraints have been explored under low-labeled data regime

- Room for improvement (many opportunities beyond weakly supervised segmentation)

Thank you!