Back to Rectal Cancer

AUTOMATIC SEGMENTATION OF RECTAL CANCER FOR THE ORIGINAL DATASET

Brief Summary of the Last Episode

- **MRI images of patients affected by colon cancer**
- Manual segmentation done by expert clinicians
- First approach to automatic segmentation with Fully Convolutional Neural Networks (FCNN)
- PROMISE – search for an inspiration

Brief Summary of the Last Episode

- MRI images of patients affected by colon cancer
- **Manual segmentation done by expert clinicians**
- First approach to automatic segmentation with Fully Convolutional Neural Networks (FCNN)
- PROMISE – search for an inspiration

Brief Summary of the Last Episode

- MRI images of patients affected by colon cancer
- Manual segmentation done by expert clinicians
- **First approach to automatic segmentation with Fully Convolutional Neural Networks (FCNN)**
- PROMISE – Search for an inspiration

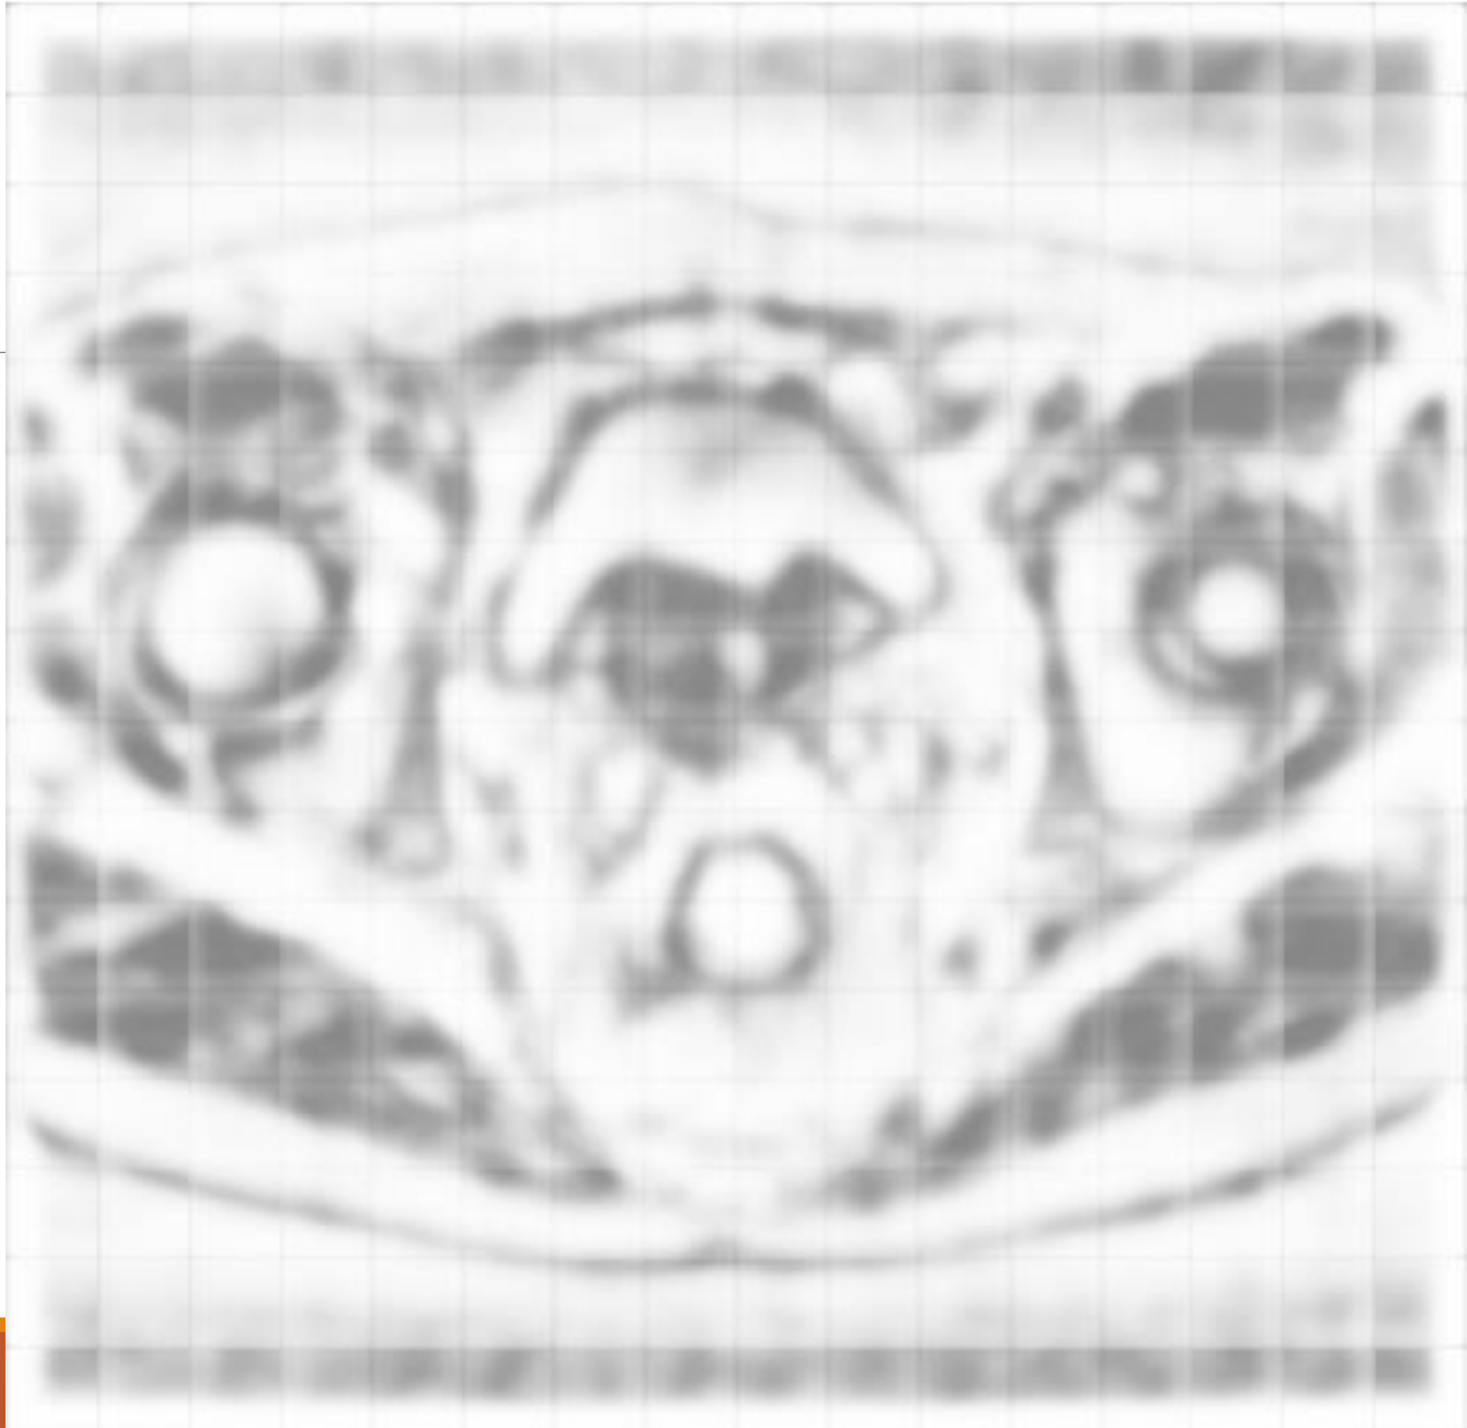

Brief Summary of the Last Episode

- MRI images of patients affected by colon cancer
- Manual segmentation done by expert clinicians
- First approach to automatic segmentation with Fully Convolutional Neural Networks (FCNN)
- **PROMISE – search for an inspiration**

Evaluation of prostate segmentation algorithms for MRI: The PROMISE12 challenge

Geert Litjens [a], Robert Toth [b], Wendy van de Ven [a], Caroline Hoeks [a], Sjoerd Kerkstra [a], Bram van Ginneken [a], Graham Vincent [e], Gwenael Guillard [e], Neil Birbeck [f], Jindang Zhang [f], Robin Strand [g], Filip Malmberg [g], Yangming Ou [h], Christos Davatzikos [h], Matthias Kirschner [i], Florian Jung [i], Jing Yuan [j], Wu Qiu [j] ... Anant Madabhushi [b]

⊞ Show more

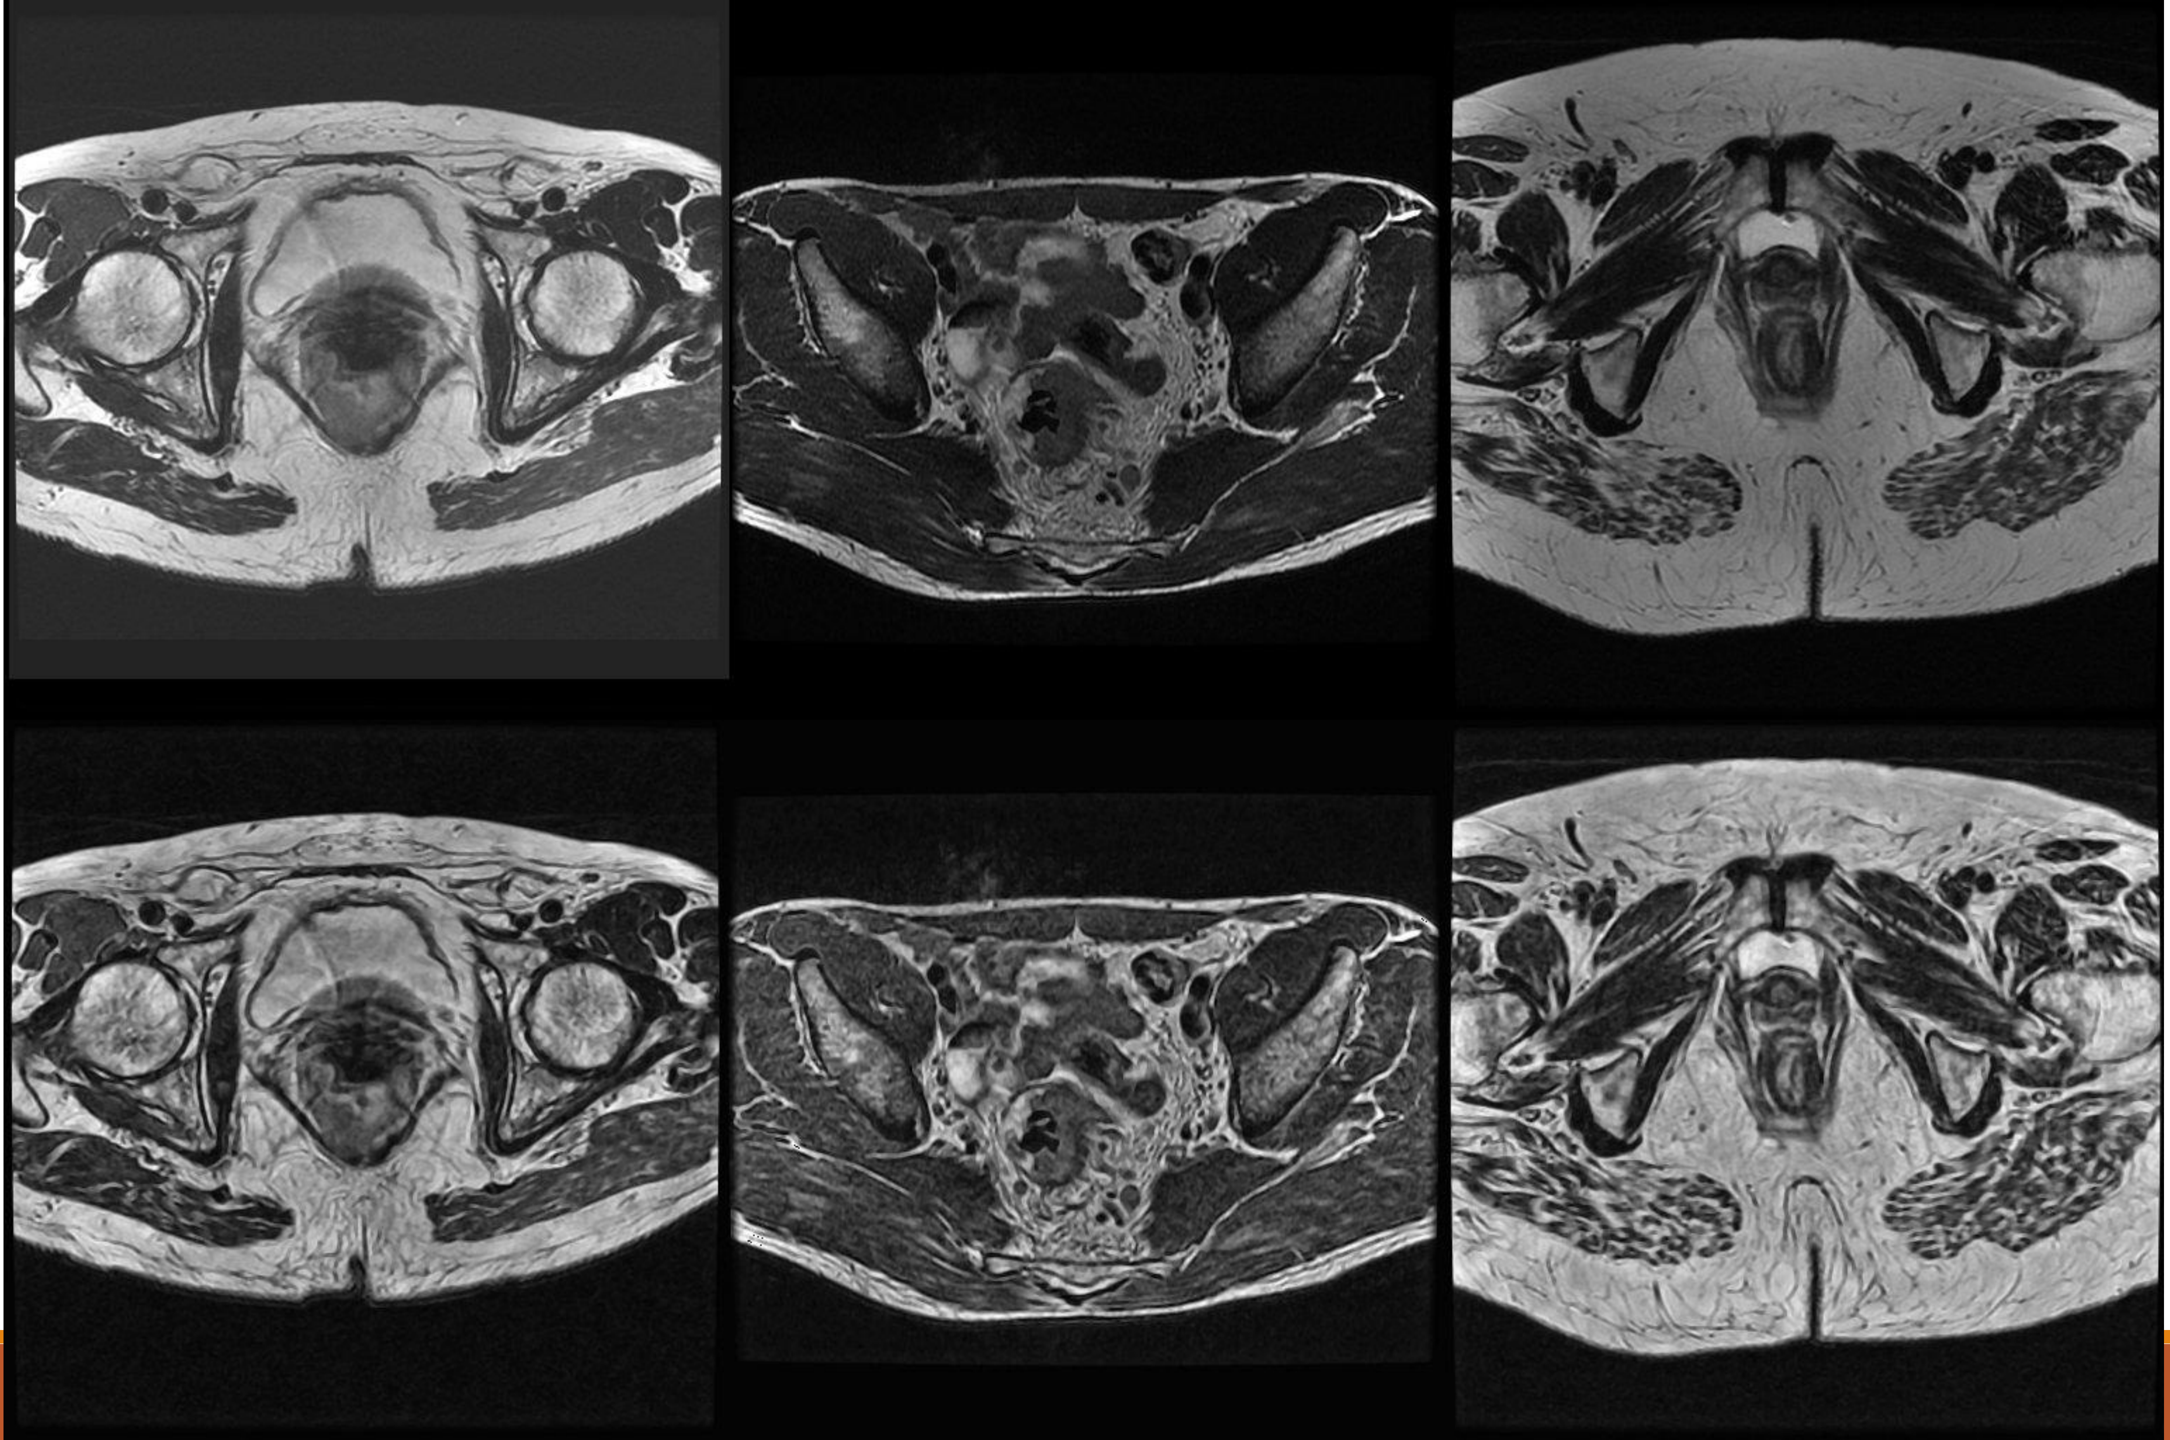

Contrast Limited Adaptive Histogram Equalization (CLAHE)

- Used to improve Local contrast and enhance definition of edges
- Histogram equalization for each section of the image
- Each pixel is transformed depending on the intensity of its neighbors (in a window of a given size)
- Contrast amplification is limited to avoid noise amplification

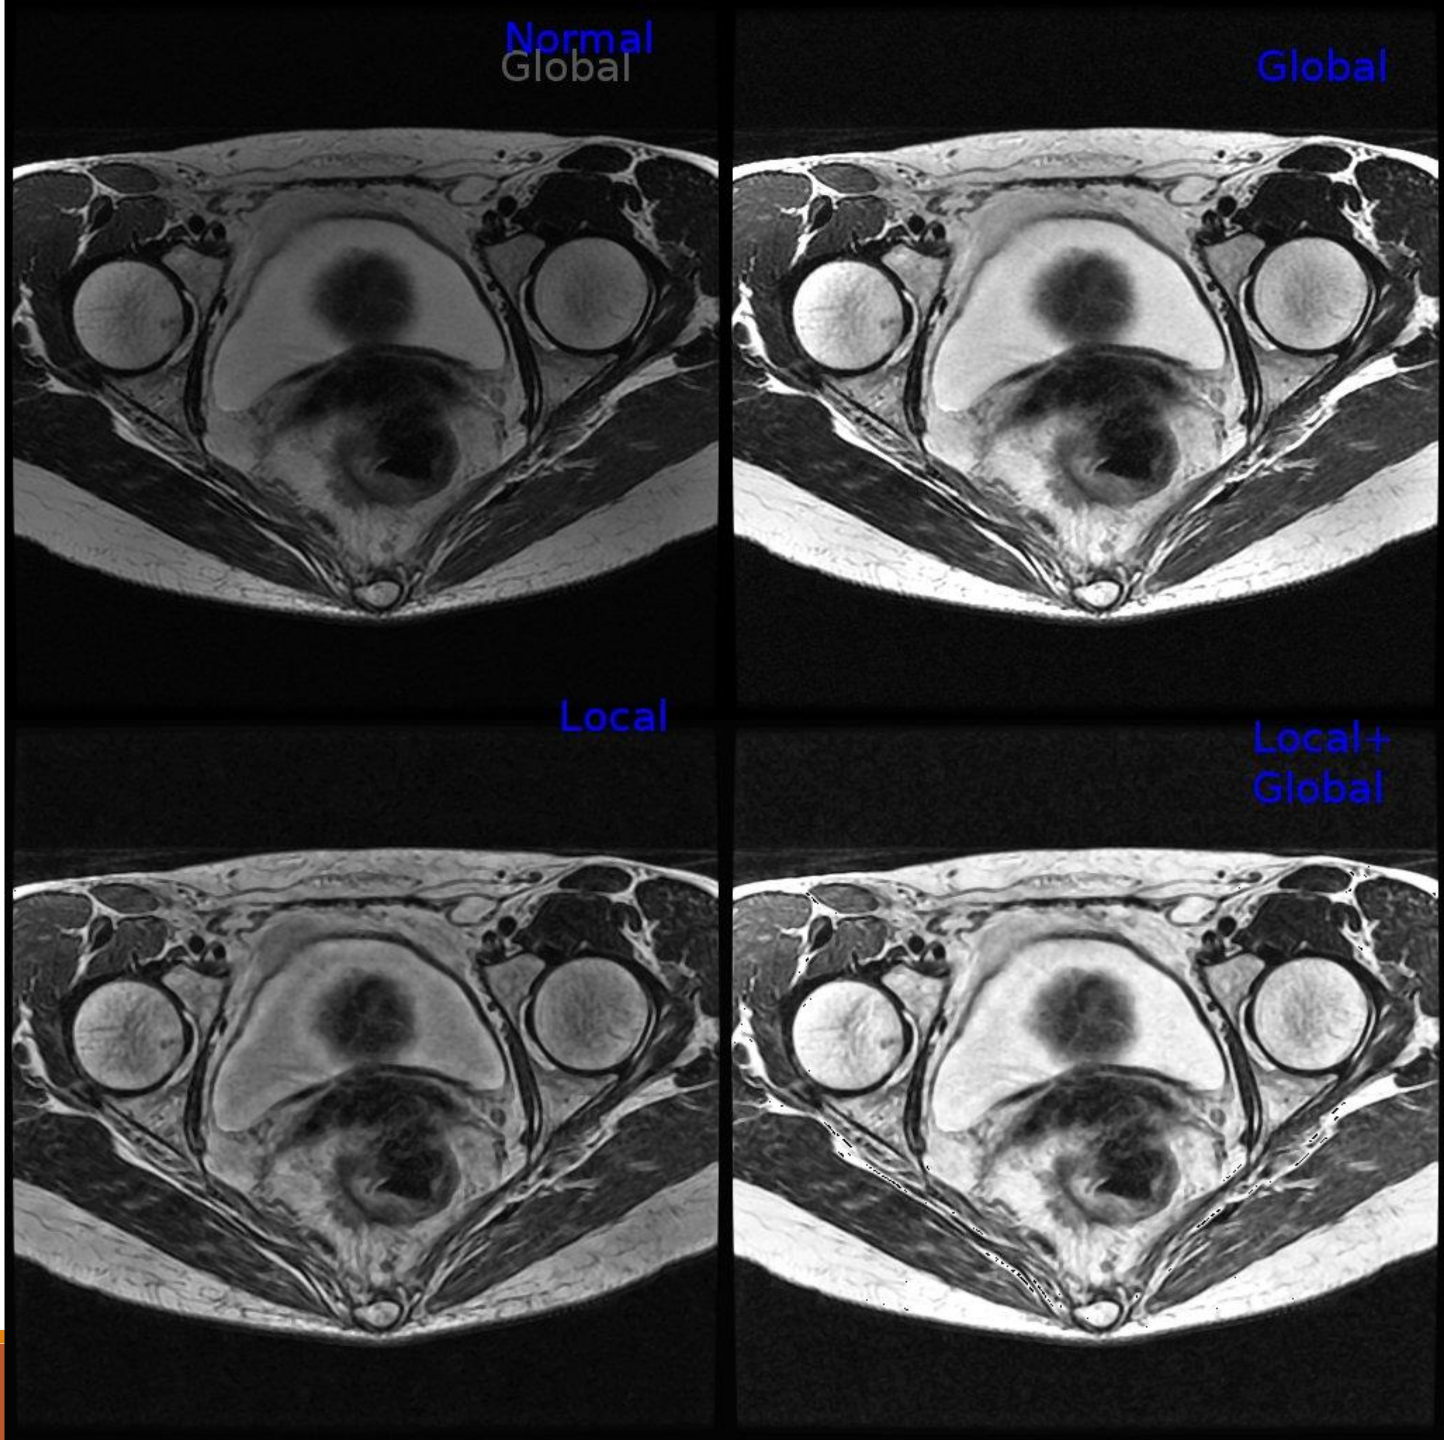


Normal
Global
Local
Global +
Local

Accuracy vs. DSC

1. The tumor is small. 2% of Total Area
2. An hard coded “no tumor” classifier has 98% accuracy
3. Different cost of Errors?

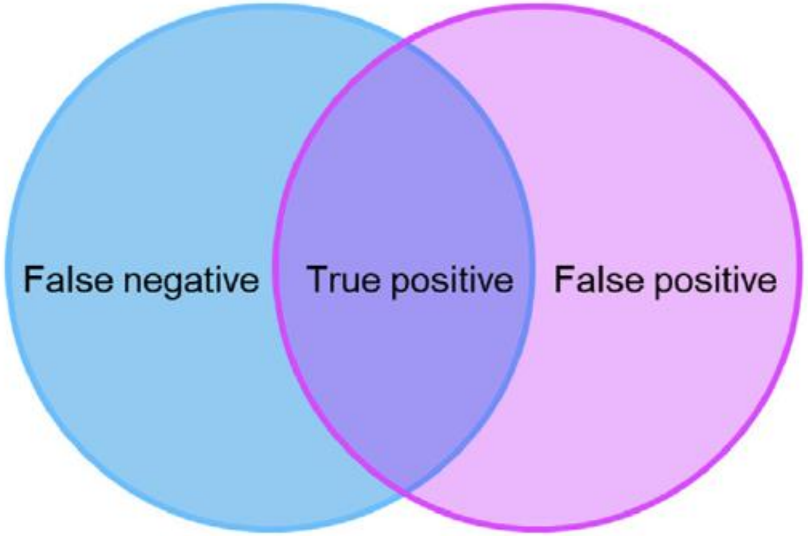

DICE Sorensen coefficient

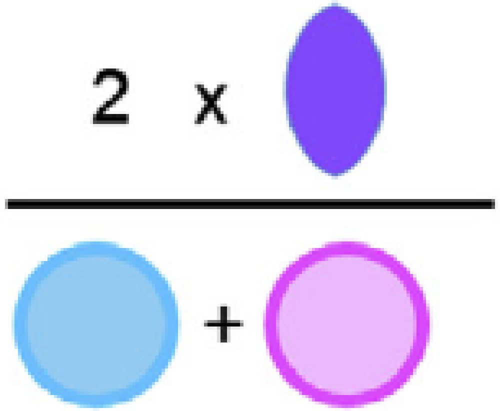

$$\text{DSC} = \frac{2TP}{2TP+FP+FN}$$

$$\text{Accuracy} = \frac{TP+TN}{TP+TN+FP+FN}$$

Low Score: separate Areas

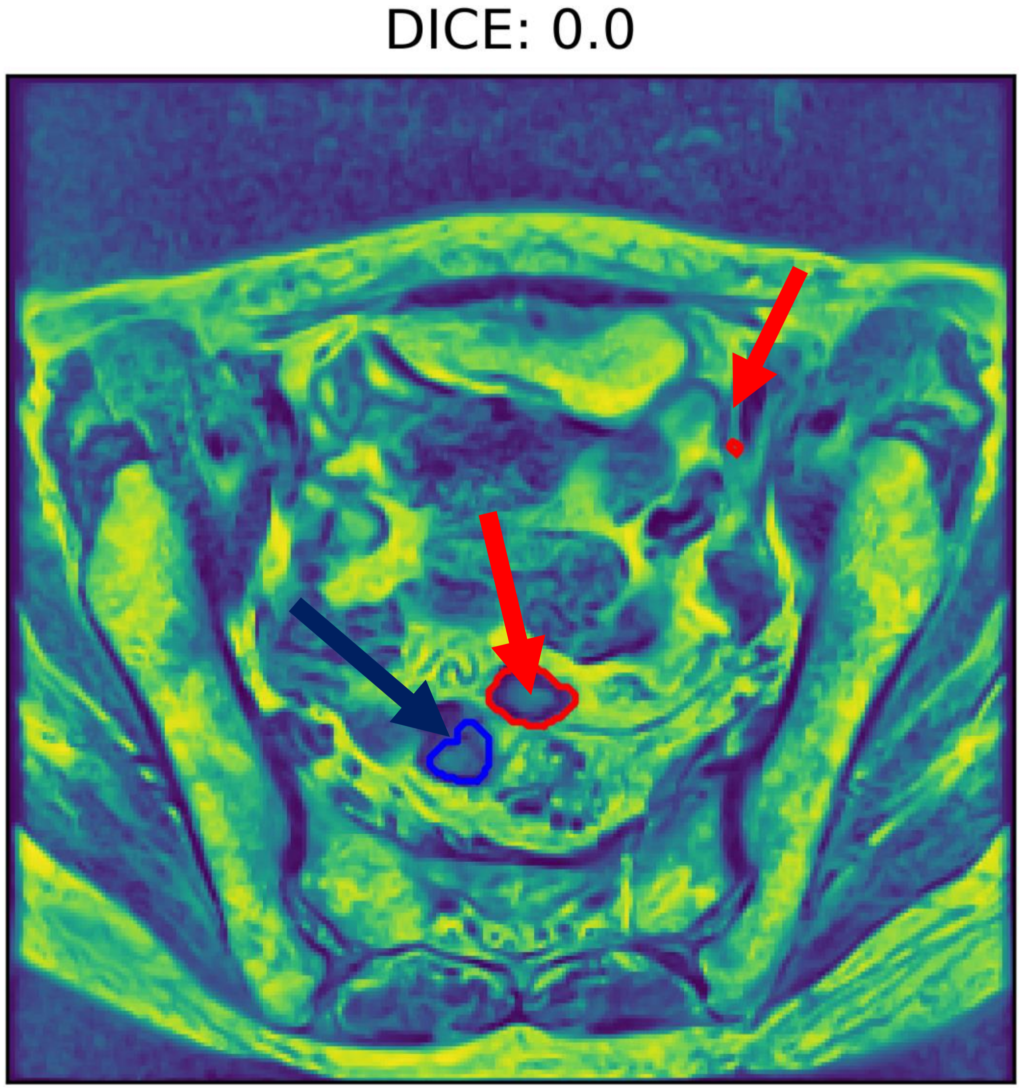

High Score: good intersection

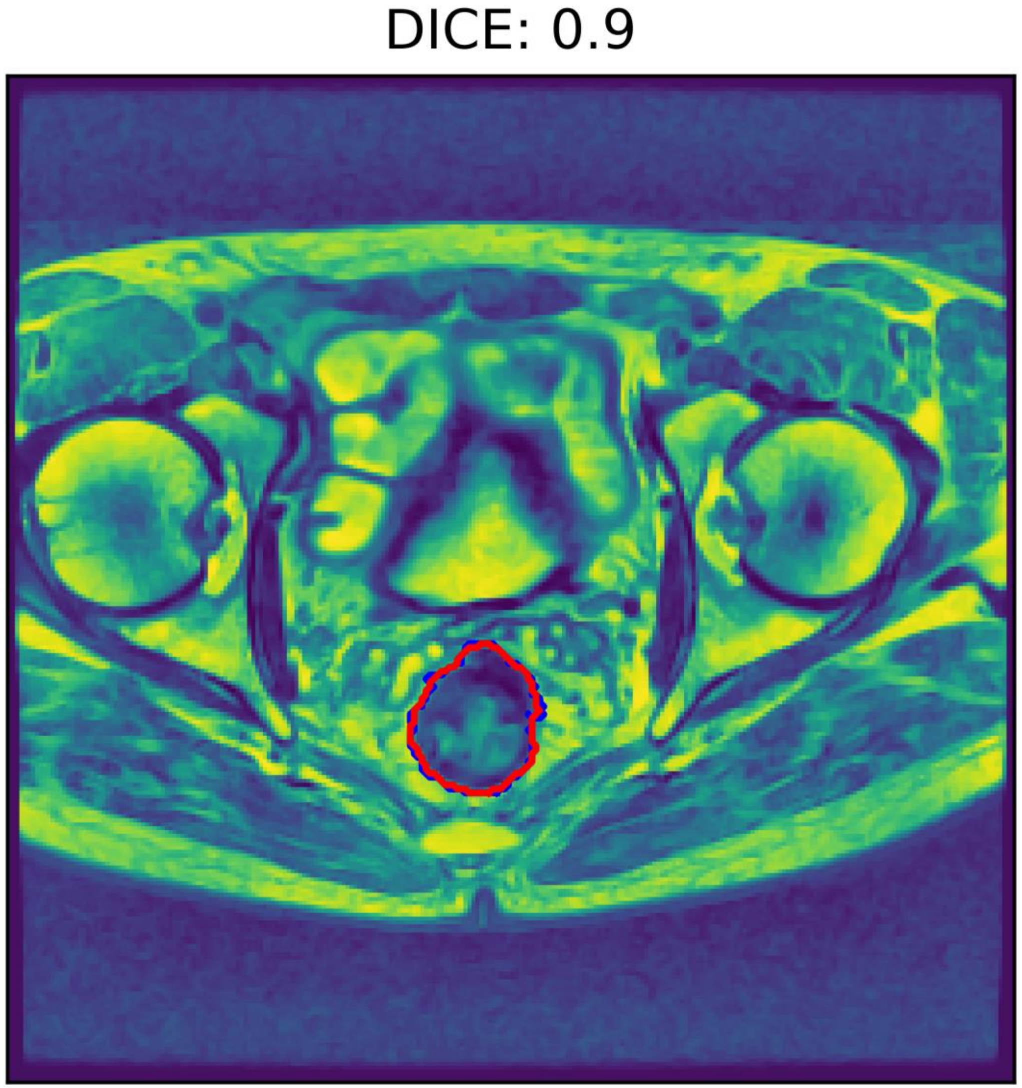

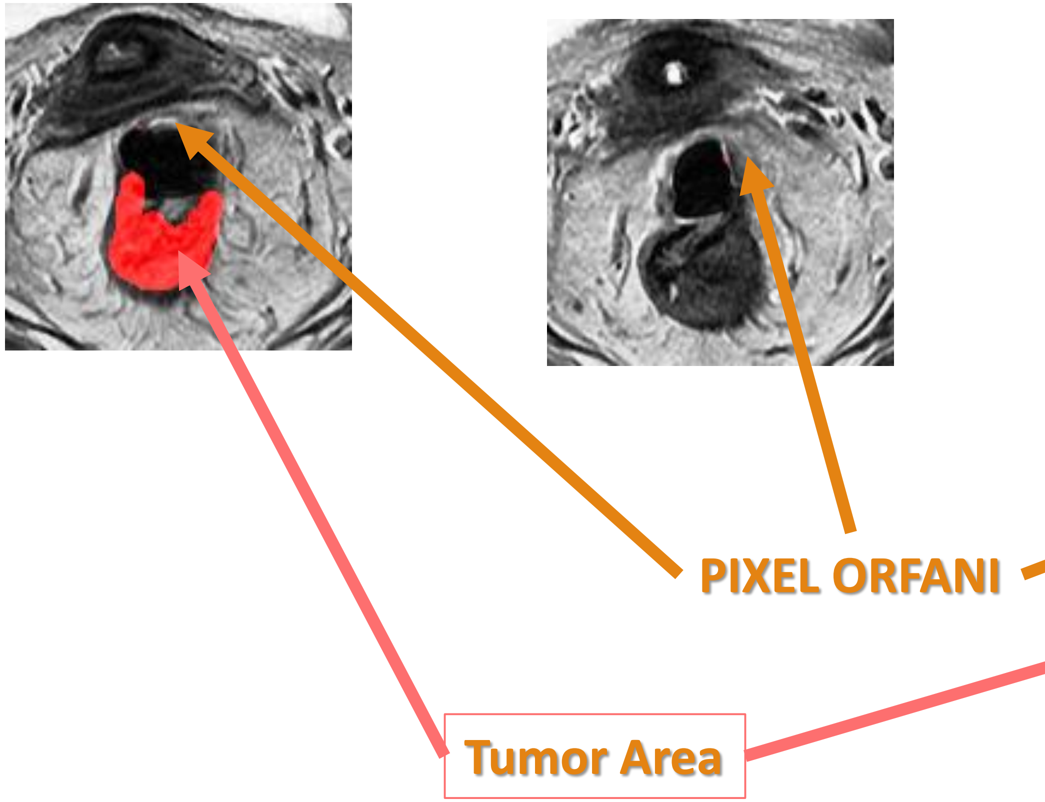

Lone pixels are difficult to treat when the measure
Is normalized to the tumor Area.
Slices with lone pixels are removed from BOTH the
training an TEST set.
This surely introduces a BIAS in the results

A question to the clinician: What are you exactly marking in the segmentation?

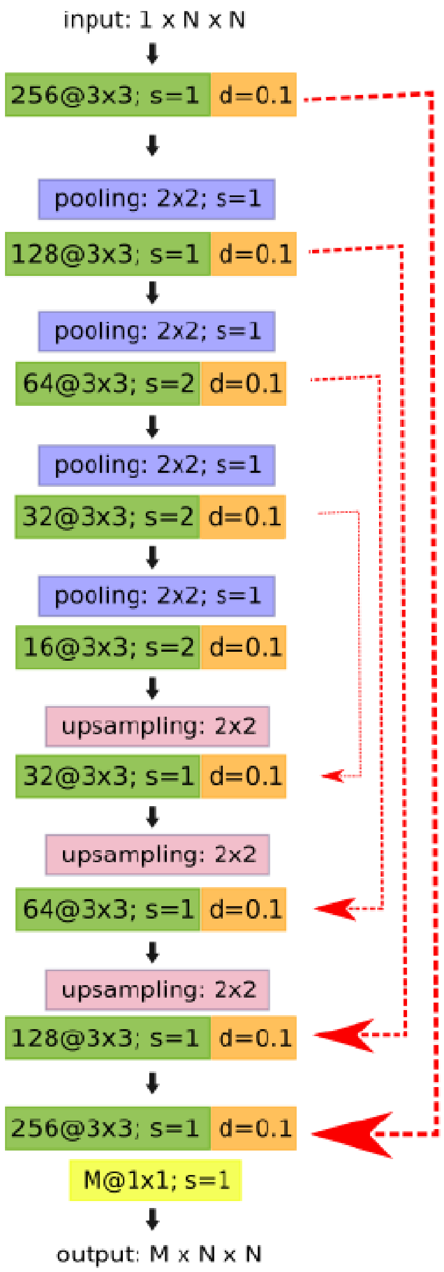

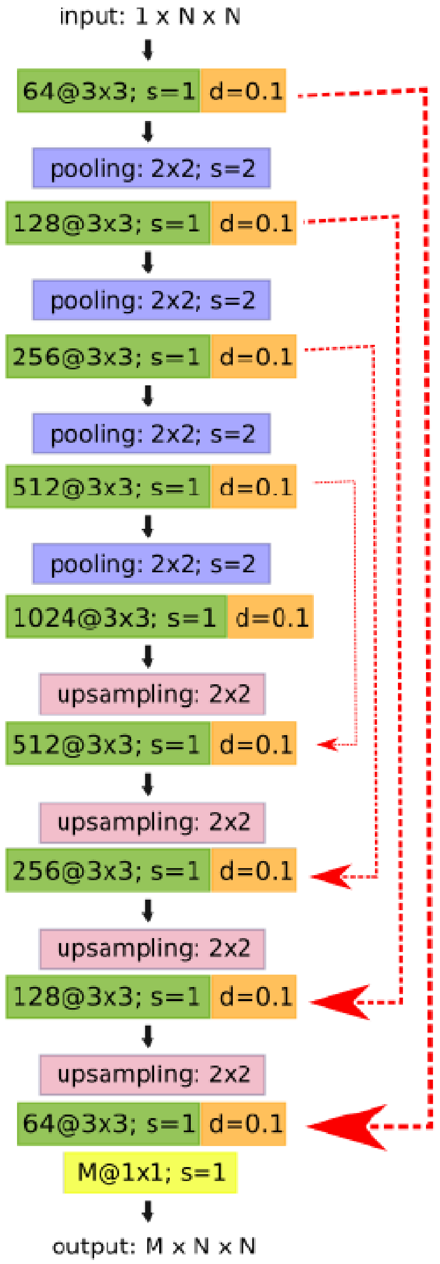

InvertedNet vs. All Drop-Out

	InvertedNet	All Drop-Out
Parameters	1.4 M	31 M
AUC	0.93	0.90
Mean Dsc	0.82	0.83
Mean Volumetric Dsc	0.80	0.82
Volumetric Dsc std	0.1	0.09

Inverted Net (1.4 M parameters)

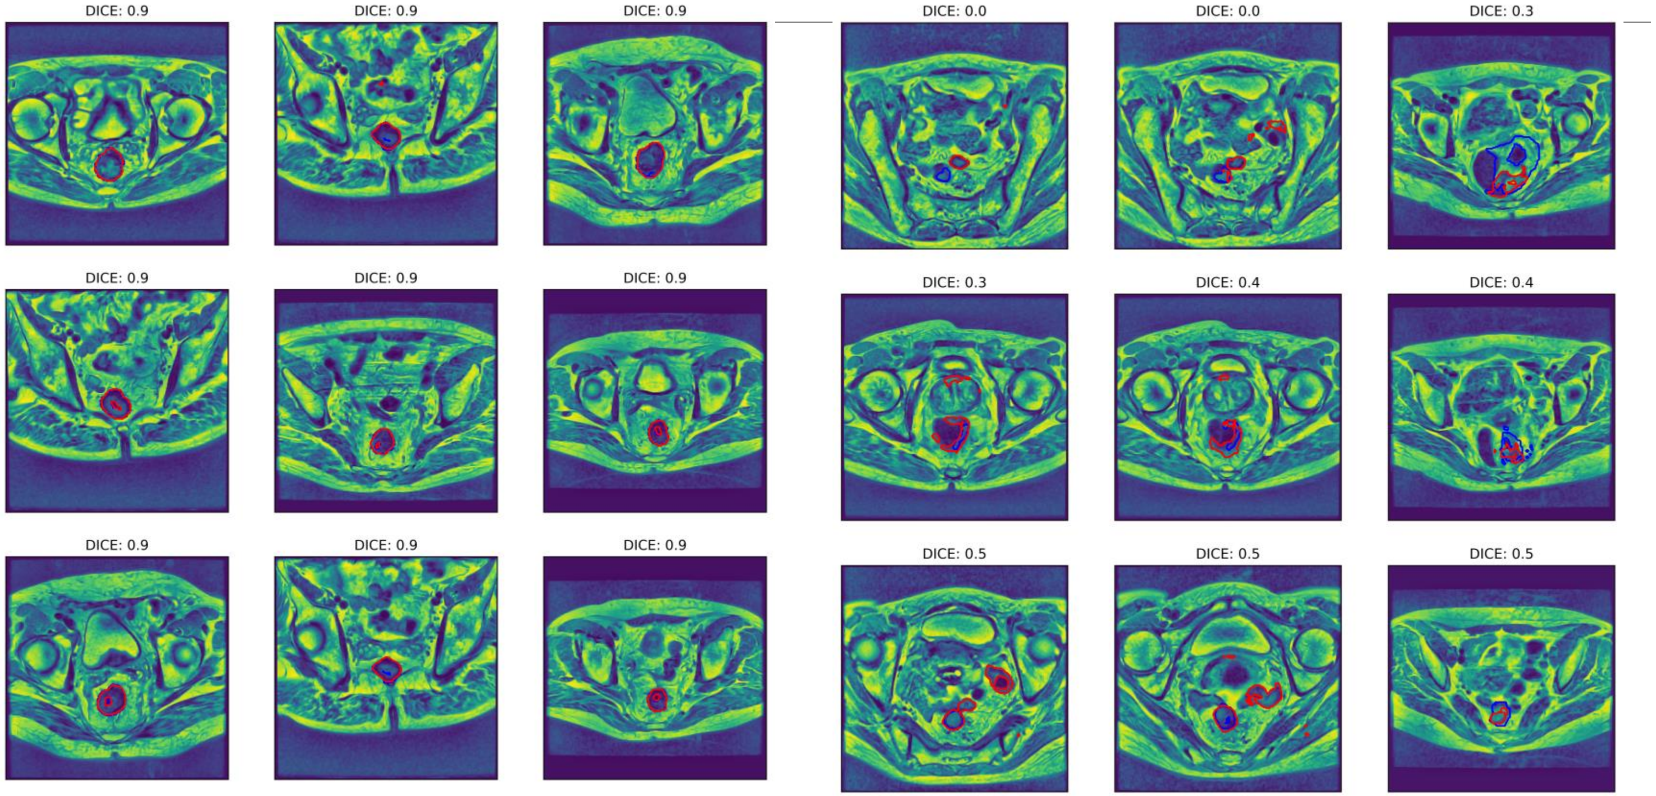

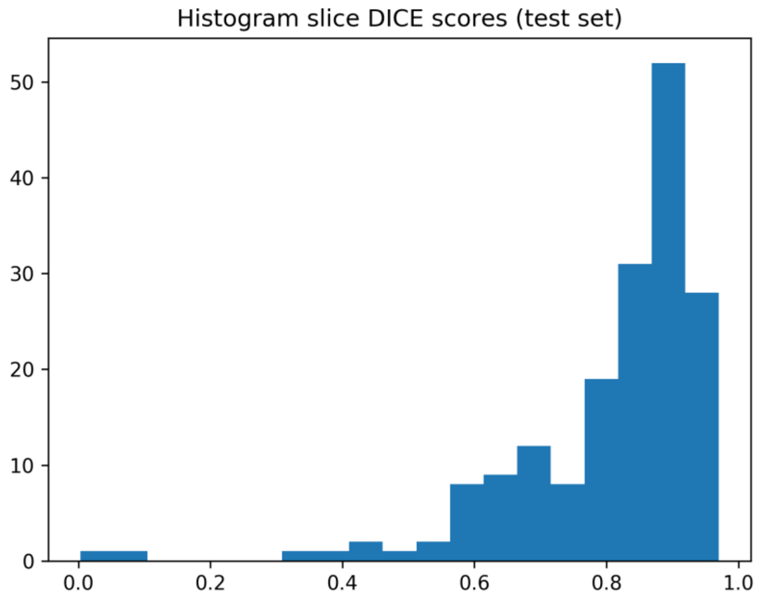

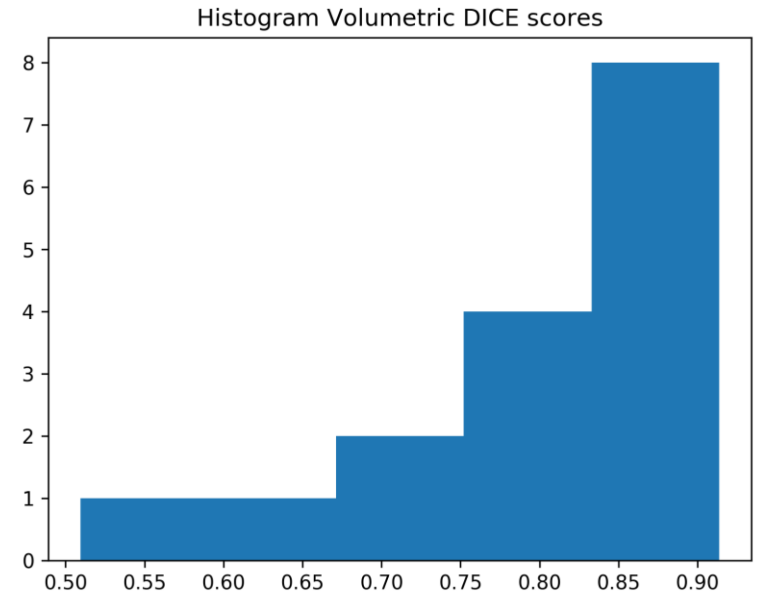

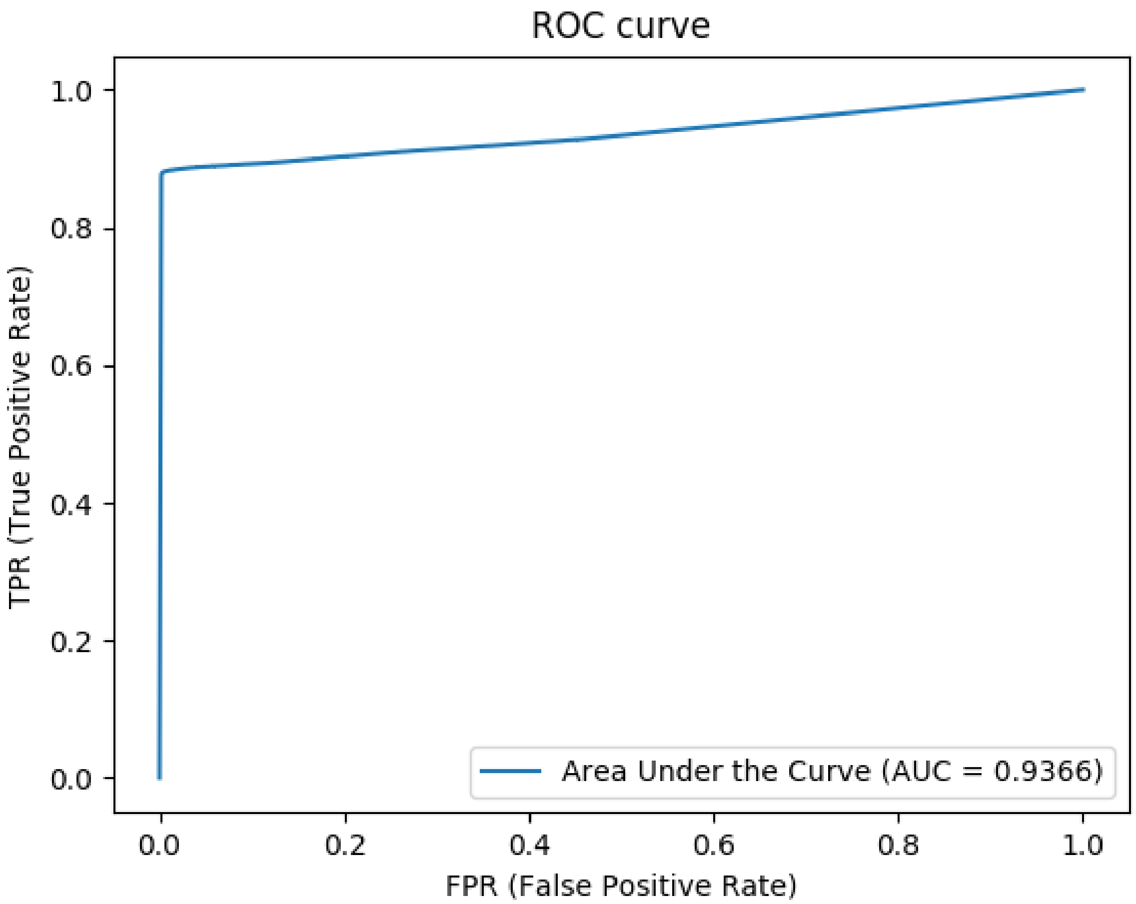
ROC curve
1.0
0.8
0.6
0.4
0.2
0.0
TPR (True Positive Rate)
Area Under the Curve (AUC = 0.9366)
0.0
0.2
0.4
0.6
0.8
1.0
FPR (False Positive Rate)

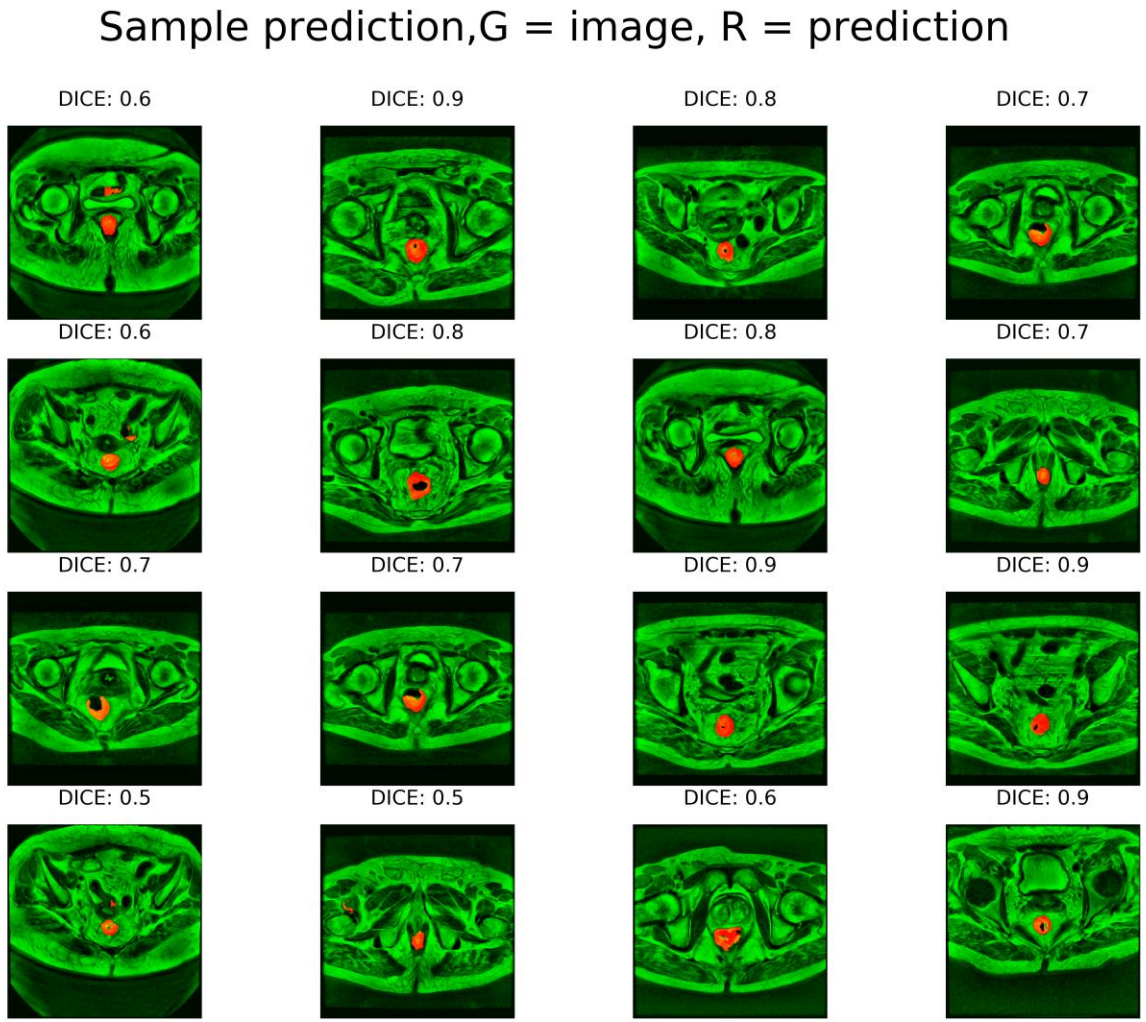
Sample prediction,G = image, R = prediction
DICE: 0.6
DICE: 0.9
DICE: 0.8
DICE: 0.7
DICE: 0.6
DICE: 0.8
DICE: 0.8
DICE: 0.7
DICE: 0.7
DICE: 0.7
DICE: 0.9
DICE: 0.9
DICE: 0.5
DICE: 0.5
DICE: 0.6
DICE: 0.9

All Drop-Out (30 M p.)

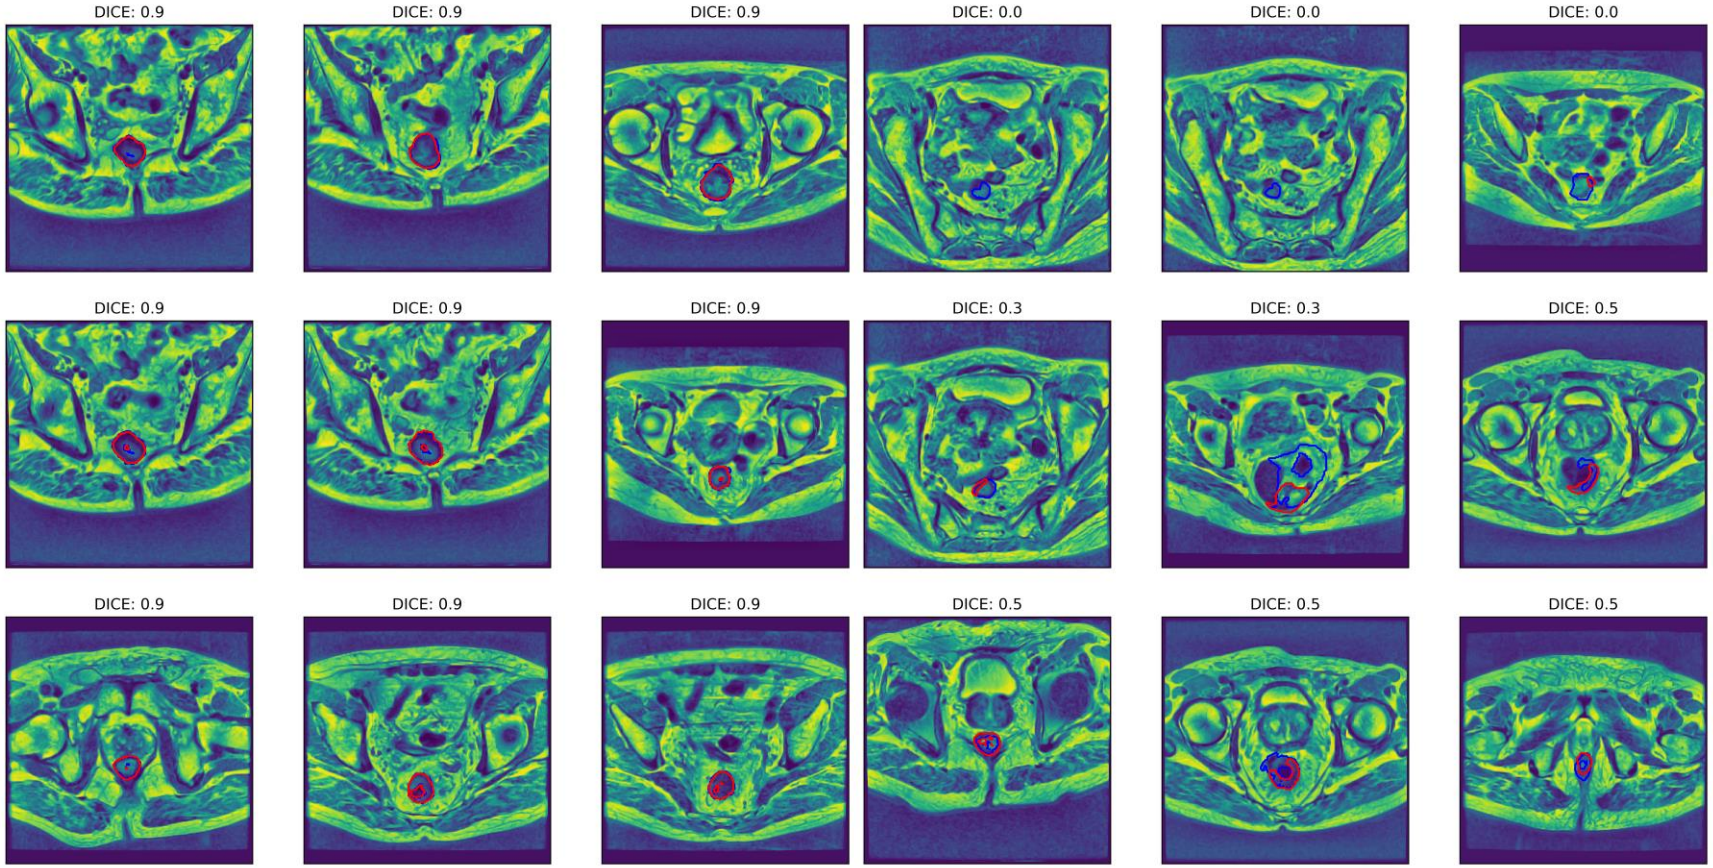

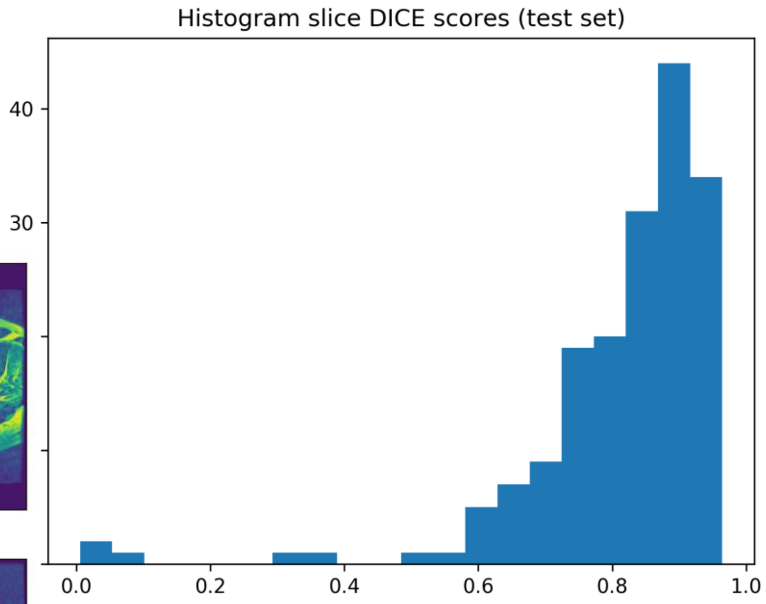

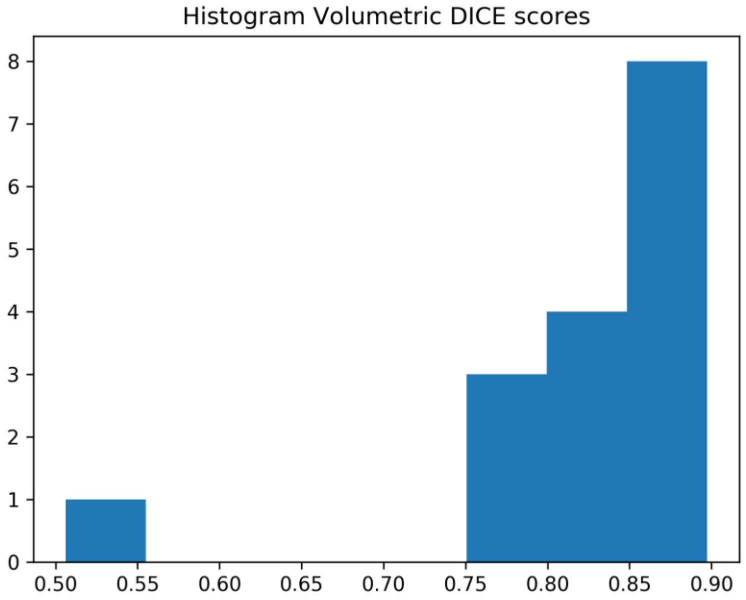

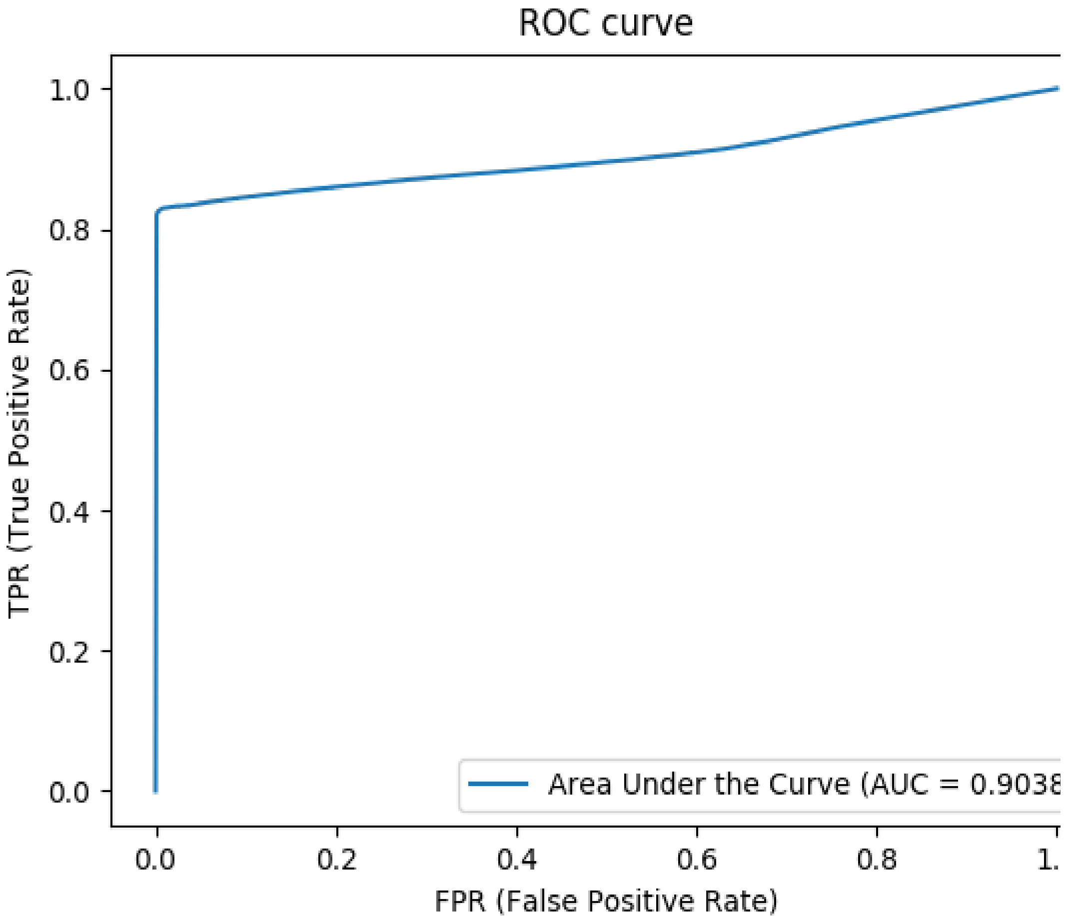
ROC curve
1.0
0.8
0.6
0.4
0.2
0.0
TPR (True Positive Rate)
0.0
0.2
0.4
0.6
0.8
1.
FPR (False Positive Rate)
Area Under the Curve (AUC = 0.9038

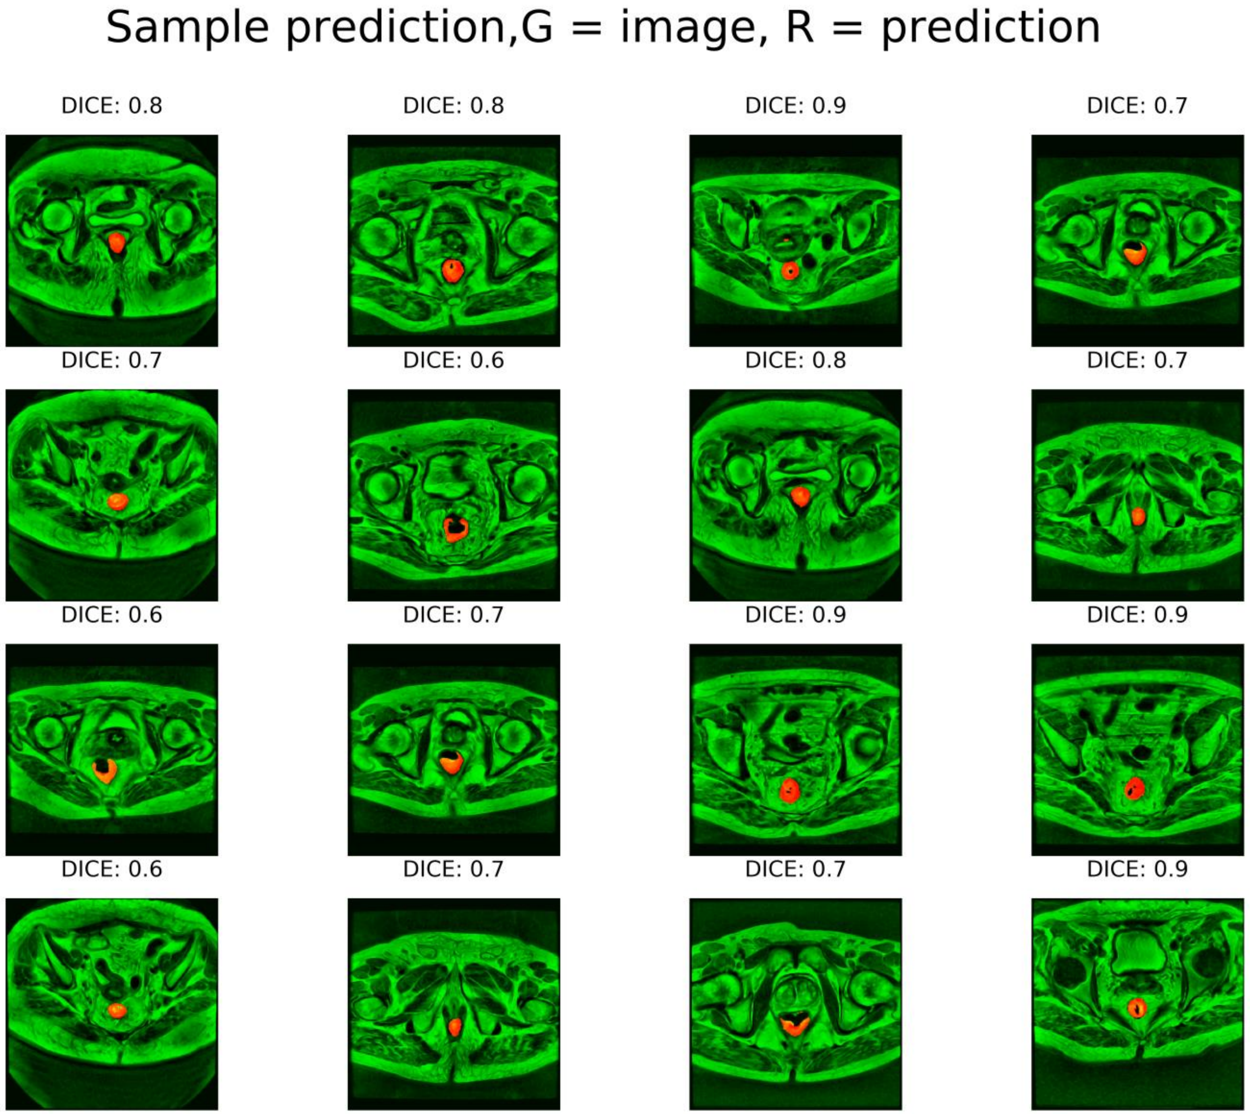
Sample prediction,G = image, R = prediction
DICE: 0.8
DICE: 0.8
DICE: 0.9
DICE: 0.7
DICE: 0.7
DICE: 0.6
DICE: 0.8
DICE: 0.7
DICE: 0.6
DICE: 0.7
DICE: 0.9
DICE: 0.9
DICE: 0.6
DICE: 0.7
DICE: 0.7
DICE: 0.9

Test on Healthy tissue

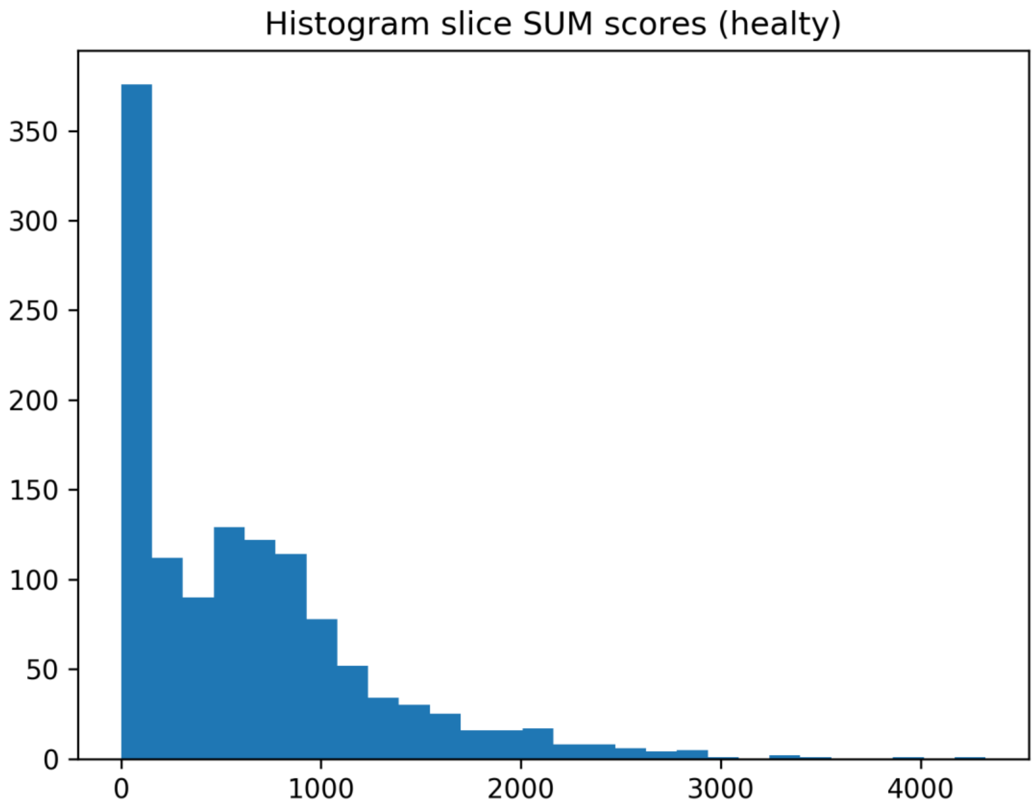

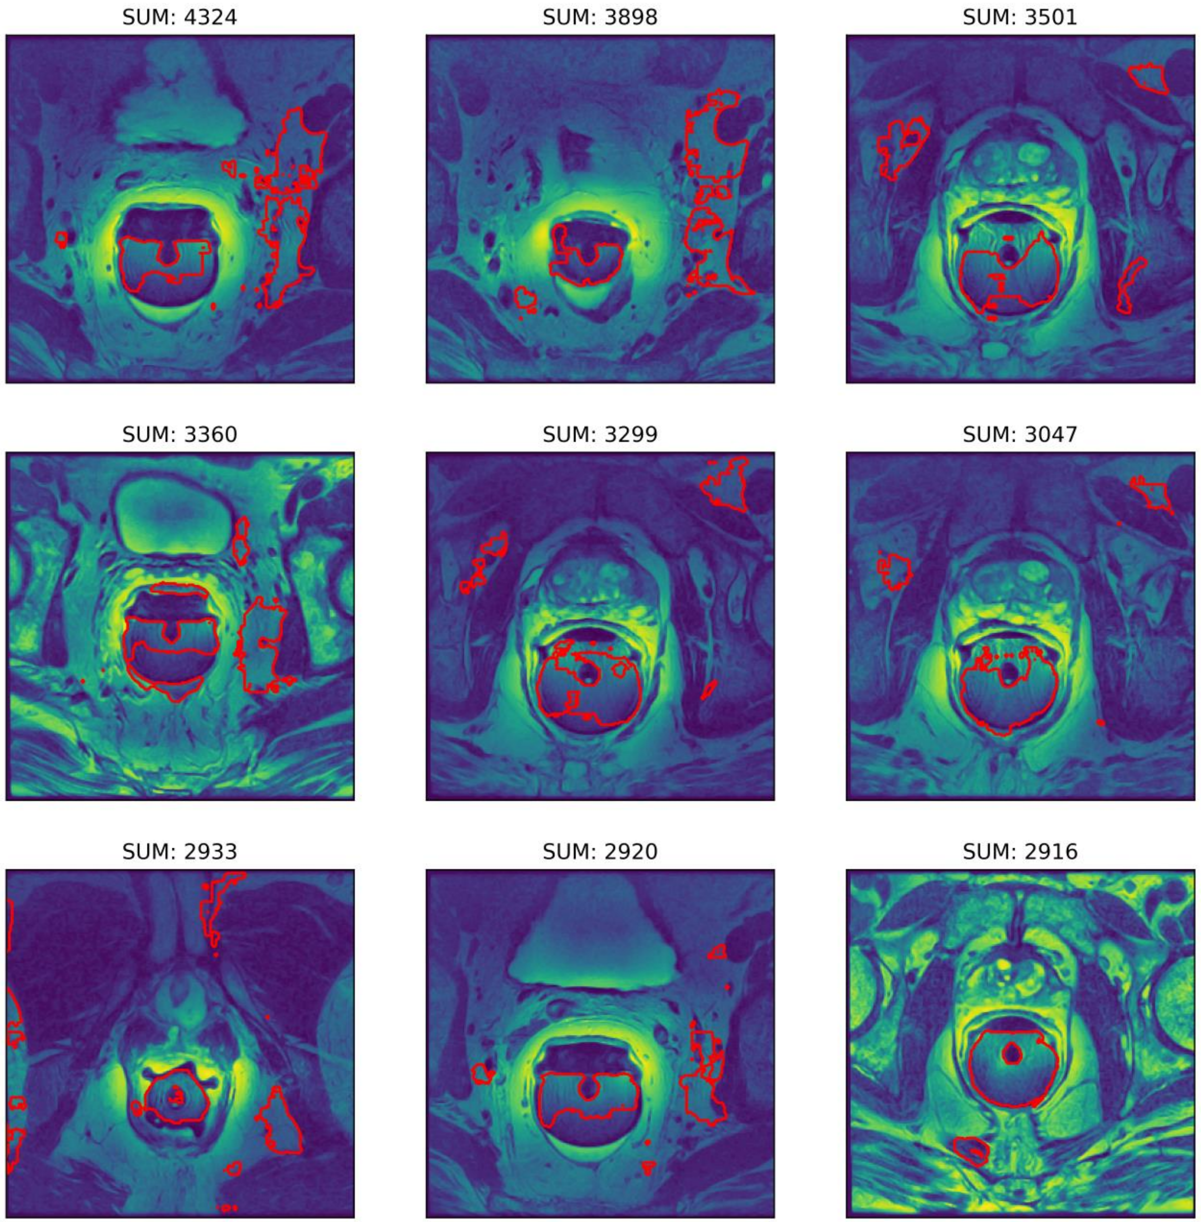

What the Net is Learning?

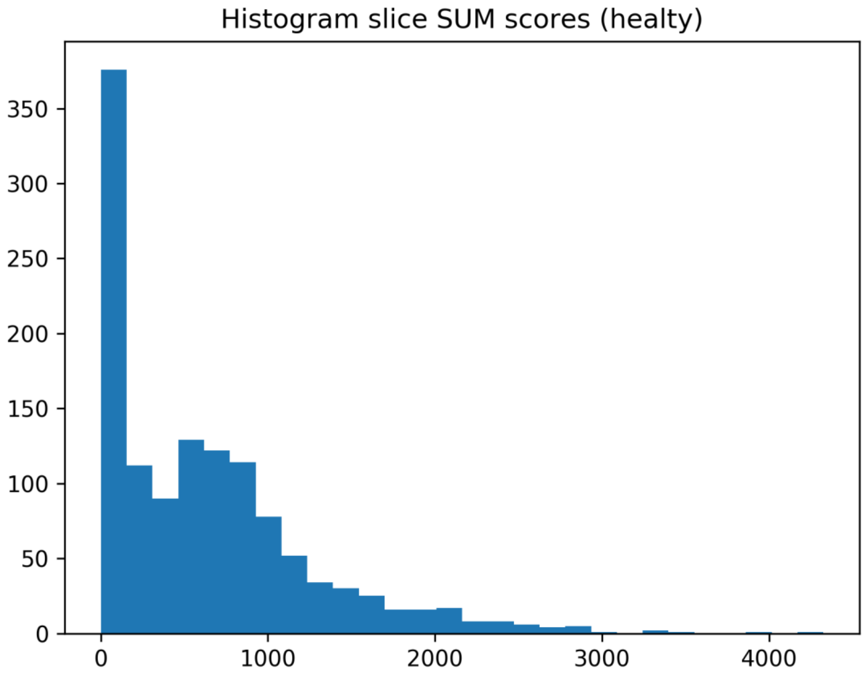

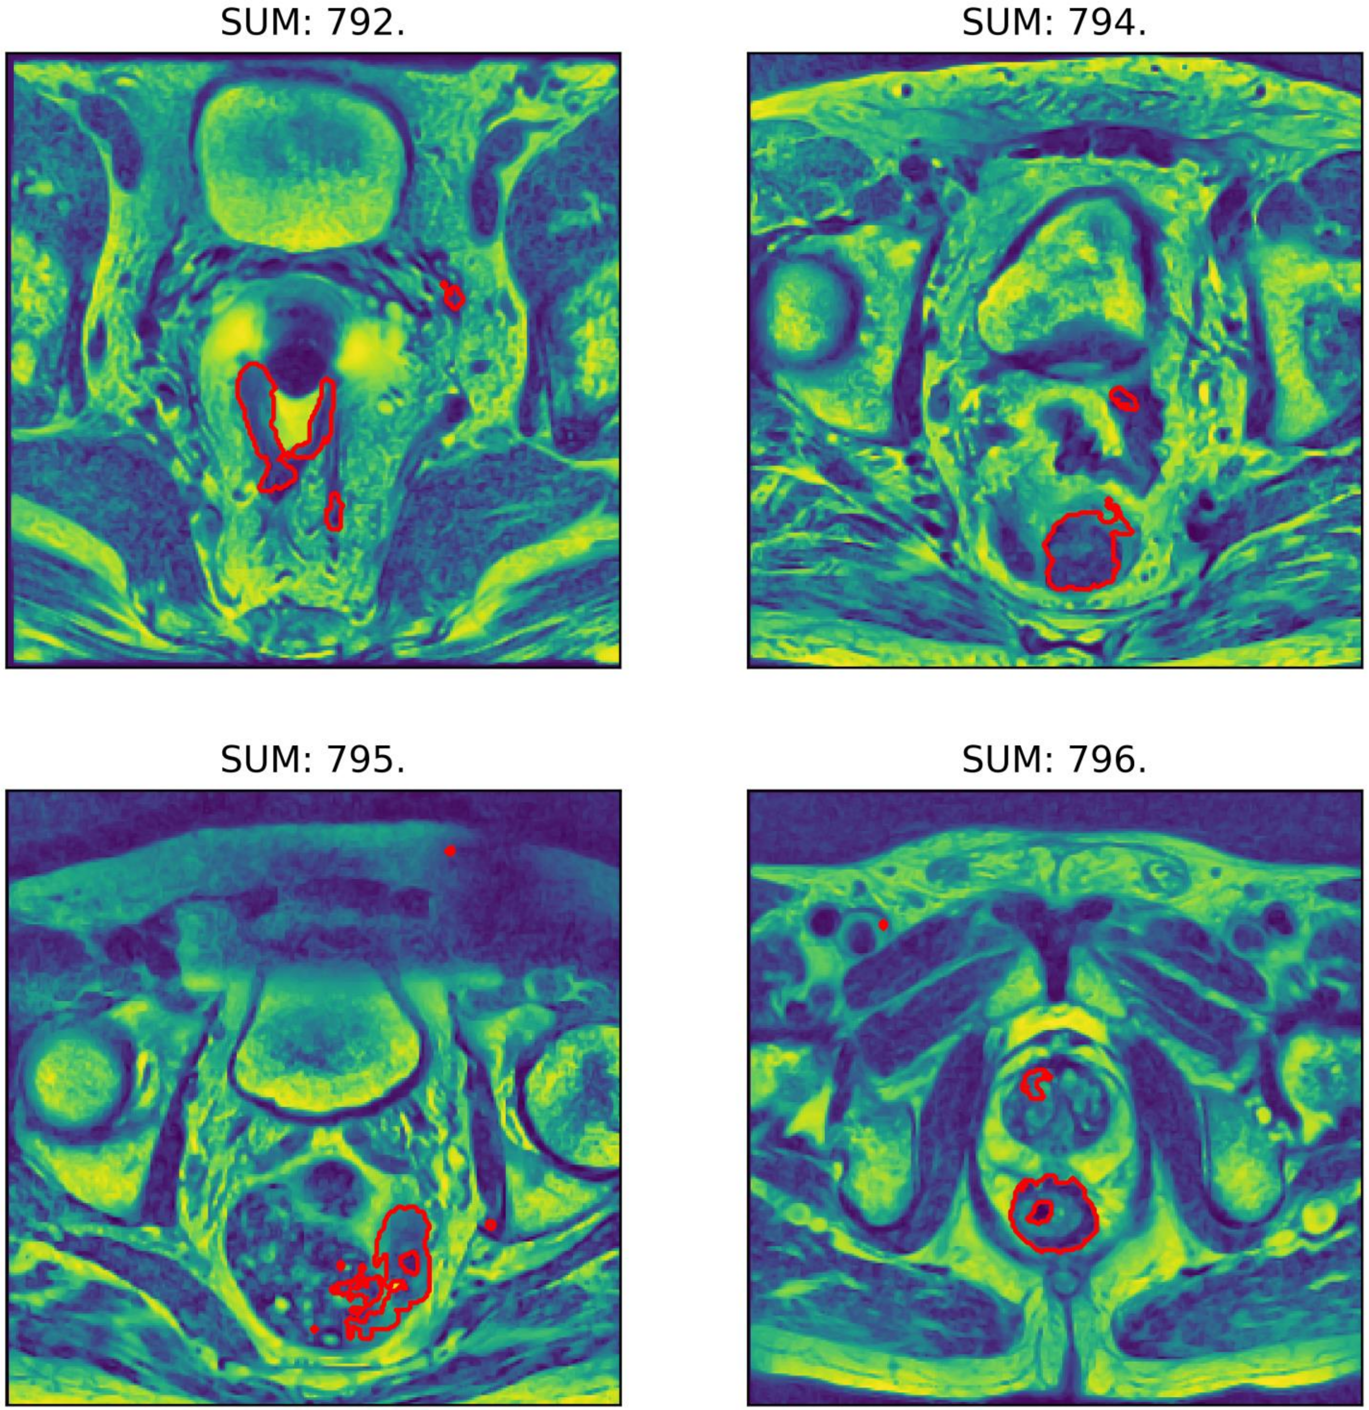

OPEN

Deep Learning for Fully-Automated Localization and Segmentation of Rectal Cancer on Multiparametric MR

Stefano Trebeschi[1,2], Joost J. M. van Griethuysen[1,2], Doenja M. J. Lambregts[1], Max J. Lahaye[1], Chintan Parmar[3], Frans C. H. Bakers[4], Nicky H. G. M. Peters[5], Regina G. H. Beets-Tan[1,2] & Hugo J. W. L. Aerts[1,3]

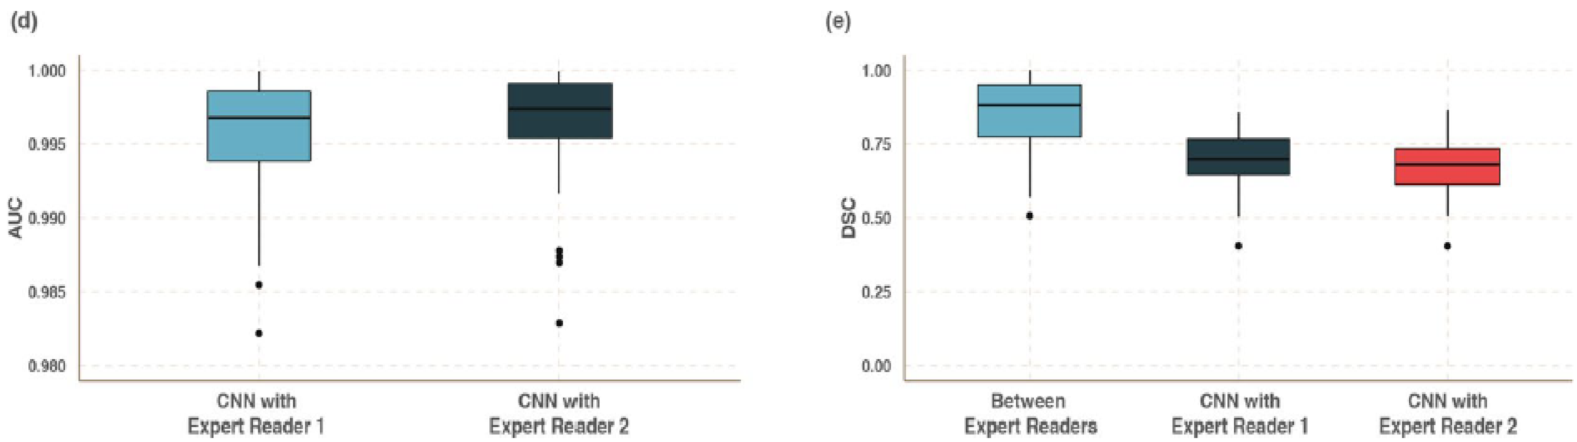

Figure 3. CNN Training and Validation. *Performance of the CNN on the discovery dataset:* (**a**) accuracy, (**b**) cross entropy and (**c**) improvement (Δ cross entropy). Improvement shown in panel (c), in computed on the test set only, preventing the model from overfitting. *Performance of the CNN on the validation dataset:* (**d**) the Area under the ROC curve (AUC) of the probability map with respect to the reader segmentation, and (**e**) Dice Similarity Coefficient (DSC) of the generated segmentations.

	Inverted Net	All Drop-Out
Parameters	1.4 M	31 M
AUC	0.93	0.90
Mean Dsc	0.82	0.83
Mean Volumetric Dsc	0.80	0.82
Volumetric Dsc std	0.1	0.09

Conclusions

We have a robust method to segment anatomical structures in the pelvis area.

This is our BASELINE model. Any development should perform better with comparable computational costs.

At the moment TRANSFER LEARNING is the best candidate to improve our results. This is what we will try next.

New data is always welcome. Improving the number of examples to address biological variance is the best strategy.

- Standard U-net architecture

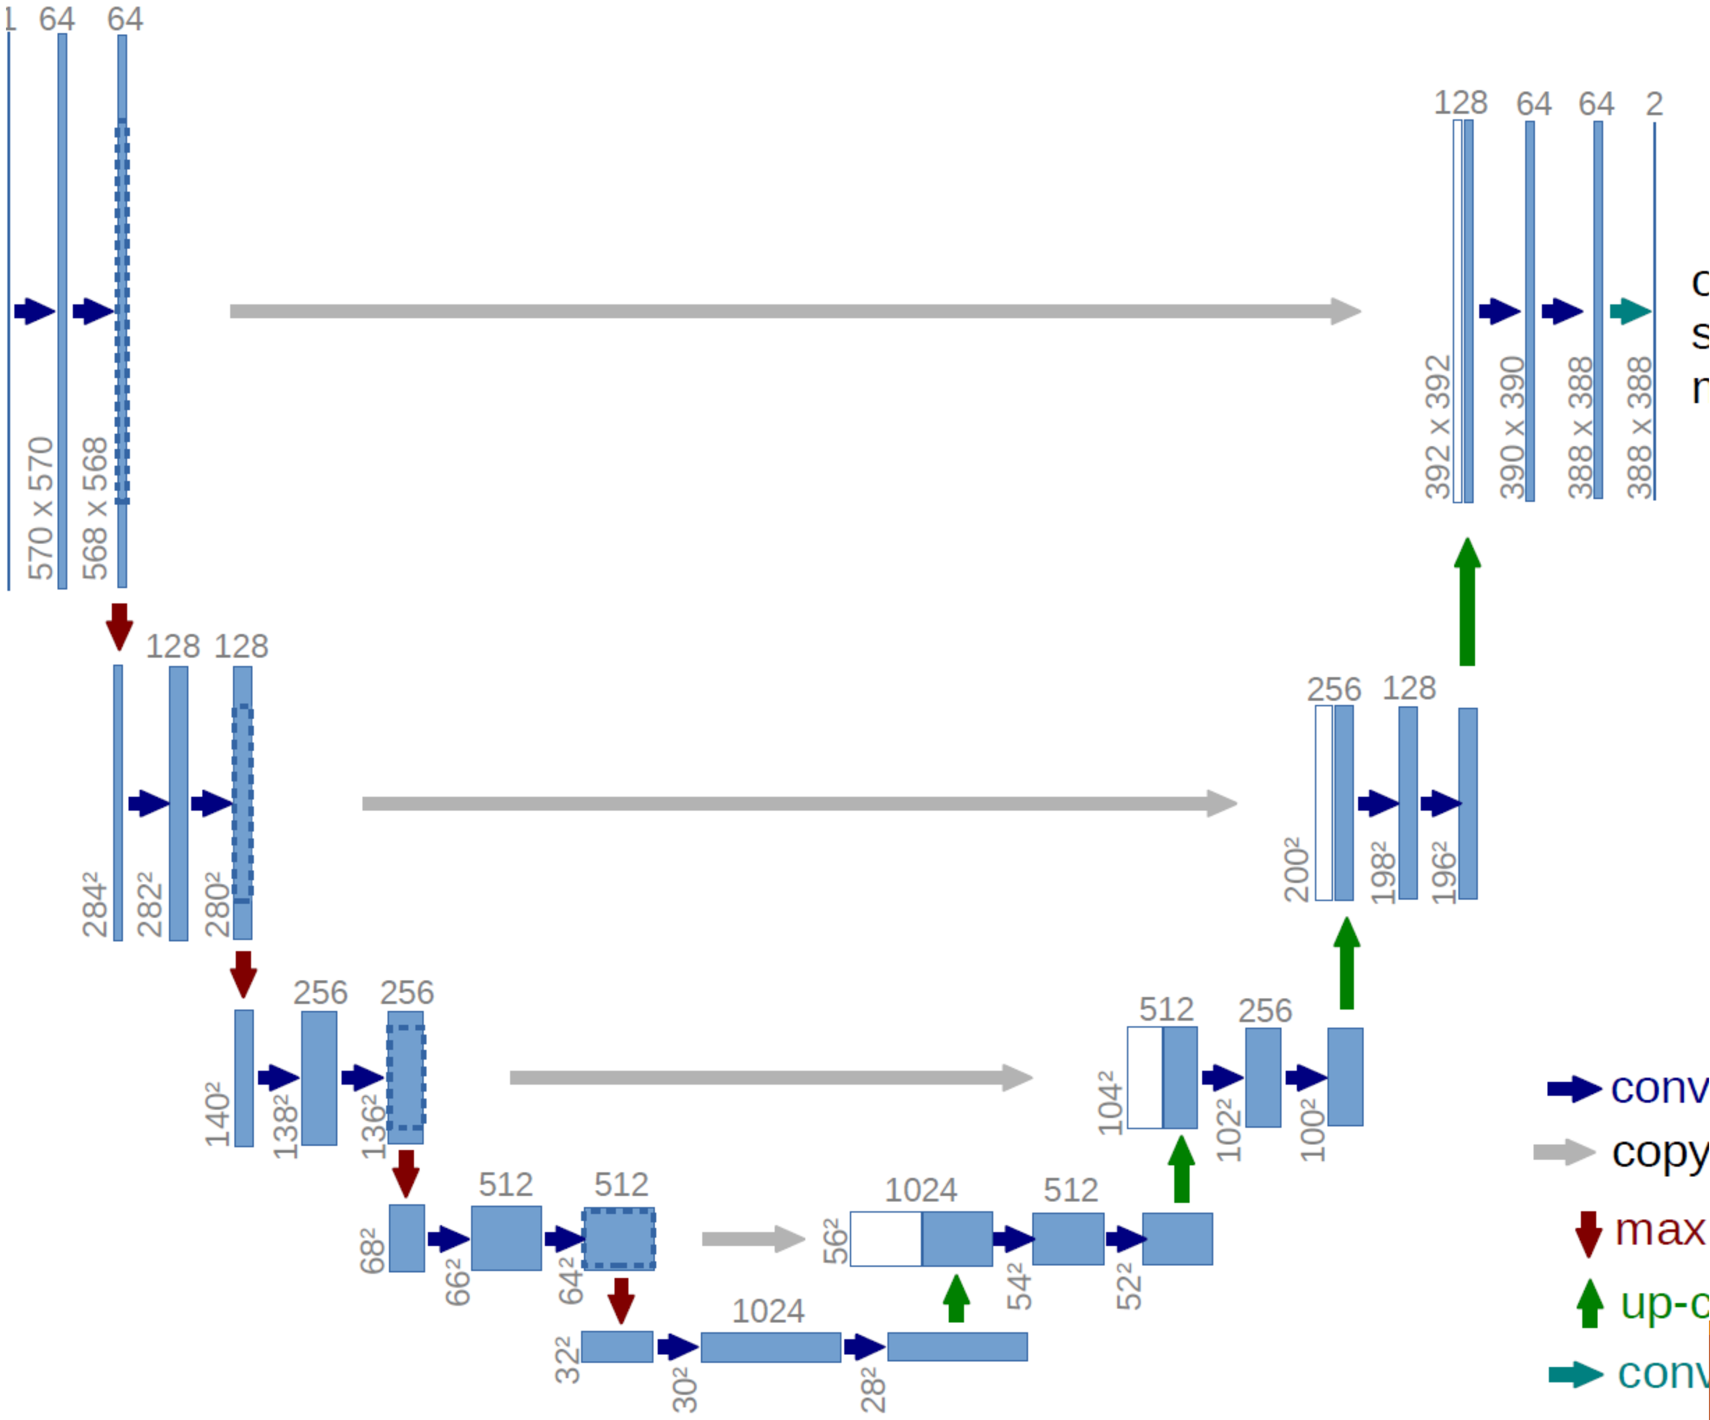

- Bad errors
- Lots of false positives
- Incorrect use of antomy (colon usually is in the center)
- Lack of generalization
- Rule probably based on grey intensity (very bad). Like a Pass-Band filter

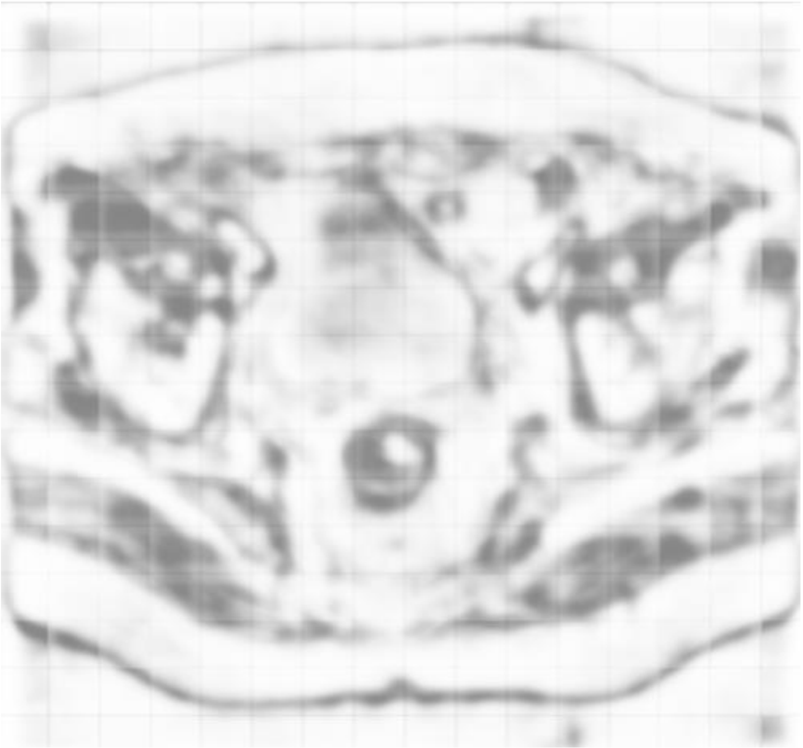

Details – what is the same

Similar to what I was alredy doing

- FCNN with skip connection (U-net like) to use information at different level of detail
- Data augmentation
- Early stopping to avoid overfitting
- Full image training

Details – What is new

Data augmentation – Elastic deformation

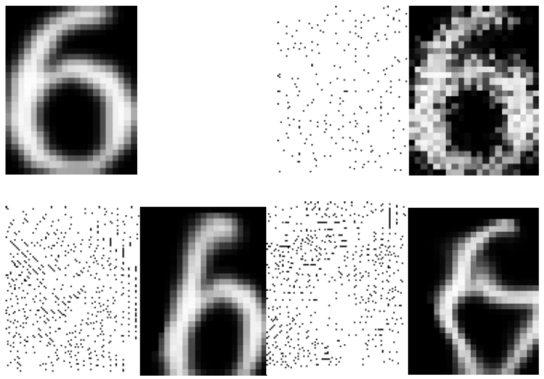

DSC as loss function

$$DSC(y,\hat{y}) = \frac{2\sum_i^N y_i \cdot \hat{y}_i}{\sum_i^N y_i + \sum_i^N \hat{y}_i}.$$

- Many more free parameters to fit:
- $5 \cdot 10^7$ (all of them necessary?)
- Dropout at each convolutional Block
- No split and stich approach! (not needed)
- Down-sampling of the image

Altamente scopiazzabile

Results on validation set	Big Net	Small Net
Total params	55,114,432	3,284,442
Accuracy	0.867	0.880
Mean volumetric DSC	0.873	0.623
Median volumetric DSC	0.862	0.722
Std volumetric DSC	0.037	0.140

Results on training set	Big Net	Small Net
Total params	55,114,432	3,284,442
Accuracy	0.955	0.965
Mean volumetric DSC	0.944	0.726
Median volumetric DSC	0.945	0.735
Std volumetric DSC	0.013	0.136

Training times	Big Net	Small Net
Training epochs	20 (1.5 hours per epoch)	25 (0.6 hours per epoch)

Results small net – Error types

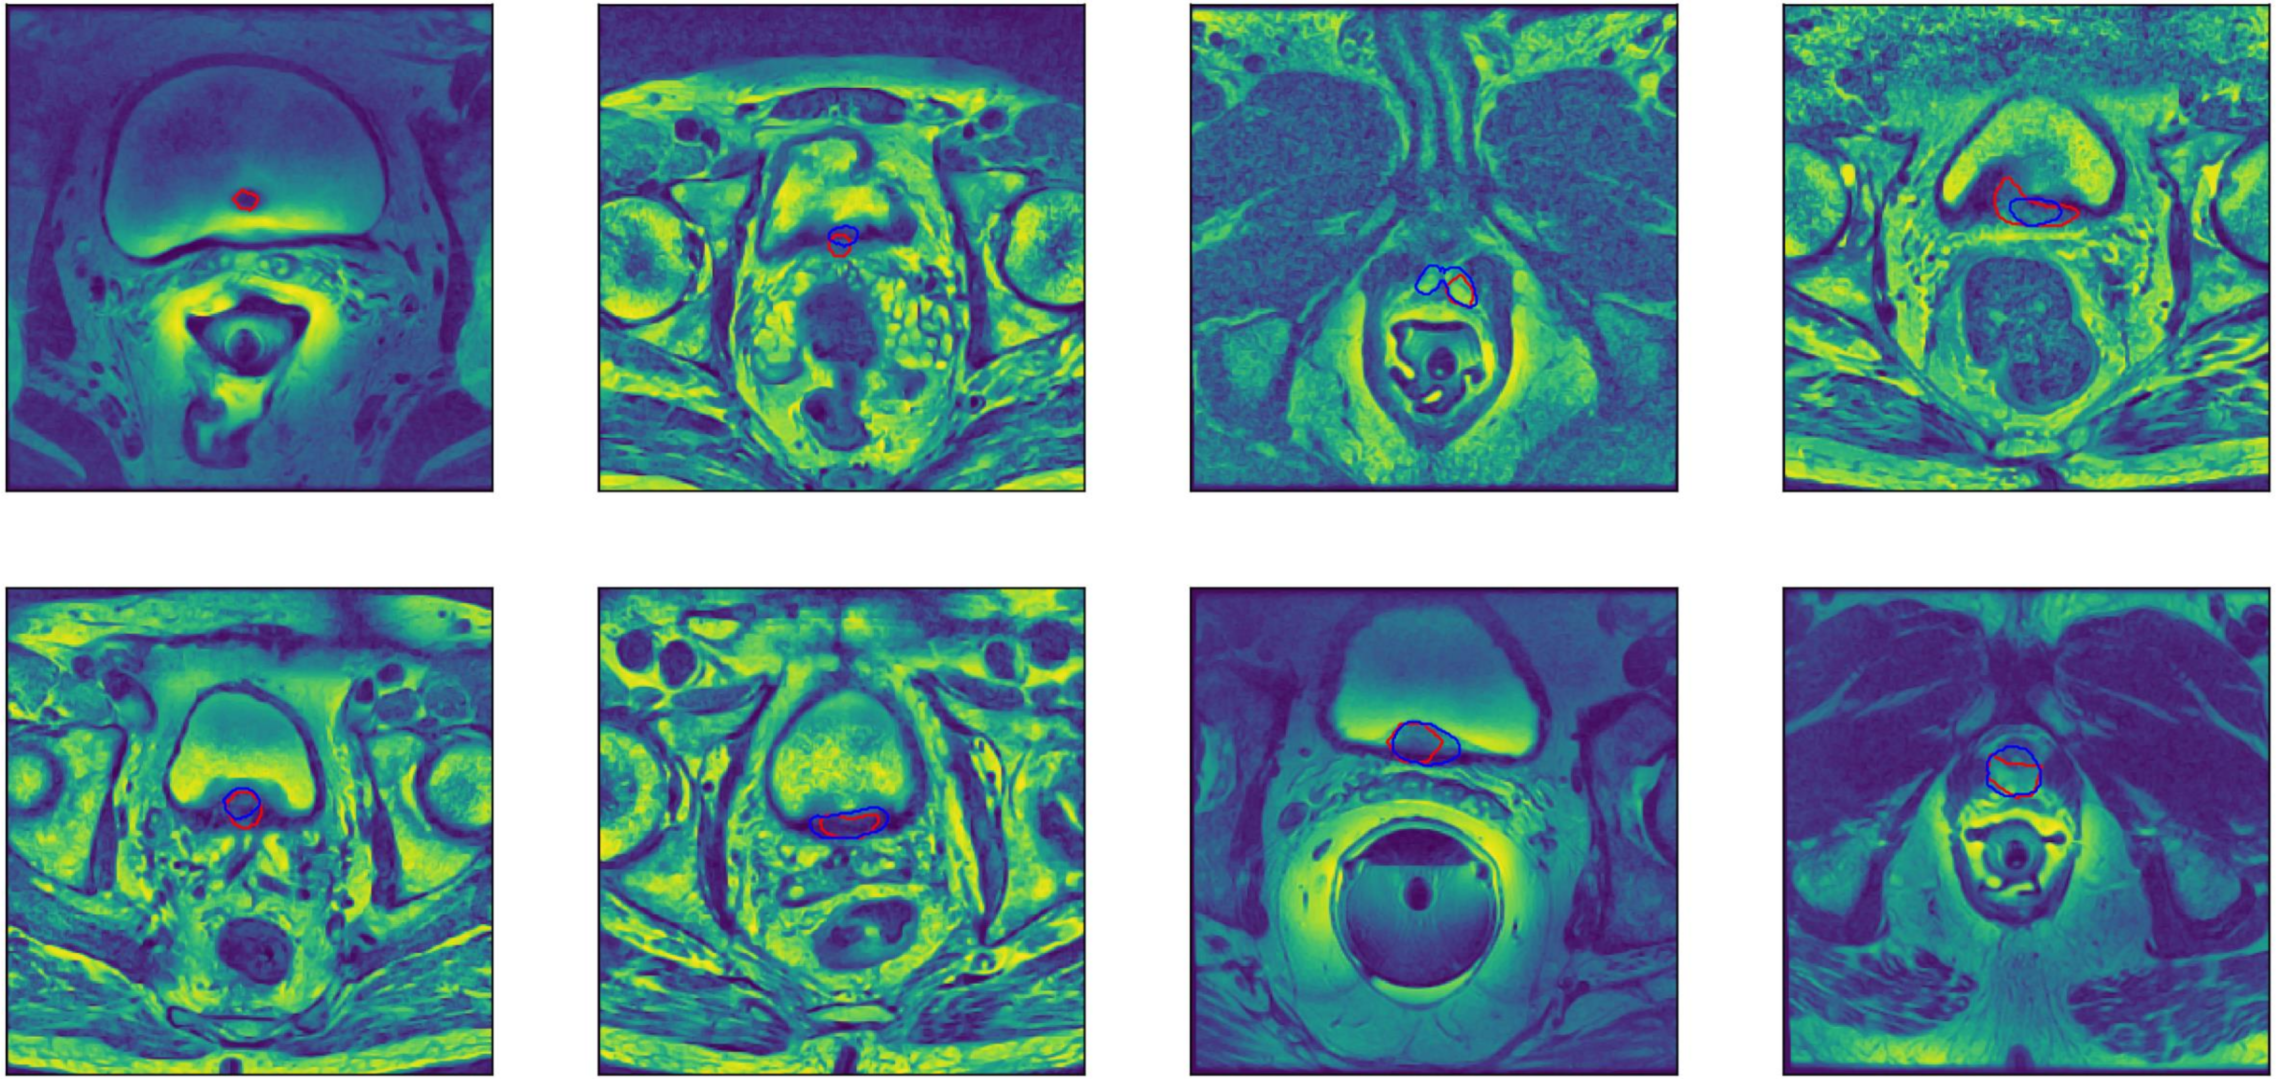

Result Big net