Shape-intensity knowledge distillation for robust medical image segmentation

Wenhui DONG, Bo DU, Yongchao XU

Frontiers of Computer Science, DOI: 10.1007/s11704-024-40462-2

Problems & Ideas

- Problems of existing deep learning-based methods:
  - The model's bias towards learning texture information leads to unsatisfied generalization performance.
  - Most existing methods do not take into account the shape-intensity prior information. This may lead to implausible segmentation results.
- Ideas: Compute class-wise average images as training data for the teacher network to extract shape and intensity priors, and utilize knowledge distillation to transfer these priors to the segmentation network..

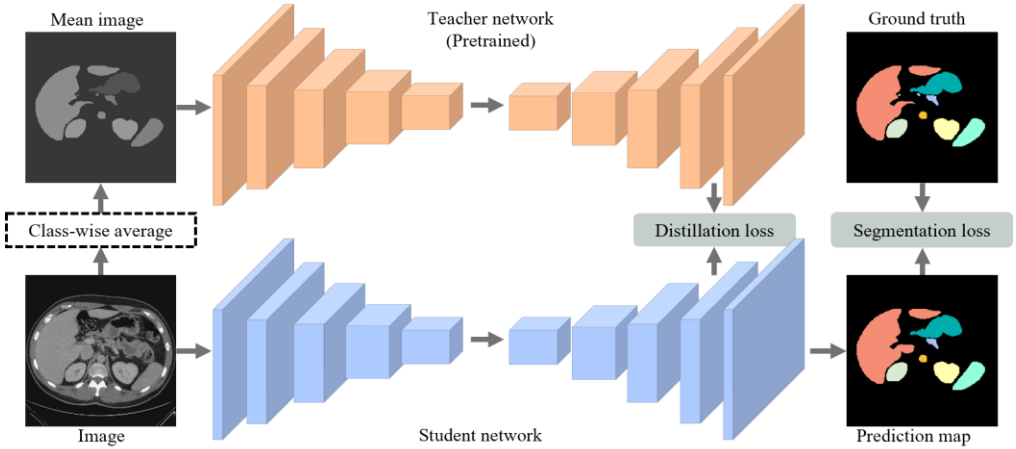

The pipeline of the proposed method. The teacher and student model have the same network architecture, trained on class-wise averaged training images and original images with segmentation loss, respectively. For the student model, we also apply the distillation loss on the penultimate layer between the teacher and student model to transfer the shape-intensity knowledge.

Main Contributions

- Contributions:
  - We propose to train a network on class-wise averaged training images to explicitly extract shape-intensity prior information;
  - We then leverage knowledge distillation to transfer the prior to segmentation network;
  - Extensive experiments demonstrate that the proposed method consistently/significantly improves the baseline models and has a better generalization ability to images of unseen datasets..

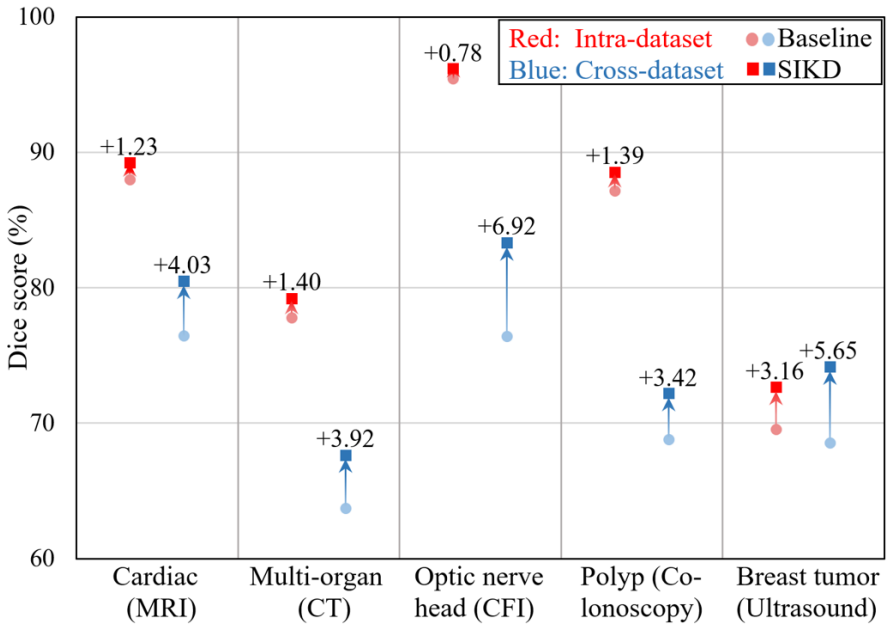

In five medical image segmentation tasks with different modalities, the proposed method outperforms the corresponding baseline models in both intra-dataset (red) and cross-dataset (blue) settings.