

Master's Thesis on Semi-Supervised Semantic Segmentation via Robust Pseudo-labels and Reversible Image Transformations

Advisor: Simon Reiß (simon.reiss@kit.edu)

1 Motivation and relevance

Deployment of semantic segmentation systems is often slowed down by the time-consuming task of data annotation. Especially when expensive experts need to provide pixel-wise annotations as, e.g., in medical image segmentation tasks. Finding strategies to get more out of a few labeled examples can help immensely when training segmentation models and bringing them to real applications.

2 Coarse topic and starting point

One component in semi-supervised segmentation methods can be inferring pseudo-labels for previously unlabeled (or weakly labeled) images [5, 6].

For a more robust pseudo-label generation, self-supervised literature [1, 3] might offer solutions. In this thesis, I would like to explore together with you how we can use tricks from self-supervised learning and combine them with sophisticated geometric image transformations as used in semi-supervised segmentation [2, 4]. Your task will be to develop, explore and evaluate different pseudo-labeling strategies for semantic segmentation to make best use of images without pixel-wise annotations. The data we work on will be either medical imagery or a standard object-centric/scene segmentation dataset.

3 Requirements

English language skills and good programming skills in Python are required. Some basic knowledge about the machine learning framework <https://pytorch.org/> is recommended. You should have a solid understanding of deep learning models, e.g., by having attended lectures like {Deep Learning for Computer Vision, Computer Vision for Human Computer Interaction, Neural Networks, Machine Learning}. Take a quick glance at the referenced papers, if the topic sounds exciting, interesting and you feel like it's a big enough challenge for you, feel free to contact me.

References

- [1] Xinlei Chen and Kaiming He. Exploring simple siamese representation learning. *arXiv preprint arXiv:2011.10566*, 2020.
- [2] Geoffrey French, Samuli Laine, Timo Aila, Michal Mackiewicz, and Graham Finlayson. Semi-supervised semantic segmentation needs strong, varied perturbations. In *British Machine Vision Conference*, number 31, 2020.
- [3] Xu Ji, João F Henriques, and Andrea Vedaldi. Invariant information clustering for unsupervised image classification and segmentation. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 9865–9874, 2019.
- [4] Viktor Olsson, Wilhelm Tranheden, Juliano Pinto, and Lennart Svensson. Classmix: Segmentation-based data augmentation for semi-supervised learning. In *Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision*, pages 1369–1378, 2021.
- [5] Qizhe Xie, Minh-Thang Luong, Eduard Hovy, and Quoc V Le. Self-training with noisy student improves imagenet classification. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 10687–10698, 2020.
- [6] Barret Zoph, Golnaz Ghiasi, Tsung-Yi Lin, Yin Cui, Hanxiao Liu, Ekin D Cubuk, and Quoc V Le. Rethinking pre-training and self-training. *arXiv preprint arXiv:2006.06882*, 2020.