RED-GAN: Attacking class imbalance via conditioned generation. Yet another medical imaging perspective.

Ahmad B Qasim*, Ivan Ezhov*, Suprosanna Shit, Oliver Schoppe, Johannes C Paetzold, Anjany Sekuboyina, Florian Kofler, Jana Lipkova, Hongwei Li, Bjoern Menze

Department of Informatics, Technical University of Munich

Motivation:

- In the medical imaging context, the GAN-based augmentation has been applied to classification [1] and segmentation problems [2].

- However, all these studies do not differentiate between a varying global attribute (e.g. MRI acquisition set-up)

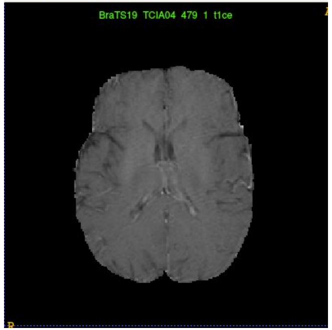

BraTS, T1ce
TCIA center

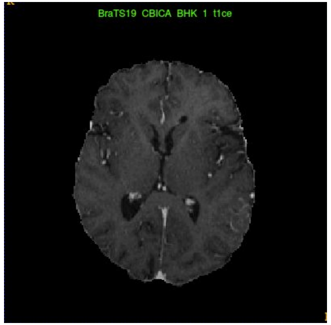

BraTS, T1ce
CBICA center

1. Frid-Adar et al., 2018; Gupta et al., 2019
2. Mok and Chung, 2018; Thomas et al., 2017; Zhang et al., 2018b; Bowles et al., 2018; Rezaei et al., 2017; Abhishek and Hamarneh, 2019

Contribution:

(1) We propose a GAN design conditioned on such *global* information in addition to the *local* one (segmentation masks). This allows controlling class specific appearance of the generated images.

(2) We incorporate a third player in the adversarial game to stimulate synthesis of the features relevant for the downstream segmentation task.

Class imbalance (semantic vs global)

- Semantic (pixel-wise) classes
  (e.g. different lesion areas, different healthy tissues)

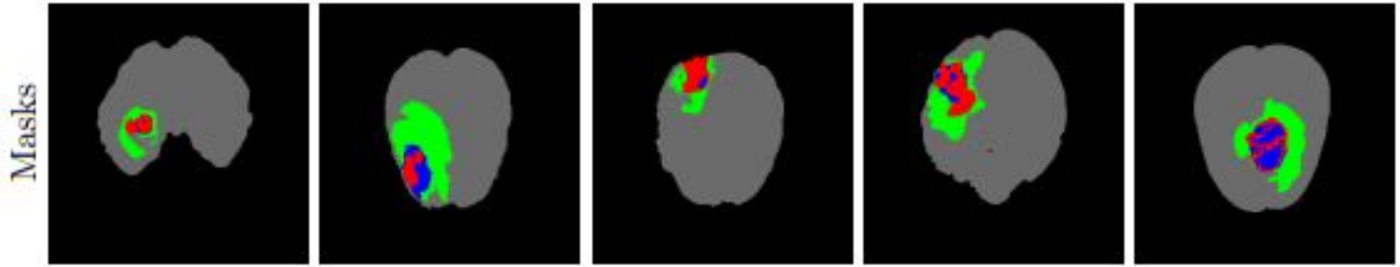

Brain tumor segmentation maps from BraTS dataset

Class imbalance (semantic vs global)

- Global classes
  (e.g. different acquisition environment)

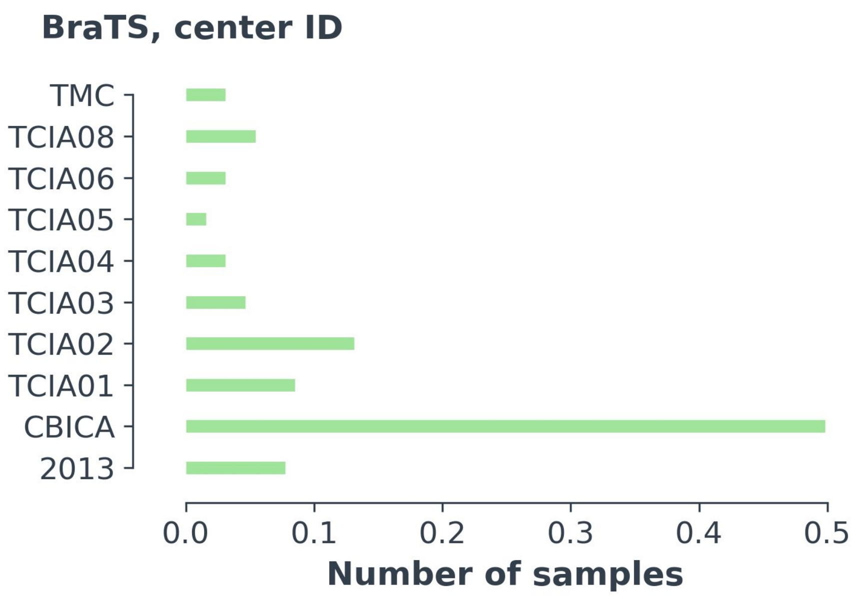

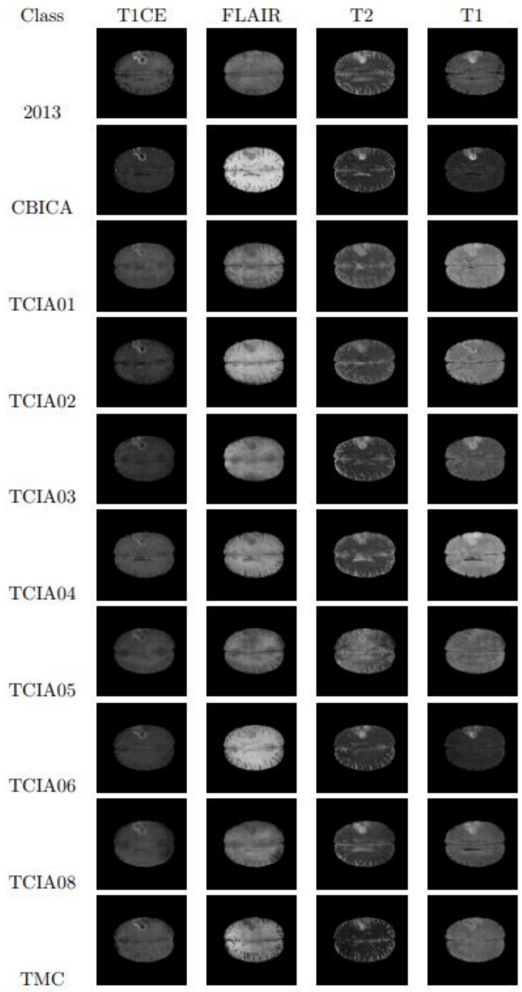

SPADE-GAN (conditioning on the mask, [3])

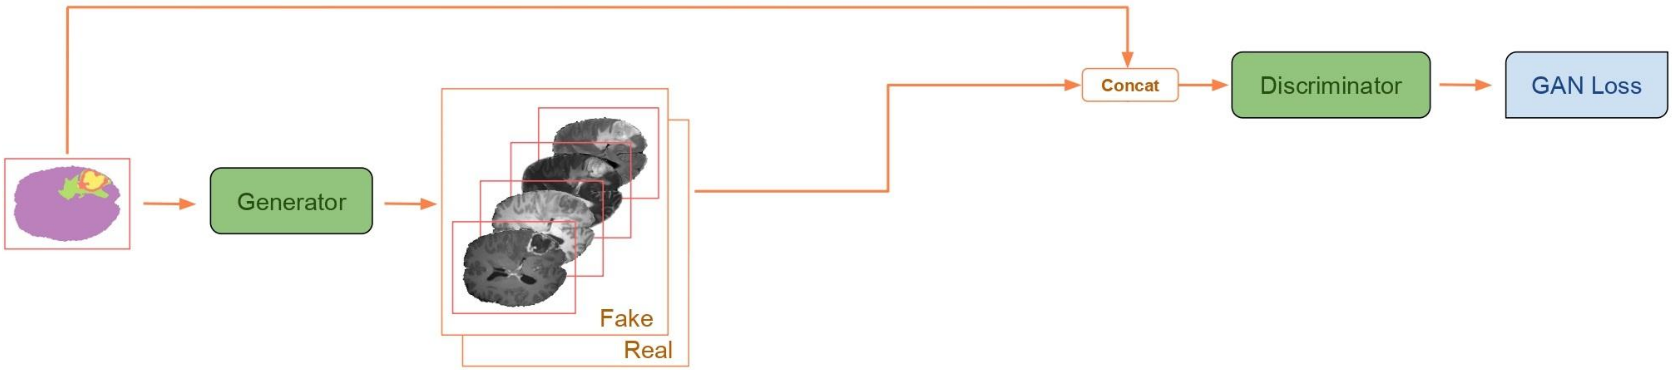

[3] Taesung Park et al., Semantic image synthesis with spatially-adaptive normalization. CVPR 2019.

Red-GAN (conditioning on both: the mask and the global label)

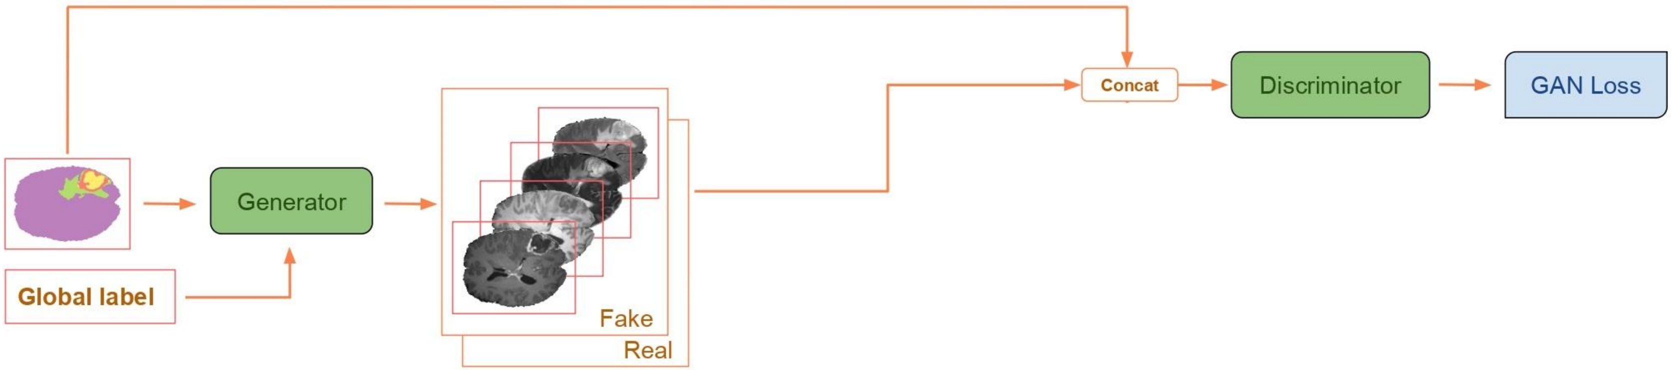

SPADE-GAN
(semantic map conditioning)

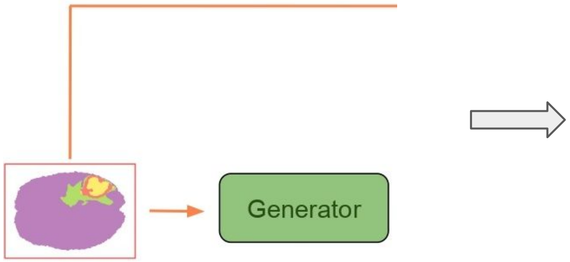

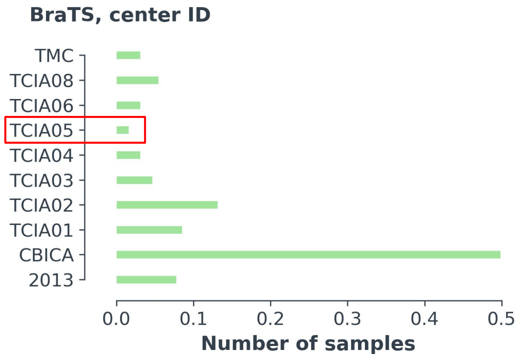

SPADE-GAN
(semantic map conditioning)

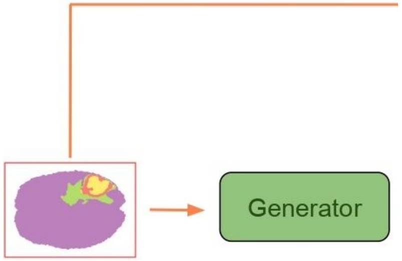

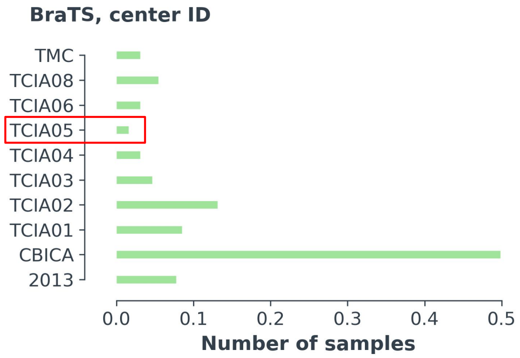

Baseline Dice: 0.74
(no augmentations)

Dice for U-Net trained on the BraTS dataset augmented with synthetic TCIA05 images: 0.683

SPADE-GAN
(semantic map conditioning)

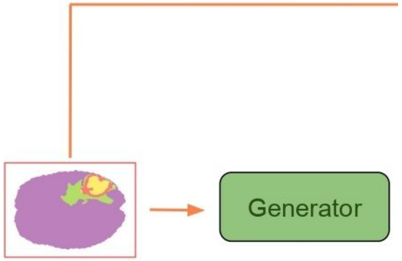

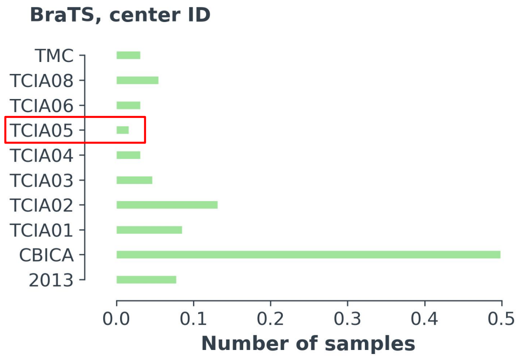

Baseline Dice: 0.74
(no augmentations)

Dice for U-Net trained on the BraTS dataset augmented with synthetic TCIA05 images: 0.683

Red-GAN (semantic map and global class conditioning)

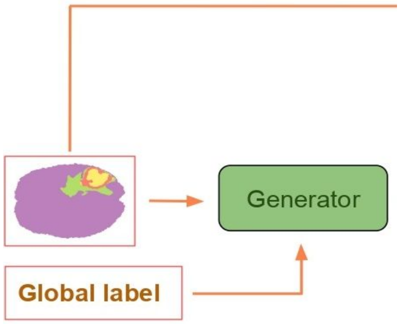

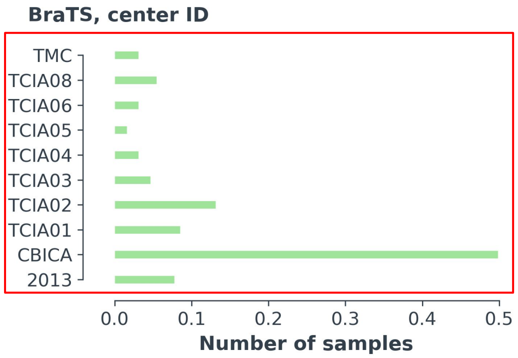

Dice for U-Net trained on the BraTS dataset augmented with synthetic TCIA05 images: 0.779

Contribution:

(1) We propose a GAN design conditioned on such *global* information in addition to the *local* one (segmentation masks). This allows controlling class specific appearance of the generated images.

(2) We incorporate a third player in the adversarial game to stimulate synthesis of the features relevant for the downstream segmentation task.

Red-GAN (adding a third player)

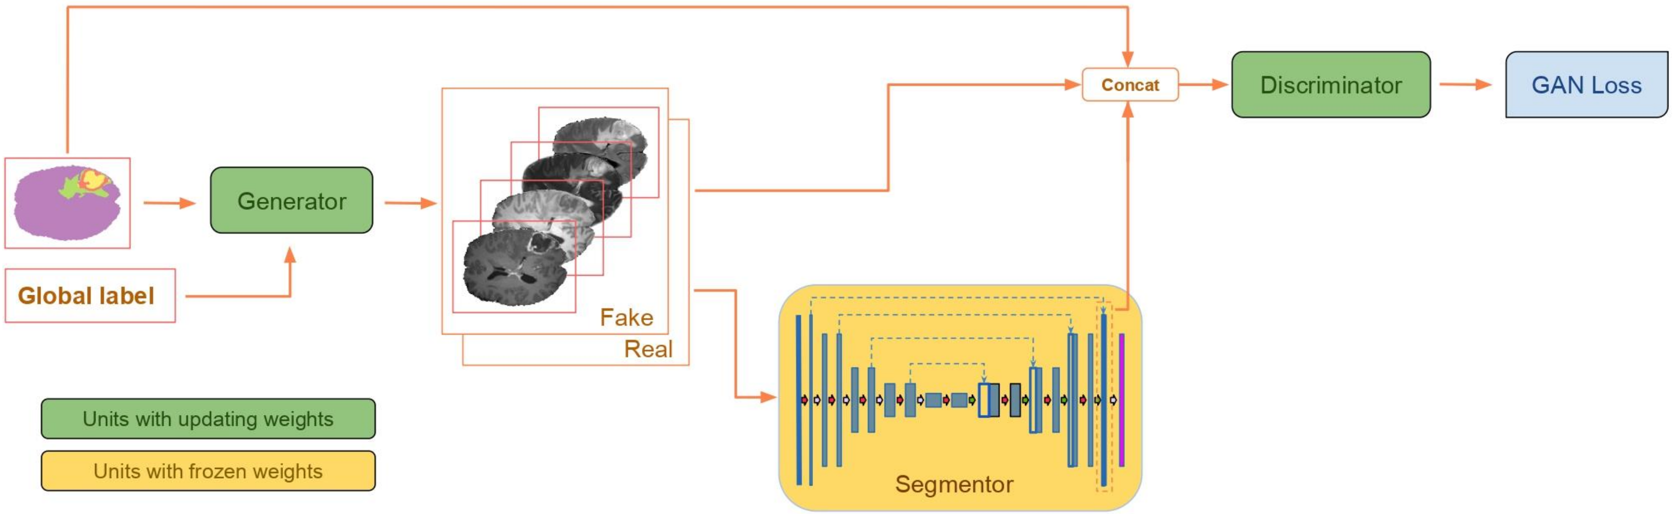

The third player (segmentor) ensures that the synthetic images lie within close proximity to the real images in the latent representation, based on which the downstream segmentation network makes its decision

Results:

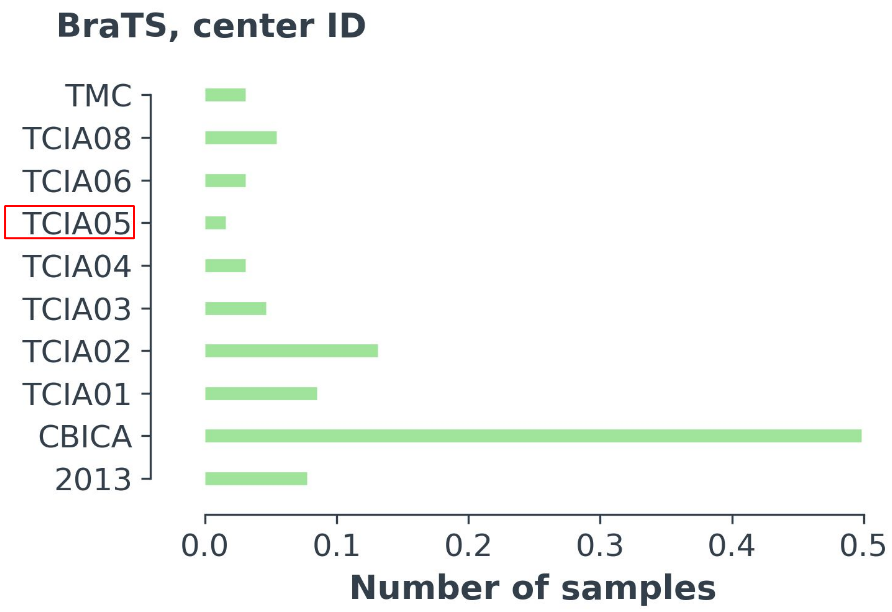
BraTS, center ID
TMC
TCIA08
TCIA06
TCIA05
TCIA04
TCIA03
TCIA02
TCIA01
CBICA
2013
0.0
0.1
0.2
0.3
0.4
0.5
Number of samples

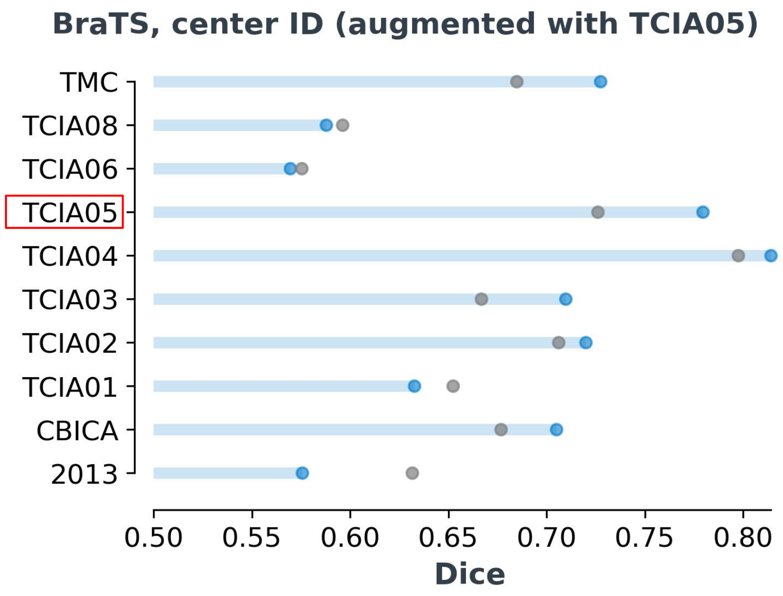
BraTS, center ID (augmented with TCIA05)
TMC
TCIA08
TCIA06
TCIA05
TCIA04
TCIA03
TCIA02
TCIA01
CBICA
2013
0.50
0.55
0.60
0.65
0.70
0.75
0.80
Dice

Results:

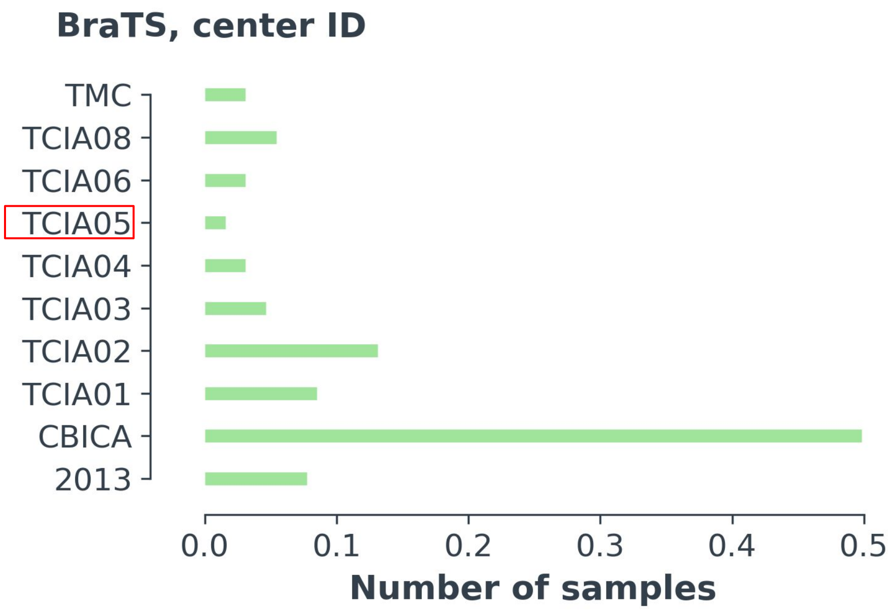
BraTS, center ID
TMC
TCIA08
TCIA06
TCIA05
TCIA04
TCIA03
TCIA02
TCIA01
CBICA
2013
0.0
0.1
0.2
0.3
0.4
0.5
Number of samples

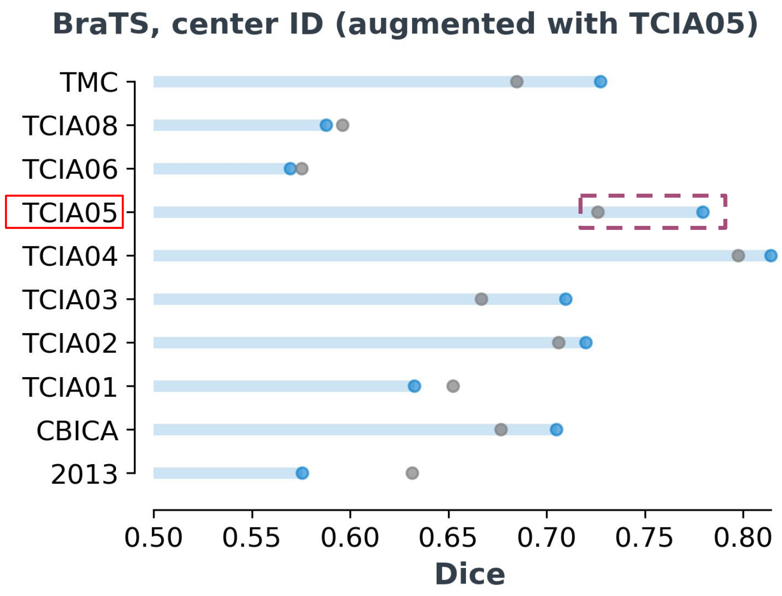
BraTS, center ID (augmented with TCIA05)
TMC
TCIA08
TCIA06
TCIA05
TCIA04
TCIA03
TCIA02
TCIA01
CBICA
2013
0.50
0.55
0.60
0.65
0.70
0.75
0.80
Dice

Thank you for your attention

https://github.com/IvanEz/Red-GAN

What’s worse visiting out there at MIDL?

Paper 182: Deep learning-based parameter mapping for joint relaxation and diffusion tensor MR Fingerprinting

Paper 128: Deep Reinforcement Learning for Organ Localization in CT