On Direct Distribution Matching for Adapting Segmentation Networks

MIDL 2020

Georg Pichler, Jose Dolz, Ismail Ben Ayed, and Pablo Piantanida

TU Wien, Austria &
ÉTS, Montreal, Canada &
CentraleSupélec-CNRS-Université Paris Sud &
Montreal Institute for Learning Algorithms (Mila), QC, Canada

Domain Adaptation in Segmentation Networks

- Source domain images X; ground truth labels Y
- A segmentation function f is trained on labeled source data $\mathcal{L} = \{(X_i, Y_i)\}_{i=1,\dots,n}$
- Images X' from a different, target domain:
  - taken with a different camera,
  - taken with a different MR/CT/X-ray machine, ...
- $f(X') \neq Y'$
- Domain Adaptation (DA): Obtain f' with good performance on X', given $\mathcal{L}$ and unlabeled pairs of source/target domain images $\mathcal{U} = \{(X_{n+1}, X'_{n+1}), \dots, (X_{n+m}, X'_{n+m})\}$

Prior art

- Previous work dominated by adversarial approaches (Goodfellow et al. (2014))
  - Y.-H. Tsai et al. (2018). "Learning to adapt structured output space for semantic segmentation". In: *Computer Vision and Pattern Recognition (CVPR)*
- Adversary can operate at output (segmentation) level
- Or image alignment at pixel/intermediate level:
  - Transform the source images into the style of the target images
  - Then train the segmentation network on artificial target images
  - Downside: only work well on narrow shifts between source and target domain

Domain Adaptation for Medical Images

- Possibility to obtain images of the same patient with different imaging methods (machines/protocols/cameras. . .)
  $\Longrightarrow$ Gap in appearance, but identical spacial layout

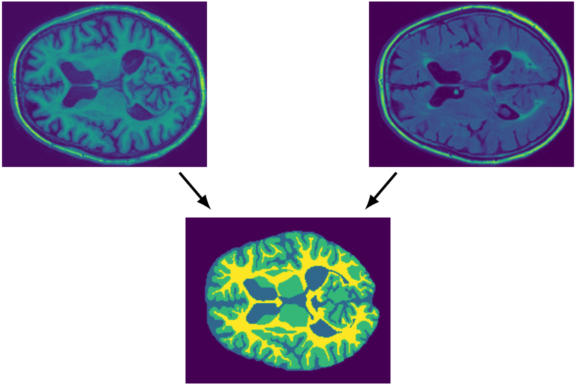

- Goal: Training one segmentation function f that works on both source and target domain
- Idea: Use $\mathcal{U}$ to enforce $f(X) \approx f(X')$

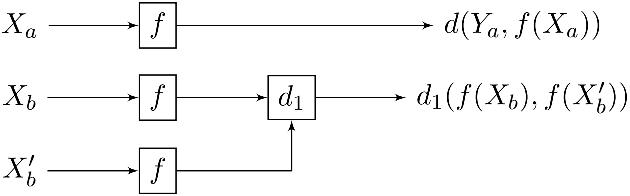

- Utilize (C)NN architecture: f_θ with parameter θ
- Loss:

$$\mathcal{F}(\theta) = \sum_{i=1}^{n} d\big(Y_i, f_\theta(X_i)\big) + \lambda \sum_{i=n+1}^{n+m} d_1\big(f_\theta(X_i), f_\theta(X'_i)\big)$$

- Choices: f_θ, d_1, d, λ

Experiments

Segmentation Network:

- f_θ : slightly modified U-Net (Ronneberger, Fischer, and Brox, 2015)

Datasets

- Human brain MR images
  - iSEG challenge dataset (Wang et al., 2019)
  - MRBrainS2013 challenge dataset (Mendrik et al., 2015)
- Segmentation in 3 classes: GM, WM, CSF
- X, X': Aligned T1/T2(-FLAIR) scans of the same patient
- d, d_1: cross entropy loss
- Three runs for cross-validation
- Figure of merit: average DICE over all three classes

Mean DICE

- Oracle: U-Net network trained on target domain
- No Adaptation: U-Net network trained on source domain only
- AdaptSegNet: (Tsai et al., 2018) with U-Net segmentation net.

Targ.	Oracle	No adaptation	AdaptSegNet	Proposed
T2*	77.35 ± 1.35	38.58 ± 1.14	56.62 ± 8.02	76.10 ± 0.45
T1*	84.71 ± 0.98	20.25 ± 3.54	73.22 ± 2.16	82.43 ± 0.50
T2†	76.89 ± 0.67	38.70 ± 10.46	63.37 ± 6.25	74.17 ± 0.78
T1†	82.28 ± 0.88	66.26 ± 0.53	70.11 ± 3.00	77.89 ± 1.15

- Asymmetry between T1 $\rightarrow$ T2 (harder) and T2 $\rightarrow$ T1 (easier) (also noted by Dou et al., 2018)

* MRBrainS 2013
† iSEG

Summary

Domain adaptation in semantic segmentation of MR images

- Additional structure in data (e.g. alignment) should be utilized!

In the paper:

- Stability during training
- Violation of alignment assumption
- Impact of distance function d_1 and Lagrangian λ

References I

Dou, Q., C. Ouyang, C. Chen, H. Chen, B. Glocker, X. Zhuang, and P.-A. Heng (2018). “PnP-AdaNet: Plug-and-Play Adversarial Domain Adaptation Network with a Benchmark at Cross-modality Cardiac Segmentation”. In: *arXiv preprint arXiv:1812.07907*.

Goodfellow, I., J. Pouget-Abadie, M. Mirza, B. Xu, D. Warde-Farley, S. Ozair, A. Courville, and Y. Bengio (2014). “Generative adversarial nets”. In: *Advances in neural information processing systems*, pp. 2672–2680.

Mendrik, A. M., K. L. Vincken, H. J. Kuijf, M. Breeuwer, W. H. Bouvy, J. De Bresser, A. Alansary, M. De Bruijne, A. Carass, A. El-Baz, et al. (2015). “MRBrainS challenge: online evaluation framework for brain image segmentation in 3T MRI scans”. In: *Computational intelligence and neuroscience* 2015, p. 1.

Ronneberger, O., P. Fischer, and T. Brox (2015). “U-Net: Convolutional networks for biomedical image segmentation”. In: *International Conference on Medical image computing and computer-assisted intervention*. Springer, pp. 234–241.

Tsai, Y.-H., W.-C. Hung, S. Schulter, K. Sohn, M.-H. Yang, and M. Chandraker (2018). “Learning to adapt structured output space for semantic segmentation”. In: *Computer Vision and Pattern Recognition (CVPR)*.

References II

Wang, L. et al. (2019). "Benchmark on Automatic 6-month-old Infant Brain Segmentation Algorithms: The iSeg-2017 Challenge". In: *IEEE Transactions on Medical Imaging*, pp. 1–1.