Test-Time Adaptable CNNs for Robust Medical Image Segmentation

MIDL 2021 | Short Paper 76

Neerav Karani
Krishna Chaitanya
Dr. Ertunc Erdil
Prof. Ender Konukoglu

Biomedical Image Computing Group
Computer Vision Institute
ETH Zurich

The Domain Shift Issue

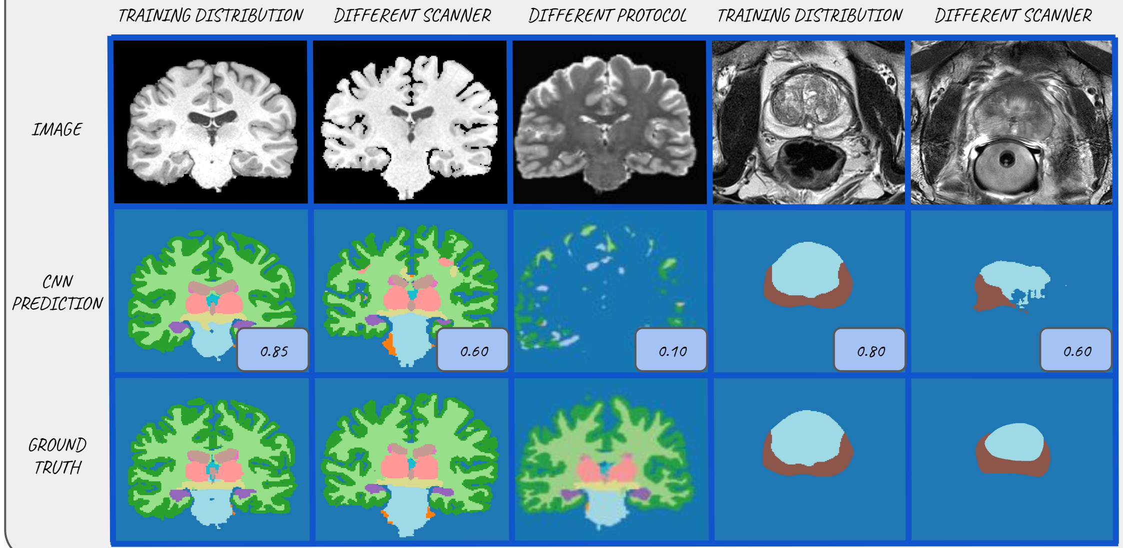

Take Home Message (1)

CNNs are not robust against domain shifts.

Relatively small differences in image statistics can cause substantial performance degradation.

HYPOTHESIS → FOR GENERALIZATION ACROSS DOMAIN SHIFTS, WE NEED PER-TEST-IMAGE ADAPTABILITY.

Q1. Which parameters to adapt?

Q2. How to drive the adaptation?

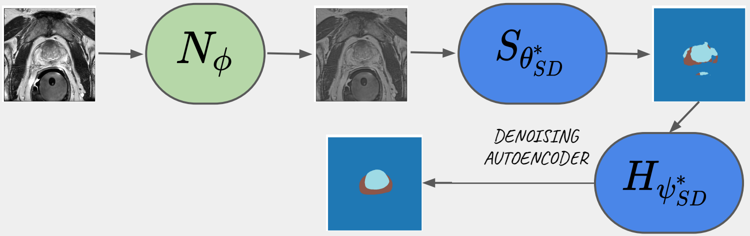

Train N, S on SD. Now, fix S and adapt N for each test image. N is shallow to allow contrast modulations without substantial structural changes.

Encourage predicted labels to become *plausible*.

$min_\phi L(y_n, H_{\psi^*_{SD}}(y_n)) \qquad y_n = S_{\theta^*_{SD}}(N_\phi(x_{TI}))$

Test-Time Adaptation Iterations

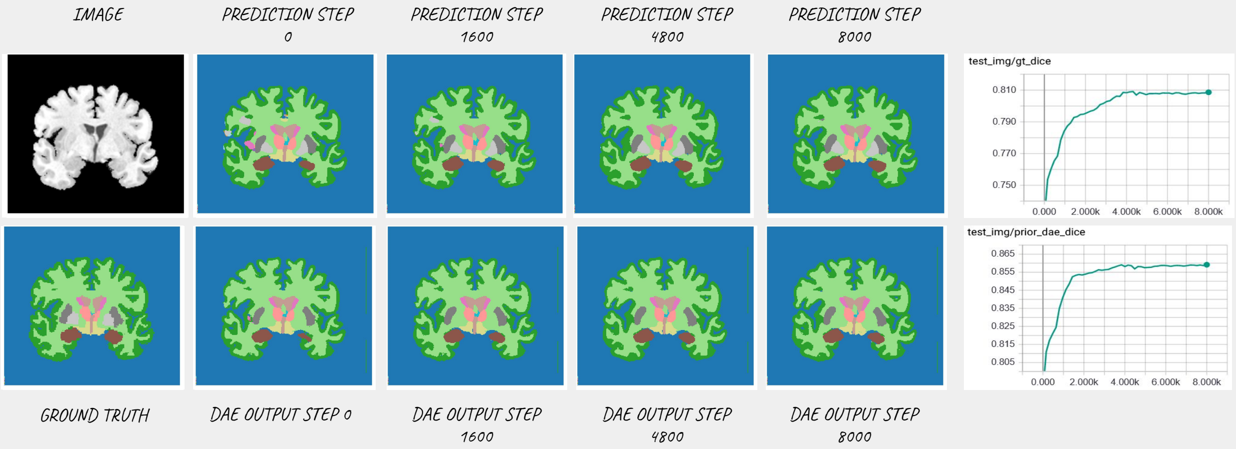

Concurrent Works

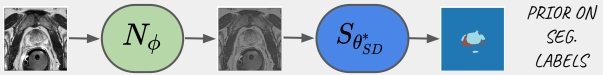

Karani et al. "Test-time adaptable neural networks for robust medical image segmentation." Medical Image Analysis 2021.

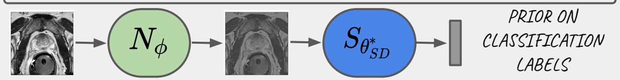

Wang et al. "TENT: Fully Test-time Adaptation by Entropy Minimization." ICLR 2021

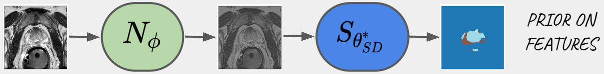

Sun et al. "Test-Time Training with Self-Supervision for Generalization under Distribution Shifts." ICML 2020

Yufan et al. Self domain adapted network. MICCAI 2020

ML settings for tackling domain shifts

Setup	Source domain		Target domain	
	Data	Algorithm	Data	Algorithm
Learning a new CNN in each new hospital	$\{x_{SD}, y_{SD}\}(many)$	$\min_\theta L^{seg}_{SD}$	$\{x_{TD}, y_{TD}\}(many)$	$\min_\theta L^{seg}_{TD}$
Transfer Learning	$\{x_{SD}, y_{SD}\}(many)$	$\min_\theta L^{seg}_{SD}$	$\{x_{TD}, y_{TD}\}(few)$	Init. at θ^*_{SD}, $\min_\theta L^{seg}_{TD}$
Unsupervised Domain Adaptation	-	-	$\{x_{SD}, y_{SD}, x_{TD}\}(many)$	$\min_\theta L^{seg}_{SD} + L^{inv}_{SD,TD}$
Domain Generalization	$\{x_{SD}, y_{SD}\}(many)$	$\min_\theta L^{seg}_{SD} + L^{inv}_{SD}$	x_{TI}	$\hat{y} = S_{\theta^*_{SD}}(x_{TI})$
Domain Generalization with test time adaptation	$\{x_{SD}, y_{SD}\}(many)$	$\min_\theta L^{seg}_{SD} + L^{inv}_{SD}$	x_{TI}	Adapt θ for each test image. $\hat{y} = S_{\theta^*_{TI}}(x_{TI})$
Unsupervised Learning	-	-	x_{TI}	Optimize θ for each test image. $\hat{y} = \mathrm{argmax}_y\, P(y)P(x_{TI}\|y)$

Take Home Message (2)

Adapting a CNN for each test image increases robustness to domain shifts.

For image segmentation, a prior on output labels can be used to drive such adaptation.