universität innsbruck AppliedMathematics

PiNet: Deep Structure Learning using Feature Extraction in Trained Projection Space

Christoph Angermann

https://applied-math.uibk.ac.at/

- **Motivation**
- **Methods**
  - Multiplanar U-net (**MPUnet**)
  - Projection-based method (**PiNet**)
- **Experiments**
- **Advantages & Drawbacks**
- **Future work**

Volumetric end-to-end segmentation

[1]Perslev et al., "One network to segment them all: A general, lightweight system for accurate 3d medical image segmentation".

[2]Milletari, Navab, and Ahmadi, "V-net: Fully convolutional neural networks for volumetric medical image segmentation".

Volumetric end-to-end segmentation

- **Naive approach:** Propagate 2D slices independently through a 2D segmentation framework $\Rightarrow$ loss of connections

[1] Perslev et al., “One network to segment them all: A general, lightweight system for accurate 3d medical image segmentation”.

[2] Milletari, Navab, and Ahmadi, “V-net: Fully convolutional neural networks for volumetric medical image segmentation”.

Volumetric end-to-end segmentation

- **Naive approach:** Propagate 2D slices independently through a 2D segmentation framework $\Rightarrow$ loss of connections
- **More advanced:** Sample volumes on 2D isotropic grids along multiple view axis $\Rightarrow$ **MPUnet**[1]

[1] Perslev et al., “One network to segment them all: A general, lightweight system for accurate 3d medical image segmentation”.

[2] Milletari, Navab, and Ahmadi, “V-net: Fully convolutional neural networks for volumetric medical image segmentation”.

Volumetric end-to-end segmentation

- **Naive approach:** Propagate 2D slices independently through a 2D segmentation framework $\Rightarrow$ loss of connections
- **More advanced:** Sample volumes on 2D isotropic grids along multiple view axis $\Rightarrow$ **MPUnet**[1]
- Extend U-net structure to higher dimension utilizing 3D convolutions[2] $\Rightarrow$ need of big memory resources & therefore restrictions to model complexity.

[1] Perslev et al., “One network to segment them all: A general, lightweight system for accurate 3d medical image segmentation”.

[2] Milletari, Navab, and Ahmadi, “V-net: Fully convolutional neural networks for volumetric medical image segmentation”.

Volumetric end-to-end segmentation

- **Naive approach:** Propagate 2D slices independently through a 2D segmentation framework $\Rightarrow$ loss of connections
- **More advanced:** Sample volumes on 2D isotropic grids along multiple view axis $\Rightarrow$ **MPUnet**[1]
- Extend U-net structure to higher dimension utilizing 3D convolutions[2] $\Rightarrow$ need of big memory resources & therefore restrictions to model complexity.
- Enable feature extraction in a learned projection space, i.e. transfer segmentation task into a bundle of regression problems in lower dimension $\Rightarrow$ **PiNet**

[1] Perslev et al., “One network to segment them all: A general, lightweight system for accurate 3d medical image segmentation”.

[2] Milletari, Navab, and Ahmadi, “V-net: Fully convolutional neural networks for volumetric medical image segmentation”.

MPUnet

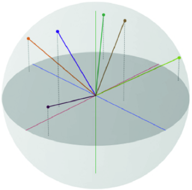 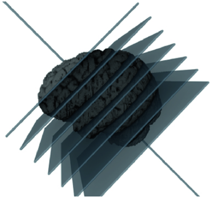 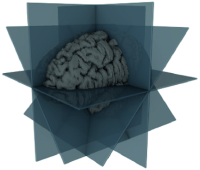

MPUnet

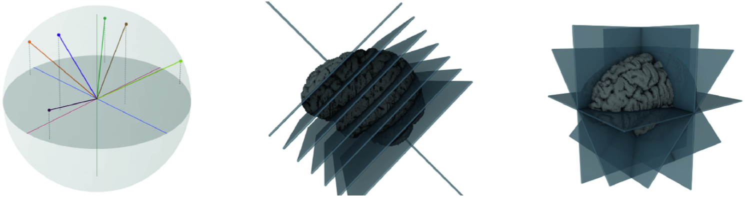

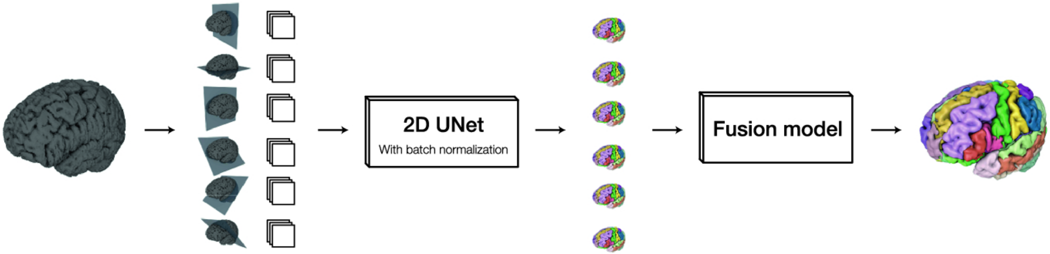

MPUnet

MPUnet

- Random samples unit vectors $V = \{v_1, \ldots, v_M\}$ in $\mathbb{R}^3$ (**orientations**) $\Rightarrow M$ different segmentation volumes $\{P_v \in \mathbb{R}^{d_1 \times d_2 \times d_3 \times K} \mid v \in V\}$

MPUnet

- Random samples unit vectors $V = \{v_1, \ldots, v_M\}$ in $\mathbb{R}^3$ (**orientations**) $\Rightarrow M$ different segmentation volumes $\{P_v \in \mathbb{R}^{d_1 \times d_2 \times d_3 \times K} \mid v \in V\}$
- **Fusion model:** For all $d_1 \cdot d_2 \cdot d_3$ voxels x and each class $k \in \{1, ..., K\}$, the fusion model

$$f_{\text{fusion}} : \mathbb{R}^{M \times K} \longrightarrow \mathbb{R}^K$$

calculates

$$z(x)_k = \sum_{n=1}^{M} W_{n,k} \cdot p_{n,x,k} + \beta_k,$$

where $p_{n,x,k}$ is the U-net prediction of voxel x for class k in P_{v_n} and the weight and bias matrices $W_{n,k}$ and β_k are adjusted via a small training run.

MPUnet

Advantages:

[3]Simpson et al., “A large annotated medical image dataset for the development and evaluation of segmentation algorithms”.

MPUnet

Advantages:

- Augmentation: Processing the input data from different views has the same effect as applying affine transformations to the 3D input and presenting the transformed images to a (single-view) network.

[3] Simpson et al., "A large annotated medical image dataset for the development and evaluation of segmentation algorithms".

MPUnet

Advantages:

- Augmentation: Processing the input data from different views has the same effect as applying affine transformations to the 3D input and presenting the transformed images to a (single-view) network.
- Generalizibility: Without any hyperparameter fine-tuning, MPUnet was able to compete with SOTA frameworks on the Medical Segmentation Decathlon 2018[3] (10 different segmentation tasks with one architecture).

[3] Simpson et al., “A large annotated medical image dataset for the development and evaluation of segmentation algorithms”.

MPUnet

Advantages:

- Augmentation: Processing the input data from different views has the same effect as applying affine transformations to the 3D input and presenting the transformed images to a (single-view) network.
- Generalizibility: Without any hyperparameter fine-tuning, MPUnet was able to compete with SOTA frameworks on the Medical Segmentation Decathlon 2018[3] (10 different segmentation tasks with one architecture).

Disadvantages:

[3] Simpson et al., "A large annotated medical image dataset for the development and evaluation of segmentation algorithms".

MPUnet

Advantages:

- Augmentation: Processing the input data from different views has the same effect as applying affine transformations to the 3D input and presenting the transformed images to a (single-view) network.
- Generalizibility: Without any hyperparameter fine-tuning, MPUnet was able to compete with SOTA frameworks on the Medical Segmentation Decathlon 2018[3] (10 different segmentation tasks with one architecture).

Disadvantages:

- Still high network complexity: 2D segmentation U-net consists of $\approx 62 \times 10^6$ parameters. Sufficient amount of orientations has to be considered to ensure reasonable segmentation.

[3] Simpson et al., "A large annotated medical image dataset for the development and evaluation of segmentation algorithms".

MPUnet

Advantages:

- Augmentation: Processing the input data from different views has the same effect as applying affine transformations to the 3D input and presenting the transformed images to a (single-view) network.
- Generalizibility: Without any hyperparameter fine-tuning, MPUnet was able to compete with SOTA frameworks on the Medical Segmentation Decathlon 2018[3] (10 different segmentation tasks with one architecture).

Disadvantages:

- Still high network complexity: 2D segmentation U-net consists of $\approx 62 \times 10^6$ parameters. Sufficient amount of orientations has to be considered to ensure reasonable segmentation.
- Redundancy: Lot of redundant computations during inference due to evaluation along every slice for all orientations.

[3] Simpson et al., "A large annotated medical image dataset for the development and evaluation of segmentation algorithms".

PiNet

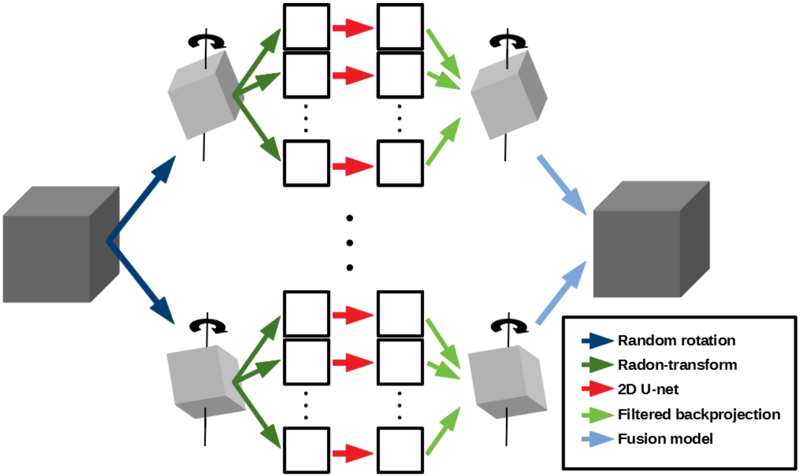

Figure: Description in 3 spatial dimensions.

PiNet

For 3D input x, the proposed PiNet takes the following form:

$$\Pi(x) = \mathcal{F} \circ \begin{bmatrix} \mathcal{B} \circ \begin{bmatrix} \mathcal{X} \circ \Phi \circ \mathcal{P}_1 \circ \mathcal{Q}_1(x) \\ \vdots \\ \mathcal{X} \circ \Phi \circ \mathcal{P}_p \circ \mathcal{Q}_1(x) \end{bmatrix} \\ \vdots \\ \mathcal{B} \circ \begin{bmatrix} \mathcal{X} \circ \Phi \circ \mathcal{P}_1 \circ \mathcal{Q}_v(x) \\ \vdots \\ \mathcal{X} \circ \Phi \circ \mathcal{P}_p \circ \mathcal{Q}_v(x) \end{bmatrix} \end{bmatrix}.$$

PiNet

For 3D input x, the proposed PiNet takes the following form:

$$\Pi(x) = \mathcal{F} \circ \begin{bmatrix} \mathcal{B} \circ \begin{bmatrix} \mathcal{X} \circ \Phi \circ \mathcal{P}_1 \circ \mathcal{Q}_1(x) \\ \vdots \\ \mathcal{X} \circ \Phi \circ \mathcal{P}_p \circ \mathcal{Q}_1(x) \end{bmatrix} \\ \vdots \\ \mathcal{B} \circ \begin{bmatrix} \mathcal{X} \circ \Phi \circ \mathcal{P}_1 \circ \mathcal{Q}_v(x) \\ \vdots \\ \mathcal{X} \circ \Phi \circ \mathcal{P}_p \circ \mathcal{Q}_v(x) \end{bmatrix} \end{bmatrix}.$$

Random rotation $\mathcal{Q}$:
$\mathcal{Q}_i : \mathbb{R}^{d_1 \times d_2 \times d_3} \to \mathbb{R}^{d_1 \times d_2 \times d_3}$, $i = 1, \ldots, v$ induces independent random rotations in 3D, i.e. the volume is rotated in the 3 main planes using uniformly at random chosen angles $(\alpha_{i_1}, \alpha_{i_2}, \alpha_{i_3}) \in [30, 150]$, $i = 1, \ldots, v$.

PiNet

$$\Pi(x) = \mathcal{F} \circ \begin{bmatrix} \mathcal{B} \circ \begin{bmatrix} \mathcal{X} \circ \Phi \circ \mathcal{P}_1 \circ \mathcal{Q}_1(x) \\ \vdots \\ \mathcal{X} \circ \Phi \circ \mathcal{P}_p \circ \mathcal{Q}_1(x) \end{bmatrix} \\ \vdots \\ \mathcal{B} \circ \begin{bmatrix} \mathcal{X} \circ \Phi \circ \mathcal{P}_1 \circ \mathcal{Q}_v(x) \\ \vdots \\ \mathcal{X} \circ \Phi \circ \mathcal{P}_p \circ \mathcal{Q}_v(x) \end{bmatrix} \end{bmatrix}.$$

PiNet

$$\Pi(x) = \mathcal{F} \circ \begin{bmatrix} \mathcal{B} \circ \begin{bmatrix} \mathcal{X} \circ \Phi \circ \mathcal{P}_1 \circ \mathcal{Q}_1(x) \\ \vdots \\ \mathcal{X} \circ \Phi \circ \mathcal{P}_p \circ \mathcal{Q}_1(x) \end{bmatrix} \\ \vdots \\ \mathcal{B} \circ \begin{bmatrix} \mathcal{X} \circ \Phi \circ \mathcal{P}_1 \circ \mathcal{Q}_v(x) \\ \vdots \\ \mathcal{X} \circ \Phi \circ \mathcal{P}_p \circ \mathcal{Q}_v(x) \end{bmatrix} \end{bmatrix}.$$

Learned weighted Radon-transform $\mathcal{P}$: $\mathcal{P}_i : \mathbb{R}^{d_1 \times d_2 \times d_3} \to \mathbb{R}^{d_2 \times d_3}$, $i = 1, \ldots, p$ is a trainable Radon-transform operator. For $M \in \mathbb{N}$, the volume is rotated around the z-axis for equidistant angles in $\Theta \triangleq \left\{k \times \frac{180 \cdot M}{d_1 \cdot \pi} \;\middle|\; k = 1, 2, \ldots, \left\lfloor \frac{d_1 \cdot \pi}{M} \right\rfloor \right\}$ and for each direction interpolated slices are summed over first spatial axis.

PiNet

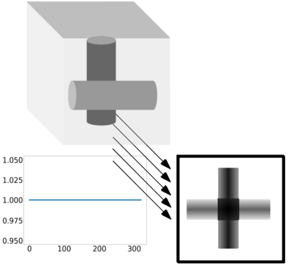
1.050
1.025
1.000
0.975
0.950
0
100
200
300

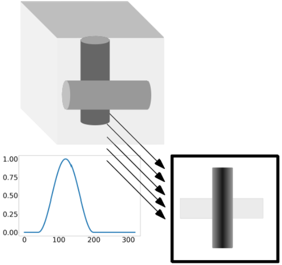
1.00
0.75
0.50
0.25
0.00
0
100
200
300

PiNet

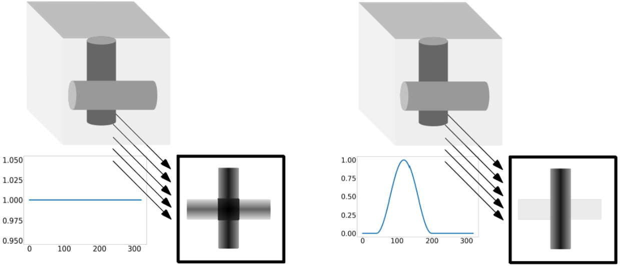

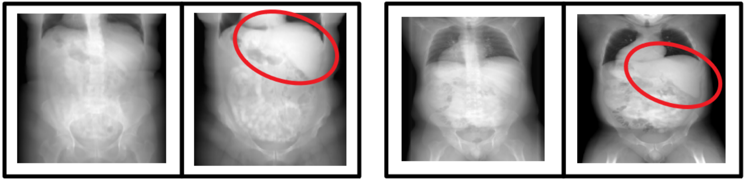

PiNet

$$\Pi(x) = \mathcal{F} \circ \begin{bmatrix} \mathcal{B} \circ \begin{bmatrix} \mathcal{X} \circ \Phi \circ \mathcal{P}_1 \circ \mathcal{Q}_1(x) \\ \vdots \\ \mathcal{X} \circ \Phi \circ \mathcal{P}_p \circ \mathcal{Q}_1(x) \end{bmatrix} \\ \vdots \\ \mathcal{B} \circ \begin{bmatrix} \mathcal{X} \circ \Phi \circ \mathcal{P}_1 \circ \mathcal{Q}_v(x) \\ \vdots \\ \mathcal{X} \circ \Phi \circ \mathcal{P}_p \circ \mathcal{Q}_v(x) \end{bmatrix} \end{bmatrix}.$$

PiNet

$$\Pi(x) = \mathcal{F} \circ \begin{bmatrix} \mathcal{B} \circ \begin{bmatrix} \mathcal{X} \circ \Phi \circ \mathcal{P}_1 \circ \mathcal{Q}_1(x) \\ \vdots \\ \mathcal{X} \circ \Phi \circ \mathcal{P}_p \circ \mathcal{Q}_1(x) \end{bmatrix} \\ \vdots \\ \mathcal{B} \circ \begin{bmatrix} \mathcal{X} \circ \Phi \circ \mathcal{P}_1 \circ \mathcal{Q}_v(x) \\ \vdots \\ \mathcal{X} \circ \Phi \circ \mathcal{P}_p \circ \mathcal{Q}_v(x) \end{bmatrix} \end{bmatrix}.$$

2D network Φ: $\Phi : \mathbb{R}^{d_2 \times d_3} \to [0,1]^{d_2 \times d_3 \times c}$ is a U-net for segmentation of 2D projection images. This U-net operator is responsible for feature extraction.

Hierarchical convolution module

Hierarchical convolution module

- Since we consider Radon-transform projection inputs $x_R \in \mathbb{R}^{d_2 \times d_3}$ $\Rightarrow$ targets $y_R \in \mathbb{R}^{d_2 \times d_3 \times c}$ are not binary any more $\Rightarrow$ transfers segmentation into a regression task for each output channel.

Hierarchical convolution module

- Since we consider Radon-transform projection inputs $x_R \in \mathbb{R}^{d_2 \times d_3} \Rightarrow$ targets $y_R \in \mathbb{R}^{d_2 \times d_3 \times c}$ are not binary any more $\Rightarrow$ transfers segmentation into a regression task for each output channel.
- Due to high sparsity of targeted objects, minimization with standard regression losses (MSE,MAE,Huber loss,...) may fail.

Hierarchical convolution module

- Since we consider Radon-transform projection inputs $x_R \in \mathbb{R}^{d_2 \times d_3} \Rightarrow$ targets $y_R \in \mathbb{R}^{d_2 \times d_3 \times c}$ are not binary any more $\Rightarrow$ transfers segmentation into a regression task for each output channel.
- Due to high sparsity of targeted objects, minimization with standard regression losses (MSE,MAE,Huber loss,...) may fail.
- Modification of U-net output layer $\Rightarrow$ discretize targets, i.e. divide targets $y_R \in \mathbb{R}^{d_2 \times d_3 \times c}$ by maximum, discretisize for the scaled versions the interval $[0, 1]$ into b equidistant bins $\Rightarrow$ new targets $y_d \in \{0, 1\}^{d_2 \times d_3 \times c \times b}$.

Hierarchical convolution module

- Since we consider Radon-transform projection inputs $x_R \in \mathbb{R}^{d_2 \times d_3} \Rightarrow$ targets $y_R \in \mathbb{R}^{d_2 \times d_3 \times c}$ are not binary any more $\Rightarrow$ transfers segmentation into a regression task for each output channel.
- Due to high sparsity of targeted objects, minimization with standard regression losses (MSE,MAE,Huber loss,...) may fail.
- Modification of U-net output layer $\Rightarrow$ discretize targets, i.e. divide targets $y_R \in \mathbb{R}^{d_2 \times d_3 \times c}$ by maximum, discretisize for the scaled versions the interval $[0, 1]$ into b equidistant bins $\Rightarrow$ new targets $y_d \in \{0, 1\}^{d_2 \times d_3 \times c \times b}$.
- The regression problem is therefore transformed again into a segmentation problem for each of the b depth channels, which follow an hierarchical order $\Rightarrow$ need to ensure dependence between output channels during optimization.

Hierarchical convolution module

Definition

Let $L \in \mathbb{R}^{m \times n \times f}$ be output of some network layer with spatial dimensions m, n and channel size f. Let $F_i : \mathbb{R}^{m \times n \times (f+i-1)} \to [0,1]^{m \times n \times 1}$, $i = 1, \ldots, b$ be trainable convolution modules with $f + i - 1$ input channels and output channel size 1. Furthermore, let C denote concatenation over the last axis. The output of hierarchical convolution module with b channels is

$$\hat{y}_d = [o_1, \ldots, o_b] \in [0,1]^{m \times n \times b},$$

where

$$\begin{aligned} o_1 &= F_1(L) \\ o_2 &= F_2(C(L, o_1)) \\ &\vdots \\ o_b &= F_b(C(L, o_{b-1})) \end{aligned}$$

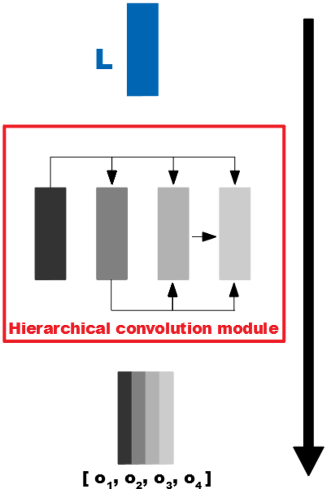

PiNet

$$\Pi(x) = \mathcal{F} \circ \begin{bmatrix} \mathcal{B} \circ \begin{bmatrix} \mathcal{X} \circ \Phi \circ \mathcal{P}_1 \circ \mathcal{Q}_1(x) \\ \vdots \\ \mathcal{X} \circ \Phi \circ \mathcal{P}_p \circ \mathcal{Q}_1(x) \end{bmatrix} \\ \vdots \\ \mathcal{B} \circ \begin{bmatrix} \mathcal{X} \circ \Phi \circ \mathcal{P}_1 \circ \mathcal{Q}_v(x) \\ \vdots \\ \mathcal{X} \circ \Phi \circ \mathcal{P}_p \circ \mathcal{Q}_v(x) \end{bmatrix} \end{bmatrix}.$$

PiNet

$$\Pi(x) = \mathcal{F} \circ \begin{bmatrix} \mathcal{B} \circ \begin{bmatrix} \mathcal{X} \circ \Phi \circ \mathcal{P}_1 \circ \mathcal{Q}_1(x) \\ \vdots \\ \mathcal{X} \circ \Phi \circ \mathcal{P}_p \circ \mathcal{Q}_1(x) \end{bmatrix} \\ \vdots \\ \mathcal{B} \circ \begin{bmatrix} \mathcal{X} \circ \Phi \circ \mathcal{P}_1 \circ \mathcal{Q}_v(x) \\ \vdots \\ \mathcal{X} \circ \Phi \circ \mathcal{P}_p \circ \mathcal{Q}_v(x) \end{bmatrix} \end{bmatrix}.$$

Filtered backprojection $\mathcal{B} \circ \mathcal{X}$: Segmented projection images $\hat{y}_d = [o_1, \ldots, o_b] \in [0,1]^{d_2 \times d_3 \times c \times b}$ are averaged over the last axis and lift again to volumetric data using filtered backprojection algorithm $\Rightarrow$ volumetric segmentation masks for each orientation $\hat{y}_1, \ldots, \hat{y}_v \in [0,1]^{d_1 \times d_2 \times d_3 \times c}$.

PiNet

$$\Pi(x) = \mathcal{F} \circ \begin{bmatrix} \mathcal{B} \circ \begin{bmatrix} \mathcal{X} \circ \Phi \circ \mathcal{P}_1 \circ \mathcal{Q}_1(x) \\ \vdots \\ \mathcal{X} \circ \Phi \circ \mathcal{P}_p \circ \mathcal{Q}_1(x) \end{bmatrix} \\ \vdots \\ \mathcal{B} \circ \begin{bmatrix} \mathcal{X} \circ \Phi \circ \mathcal{P}_1 \circ \mathcal{Q}_v(x) \\ \vdots \\ \mathcal{X} \circ \Phi \circ \mathcal{P}_p \circ \mathcal{Q}_v(x) \end{bmatrix} \end{bmatrix}.$$

PiNet

$$\Pi(x) = \mathcal{F} \circ \begin{bmatrix} \mathcal{B} \circ \begin{bmatrix} \mathcal{X} \circ \Phi \circ \mathcal{P}_1 \circ \mathcal{Q}_1(x) \\ \vdots \\ \mathcal{X} \circ \Phi \circ \mathcal{P}_p \circ \mathcal{Q}_1(x) \end{bmatrix} \\ \vdots \\ \mathcal{B} \circ \begin{bmatrix} \mathcal{X} \circ \Phi \circ \mathcal{P}_1 \circ \mathcal{Q}_v(x) \\ \vdots \\ \mathcal{X} \circ \Phi \circ \mathcal{P}_p \circ \mathcal{Q}_v(x) \end{bmatrix} \end{bmatrix}.$$

Fusion model $\mathcal{F}$: $\mathcal{F} : \left([0,1]^{d_1 \times d_2 \times d_3 \times c}\right)^v \to [0,1]^{d_1 \times d_2 \times d_3 \times c}$ to combine the different orientations to one final output $\hat{y}$:

$$\hat{y} = \frac{1}{v} \left[\sum_{l=1}^{v} \mathcal{Q}_l^{-1}\left(\hat{y}_l\right) \cdot W[l,k] \right]_{k=1}^{c}.$$

Experiments

Left cardiac atrium segmentation of mono-modal MRI scans

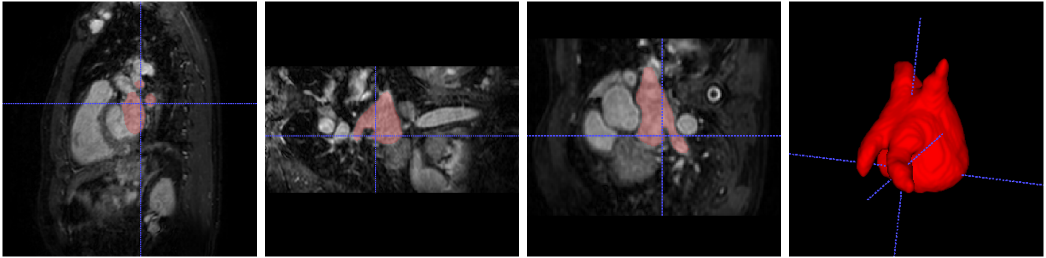

Figure: Left to right: axial, sagittal, coronal, 3D visualization

Given are 20 MRI volumes with corresponding segmentations $\Rightarrow$ evaluation via 3-fold cross-validation.

Experiments

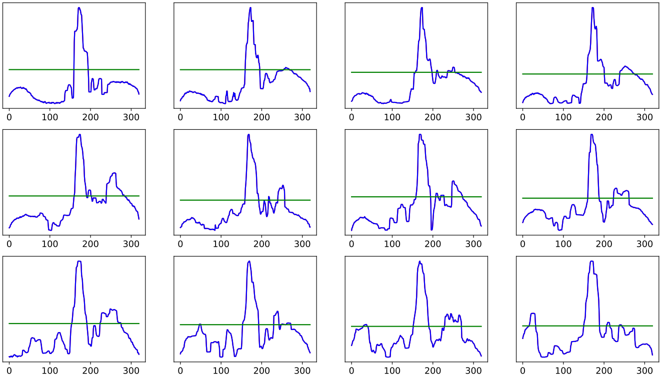

Figure: Example for self adjusting slice weights of the Radon-transform operator.

Experiments

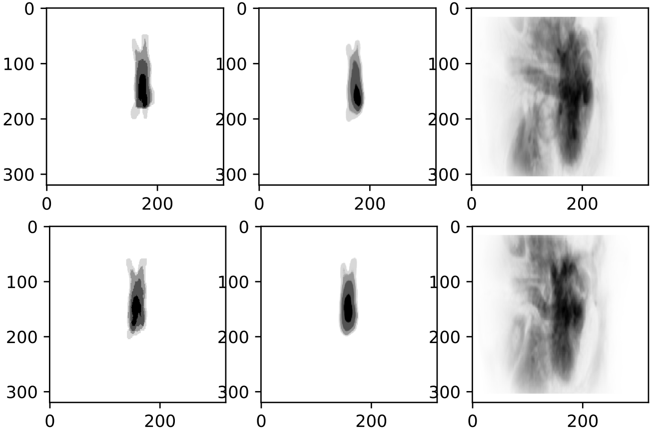

Figure: Left to right: discretised ground truth, output of hierarchical convolution module, weighted Radon-transform.

Experiments

Left cardiac atrium segmentation of mono-modal MRI scans

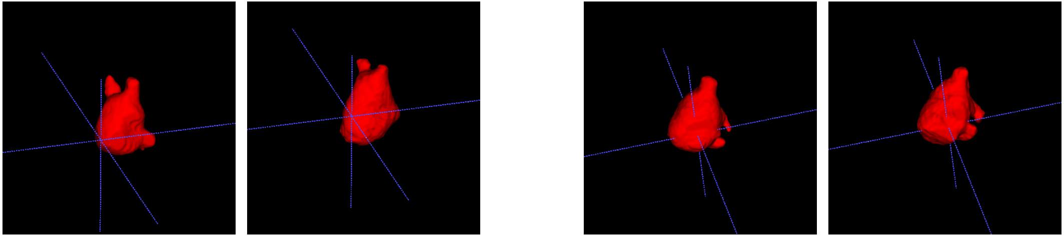

Figure: Left: Ground truth. Right: PiNet prediction.

Experiments

Table: Comparison between MPUnet and PiNet for small datasets of the 2018 Medical Segmentation Decathlon.

Task	**Size**	**Network**	**Labels**	**Dice score**	**< 4Gb GPU**
Left cardiac atrium	20	MPUnet	1	0.89 ± 0.09	no
		PiNet	1	0.88 ± 0.05	yes
Spleen	41	MPUnet	1	0.95	no
		PiNet	1	0.93 ± 0.04	yes

Advantages

Advantages

- Network architecture enabling learnable Radon-transform operator to transfer segmentation tasks to lower-dimensional space.

Advantages

- Network architecture enabling learnable Radon-transform operator to transfer segmentation tasks to lower-dimensional space.
- Projection images lifted again to original dimension deploying self-adjusting reconstruction operators.

Advantages

- Network architecture enabling learnable Radon-transform operator to transfer segmentation tasks to lower-dimensional space.
- Projection images lifted again to original dimension deploying self-adjusting reconstruction operators.
- Universality and plausibility for very small 3D datasets without need of further data augmentation.

Advantages

- Network architecture enabling learnable Radon-transform operator to transfer segmentation tasks to lower-dimensional space.
- Projection images lifted again to original dimension deploying self-adjusting reconstruction operators.
- Universality and plausibility for very small 3D datasets without need of further data augmentation.
- Trainable and applicable on hardware with low-end GPU resources.

Advantages

- Network architecture enabling learnable Radon-transform operator to transfer segmentation tasks to lower-dimensional space.
- Projection images lifted again to original dimension deploying self-adjusting reconstruction operators.
- Universality and plausibility for very small 3D datasets without need of further data augmentation.
- Trainable and applicable on hardware with low-end GPU resources.
- Total amount of training parameters: $\approx 8 \times 10^6$.

Drawbacks

Drawbacks

- Although the network is able to locate the organ's position and returns a satisfying 3D segmentation with respect to Dice metric, fine structures (e.g. outgoing arteries) are not satisfyingly captured.

Drawbacks

- Although the network is able to locate the organ's position and returns a satisfying 3D segmentation with respect to Dice metric, fine structures (e.g. outgoing arteries) are not satisfyingly captured.

- The learned slice weighting of the Radon-transform operator is adjusted by the training data but actually does not dependent on the network input during inference.

Drawbacks

- Although the network is able to locate the organ’s position and returns a satisfying 3D segmentation with respect to Dice metric, fine structures (e.g. outgoing arteries) are not satisfyingly captured.

- The learned slice weighting of the Radon-transform operator is adjusted by the training data but actually does not dependent on the network input during inference.

- **Ideas?**

Future work

Future work

- Leverage PiNet to multilabel segmentation.

Future work

- Leverage PiNet to multilabel segmentation.
- Deployment of PiNet for high-dimensional image reconstruction tasks.

Future work

- Leverage PiNet to multilabel segmentation.
- Deployment of PiNet for high-dimensional image reconstruction tasks.
- Use time and memory efficient aspects of PiNet to implement hyper-parameter fine-tuning in a fully automated manner.

Future work

- Leverage PiNet to multilabel segmentation.
- Deployment of PiNet for high-dimensional image reconstruction tasks.
- Use time and memory efficient aspects of PiNet to implement hyper-parameter fine-tuning in a fully automated manner.
- Further develop Radon-transform operator to ensure input dependency during inference.

Thank you for your attention!

Christoph Angermann
https://applied-math.uibk.ac.at/