



Treatment planning adaptation in ion beam therapy using sparse in-room ion radiographies

Markus Schöttle, Ines Butz, Julian Titze, Hector Andrade Loarca, Ozan Öktem, Gitta Kutyniok, Katia Parodi,
Chiara Gianoli

Chair of Experimental Physics – Medical Physics

Faculty for Physics

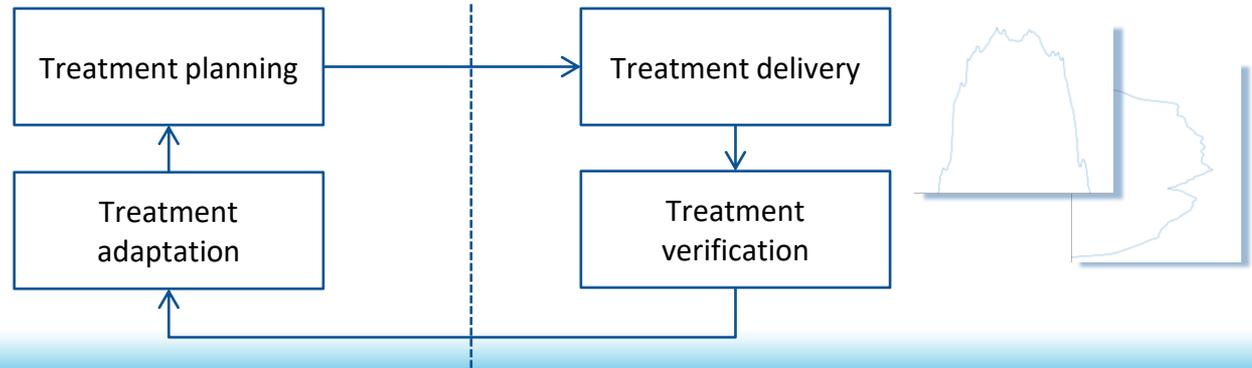
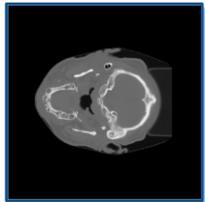
Ludwig-Maximilians-Universität München

Ion Imaging Workshop 2025 - Florence, Italy - 20-21 October 2025

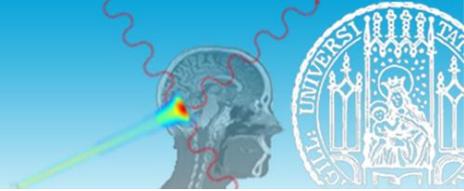
The role of ion imaging in ion beam therapy



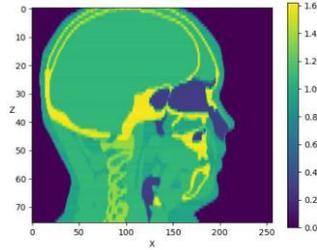
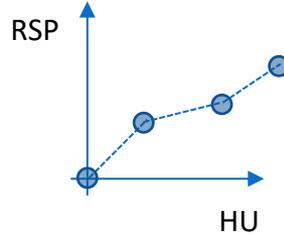
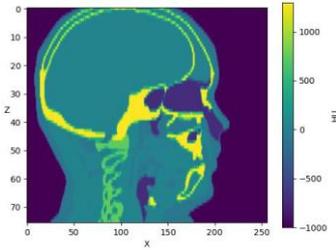
- In current clinical practice, ion beam therapy still relies on **X-ray imaging**
- X-ray imaging introduces inaccuracies in the patient model due to **semi-empirical calibration** from linear attenuation to relative stopping power
- Ion imaging provides direct relative stopping power information, potentially eliminating or reducing these inaccuracies in **treatment planning**
- As an in-room modality, ion imaging offers also potential for **adaptive radiation therapy**



The role of ion imaging in ion beam therapy

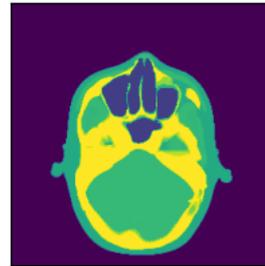
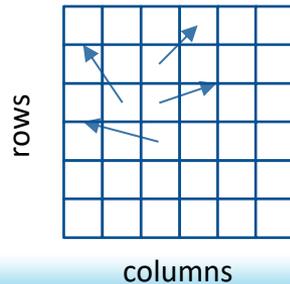
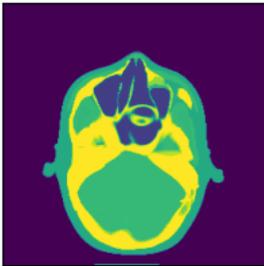


- The investigation of the role of ion imaging in **treatment planning** targets the calibration inaccuracies

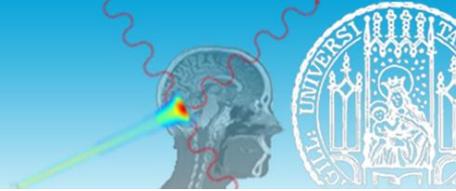


Rationale of the work of Ines Butz

- The investigation of the role of ion imaging in **treatment adaptation** targets inter-fractional anatomical changes occurring between treatment planning and treatment delivery due to physiological deformations



The role of ion imaging in ion beam therapy



- Both problems have been proposed and investigated using conventional methodologies (i.e., **model-based approaches**) within the framework of the project “Hybrid ImaGing framework in Hadrontherapy for Adaptive Radiation Therapy” (HIGH ART), financed by the *Deutsche Forschungsgemeinschaft* (DFG)
- A pathway for clinical translation from X-ray to ion imaging is envisioned through the combination of X-ray and ion imaging based on sparse ion radiographies, for both treatment planning and treatment adaptation
- Within the framework of the renewal project, the reinterpretation of both problems in terms of machine learning and deep learning (i.e., **data-driven approaches**) has been proposed and investigated



Chiara Gianoli



Katia Parodi



Prasannakumar Palaniappan



Ines Butz

Gefördert durch

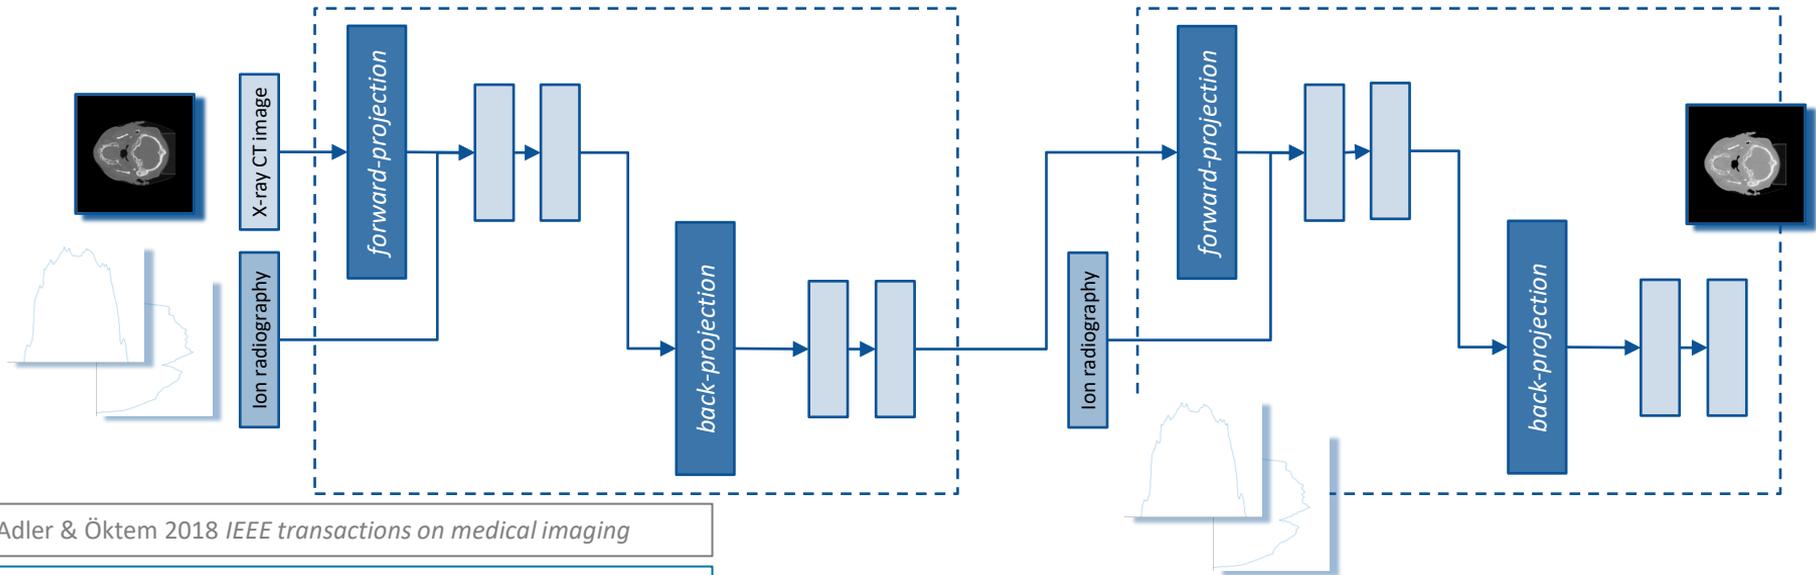
DFG Deutsche
Forschungsgemeinschaft



The Learned Primal-Dual network



- The **Learned Primal-Dual (LPD) network**, originally developed for deep image reconstruction, is investigated to solve the registration problem based on sparse ion radiographies
- The LPD network is composed by 10 Primal-Dual blocks

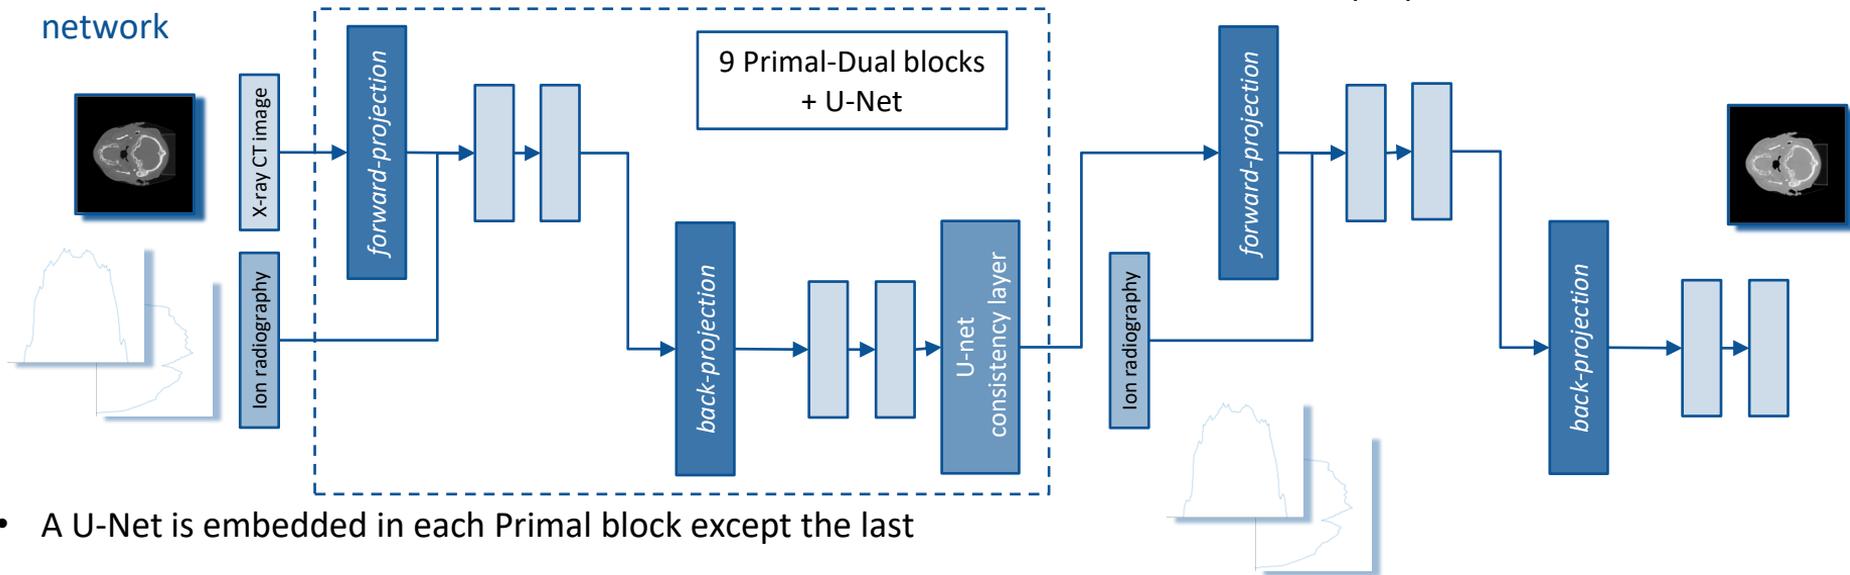


Adler & Öktem 2018 *IEEE transactions on medical imaging*

Butz, ..., Parodi, Gianoli 2025 *Phys. Med. Biol.* (under revision)



- An extension of the Learned Primal-Dual network based on a U-Net architecture is proposed, referred to as **U-LPD network**



- A U-Net is embedded in each Primal block except the last
- To enforce consistency between the U-Net output and the radiographies, a consistency layer is incorporated

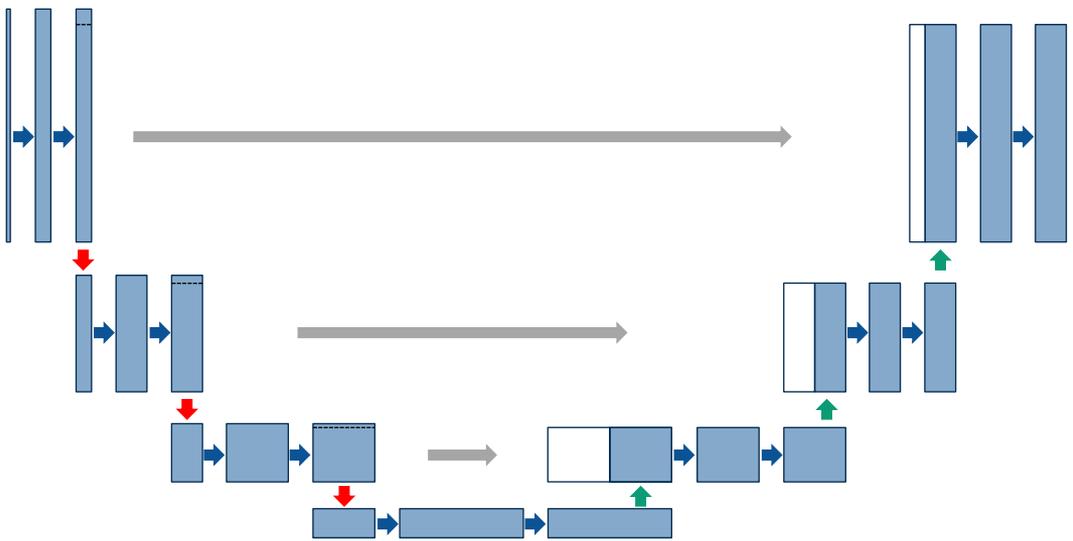
$$Primal_{consistent} = Primal_{U-net} + BP(Dual - FP(Primal_{U-net}))$$

U-net embedded in the Learned Primal-Dual network



- Base filter = 4
- Channel expansion factor = 2 ($4 \rightarrow 8 \rightarrow 16$)
- Number of layers = 3

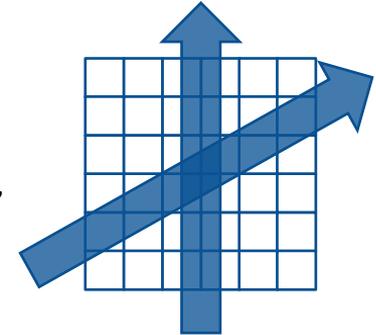
- ▲ 3D convolution ($3 \times 3 \times 3$) + ReLU
- ▶ Concatenative skip connection
- ➡ Max pooling
- ▼ 3D transposed convolution ($3 \times 3 \times 3$)



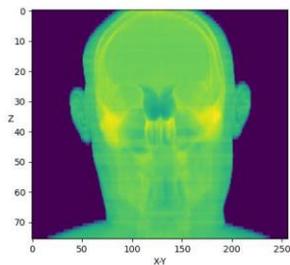
Definition of sparse ion radiographies



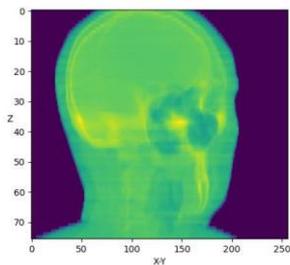
- The dataset is composed by ten original anthropomorphic phantoms, calibrated from linear attenuation to relative stopping power, assumed as ground truth images
 - Each slice is calibrated according to a randomly generated calibration curve
- Ion radiographies are generated by **forward-projecting** the ground truth images, simulating the in-room treatment delivery scenario
 - A total of **five radiographies** are considered, uniformly distributed over 180 degrees



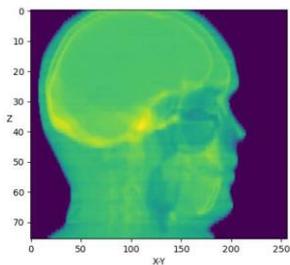
0 degree



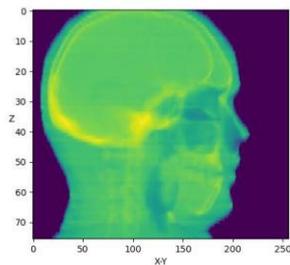
36 degrees



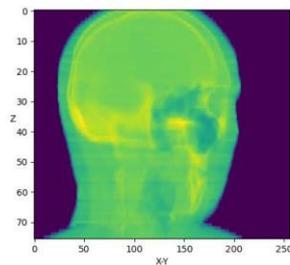
72 degrees



108 degrees



144 degrees



Definition of transformation fields



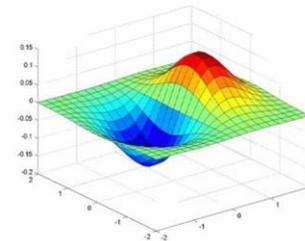
- To generate the treatment planning scenario, controlled anatomical changes are applied to the ground truth images by deforming the phantoms using **Gaussian derivative-based transformations**, simulating effects such as tissue shrinkage and dilation
- The deformation field is defined based on anisotropic Gaussian derivative functions

$$G(x, y, z) = \exp\left(-\left[\frac{(x - \mu_x)^2}{2 \sigma_x^2} + \frac{(y - \mu_y)^2}{2 \sigma_y^2} + \frac{(z - \mu_z)^2}{2 \sigma_z^2}\right]\right)$$

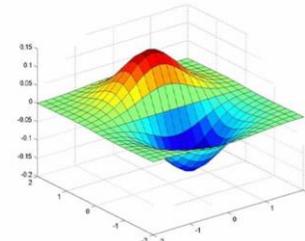
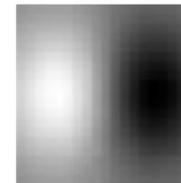
$$G_x = -\alpha_x \left(\frac{(x - \mu_x)}{\sigma_x^2}\right) G \quad G_y = -\alpha_y \left(\frac{(y - \mu_y)}{\sigma_y^2}\right) G$$

$$G_z = -\alpha_z \left(\frac{(z - \mu_z)}{\sigma_z^2}\right) G$$

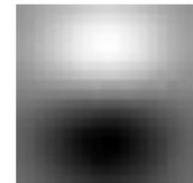
Composite of three
orthogonally
oriented lobed
patterns



x-direction



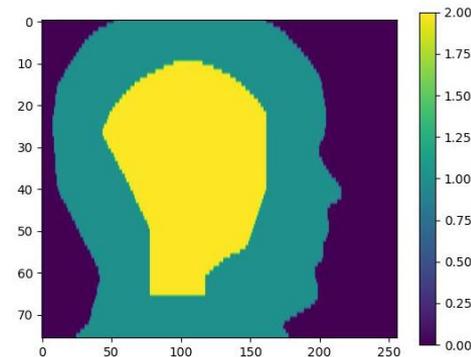
y-direction



Parameters of the Gaussian derivative functions

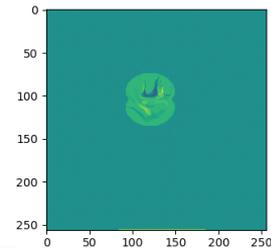


- The **center** of the Gaussian function is randomly selected within an inner mask of the original anthropomorphic phantoms
- The **standard deviation** of the Gaussian function are randomly selected within realistic ranges of transformation size
 - 7-10.5 pixels for σ_x and σ_y (pixel size 1 mm)
 - 2.5-3.5 pixels for σ_z (pixel size 3 mm)
- The **amplitude** of the Gaussian derivatives α_x , α_y , α_z are randomly scaled according to realistic ranges of transformation amplitude
 - 3-7.5 pixels for α_x and α_y (pixel size 1 mm)
 - 1-2.5 pixels for α_z (pixel size 3 mm)
- For each central slice, three adjacent slices are included on either side, resulting in an input **slab of seven consecutive slices**





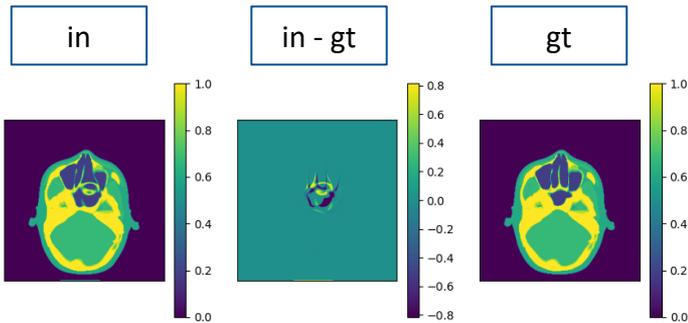
- 10 anthropomorphic phantoms, 3 transformations applied to each phantom, 5 samplings per transformation
 - 60 slabs selected for the “full training” study, 50 for training, 10 for validation
 - 10 additional slabs used for testing, sourced from phantoms excluded from training and validation
- Deterministic “subset training” studies with batch sizes equal to the training data size (1, 2, 4, and 8 slabs) and 1 validation slab (i.e., the “example slab”)
- Loss Function based on mean squared error (**MSE**), computed with respect to the ground truth image
- Task-specific evaluation metrics calculated within the region of applied transformation, defined by a threshold on the magnitude of the ground truth transformation field
 - Root mean square error (**RMSE**)
 - Structural similarity index (**SSIM**)



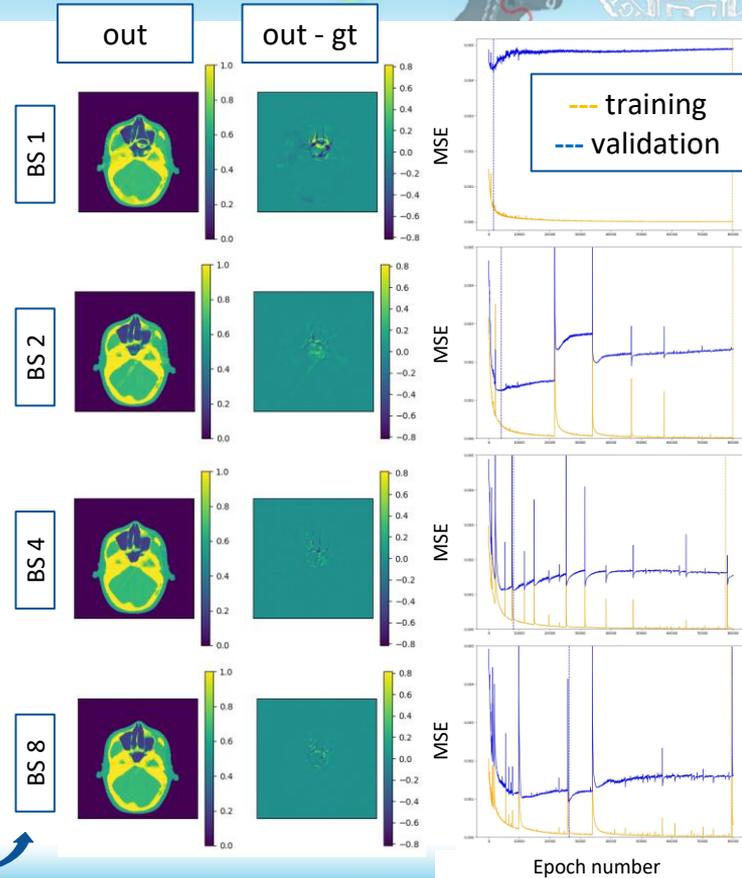
“Subset training” batch size



- Based on deterministic “subset training”, the batch size (BS) is set to 4
- BS 1: Overfitting (model capacity > size of the training data), inaccurate gradient (relative to the full-dataset gradient) but deterministic
- BS 2: Overfitting, improved generalization gap, more accurate but unstable gradient
- BS 4, 8: Overfitting, stable generalization gap



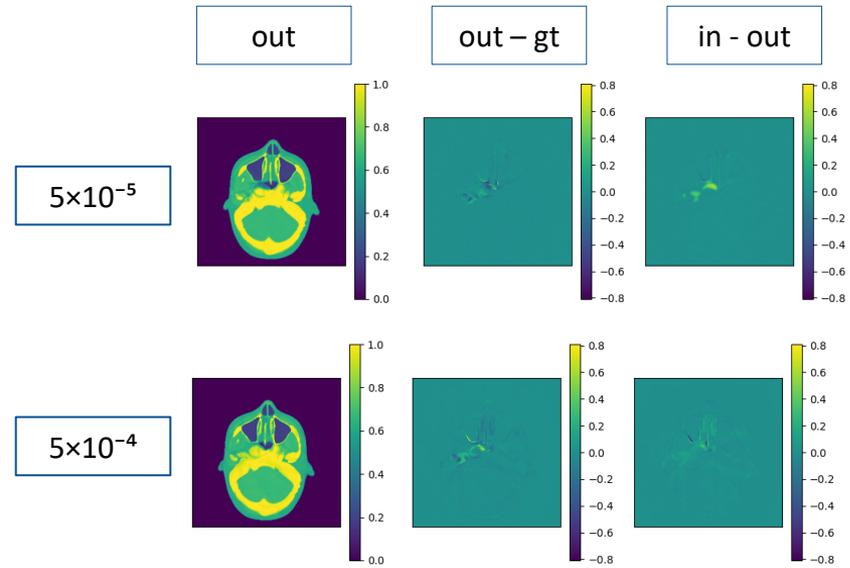
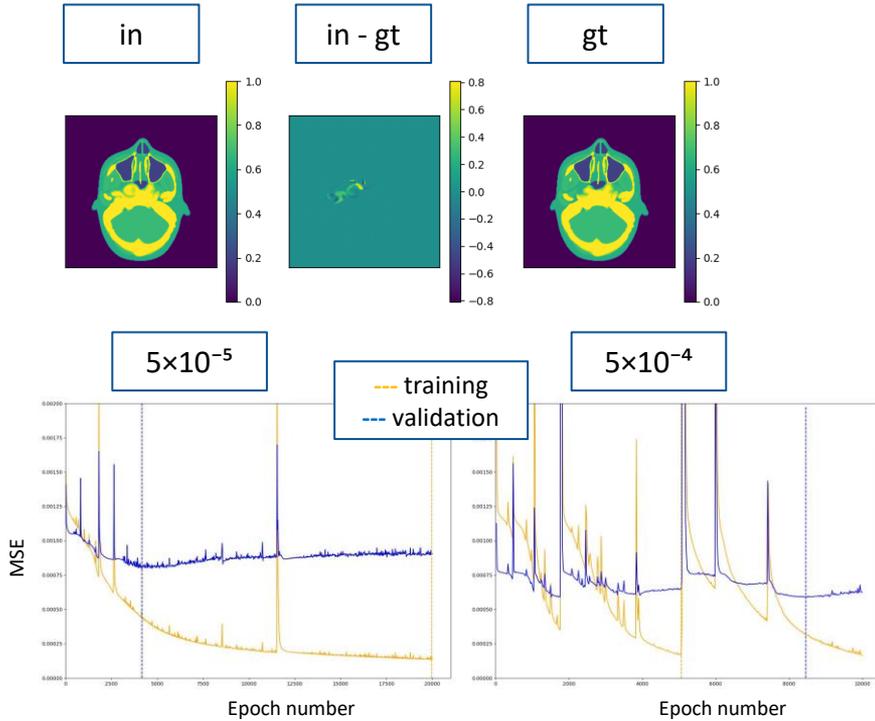
@minimum
validation MSE



“Subset training” learning rate



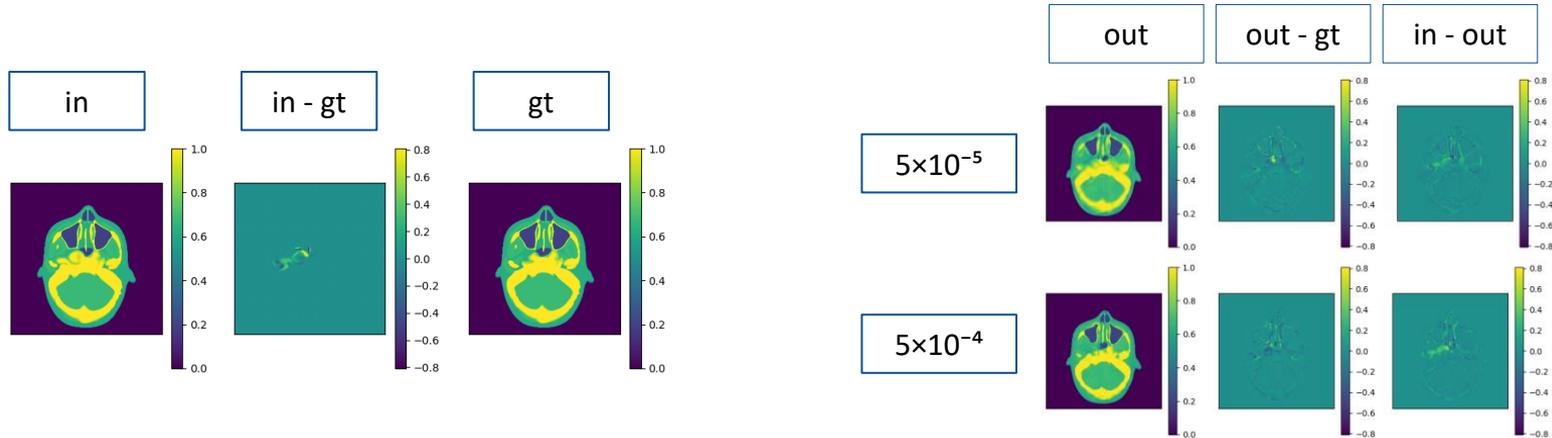
- The learning rate for the LPD network is reduced (from 5×10^{-4} to 5×10^{-5}) to achieve more stable training



“Subset training” learning rate



- The learning rate of the U-LPD network is increased (from 5×10^{-5} to 5×10^{-4}) to achieve better training



- Gradients propagate through both the LPD blocks and the U-Net, and a higher learning rate counteracts gradient attenuation caused by the increased network depth and normalization layers
- The training and validation curves of the U-LPD network exhibit high stability and low sensitivity to variations in the learning rate, with signs of overfitting observed for values greater than 1×10^{-3}

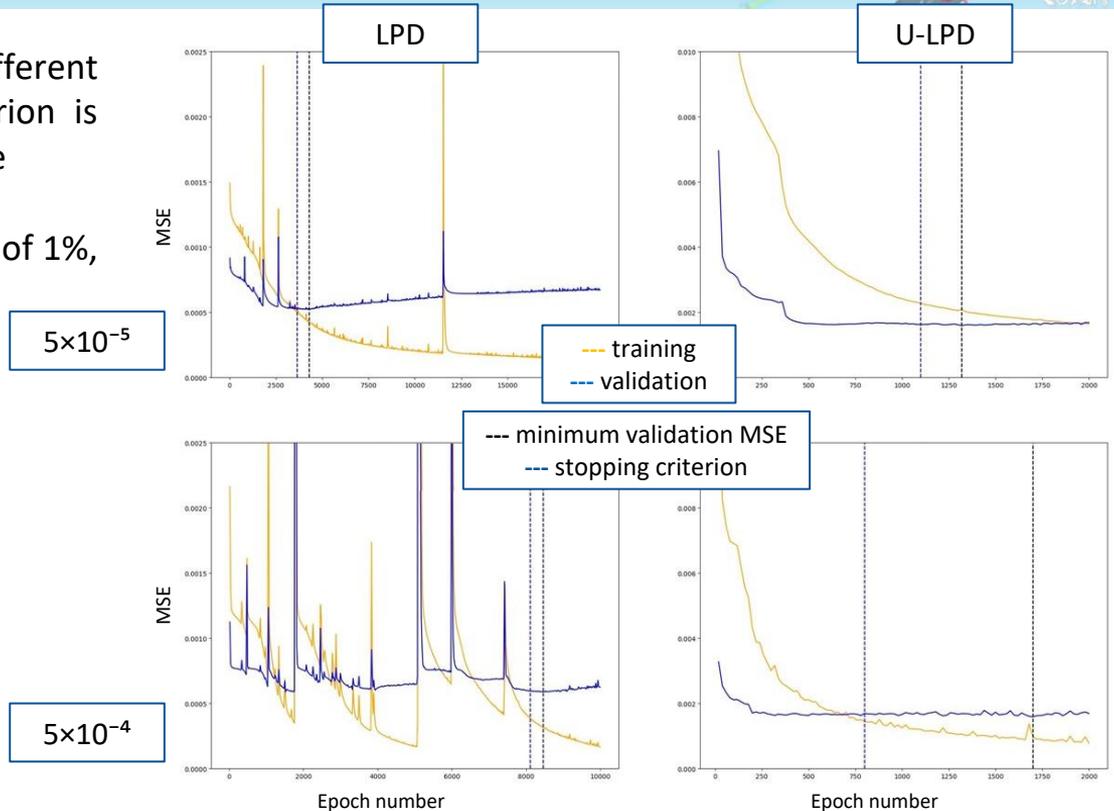
“Subset training” stopping criterion



- The LPD and U-LPD are trained using different learning rates, since the stopping criterion is generally independent of the learning rate
- Window size of 600 epochs, tolerance of 1%, patience of 800 epochs

@“example slab”

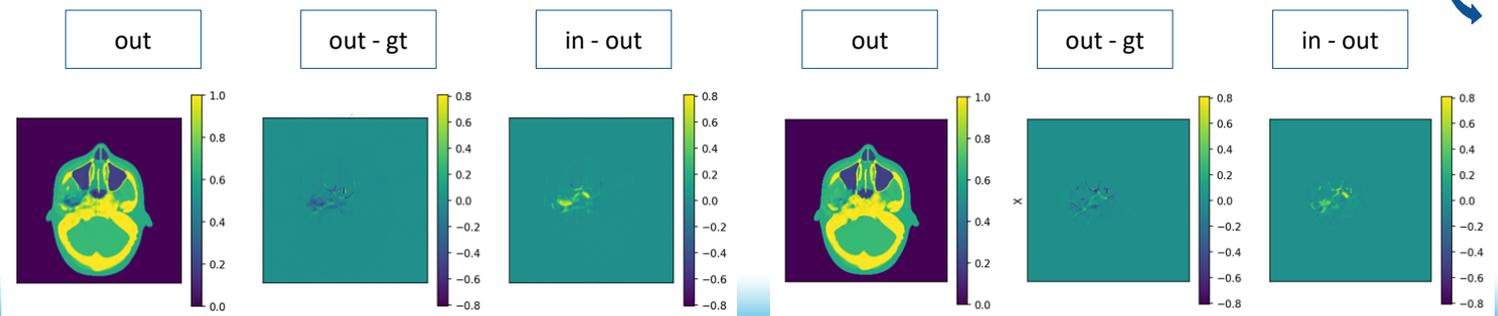
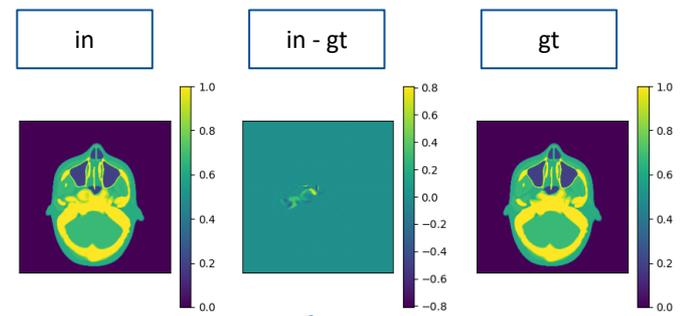
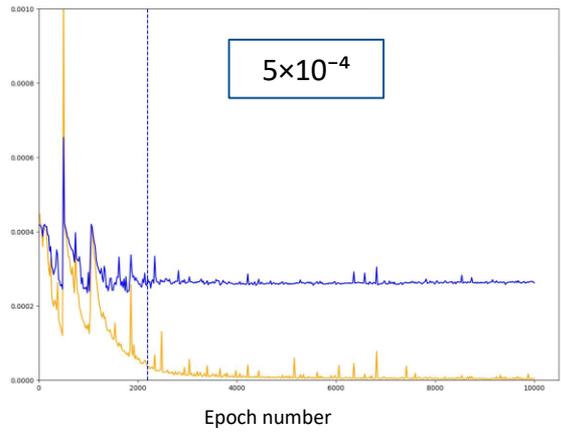
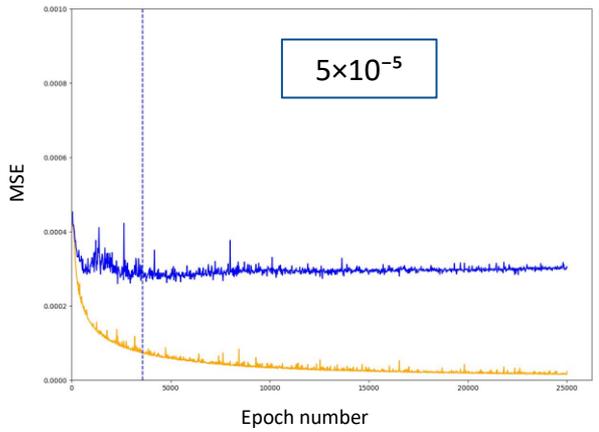
	LR	RMSE	SSIM
LPD	5×10^{-5}	2.46×10^{-2}	0.870
	5×10^{-4}	2.60×10^{-2}	0.854
U-LPD	5×10^{-5}	2.68×10^{-2}	0.829
	5×10^{-4}	2.91×10^{-2}	0.804





- The LPD is trained in “full training”

-- training
-- validation

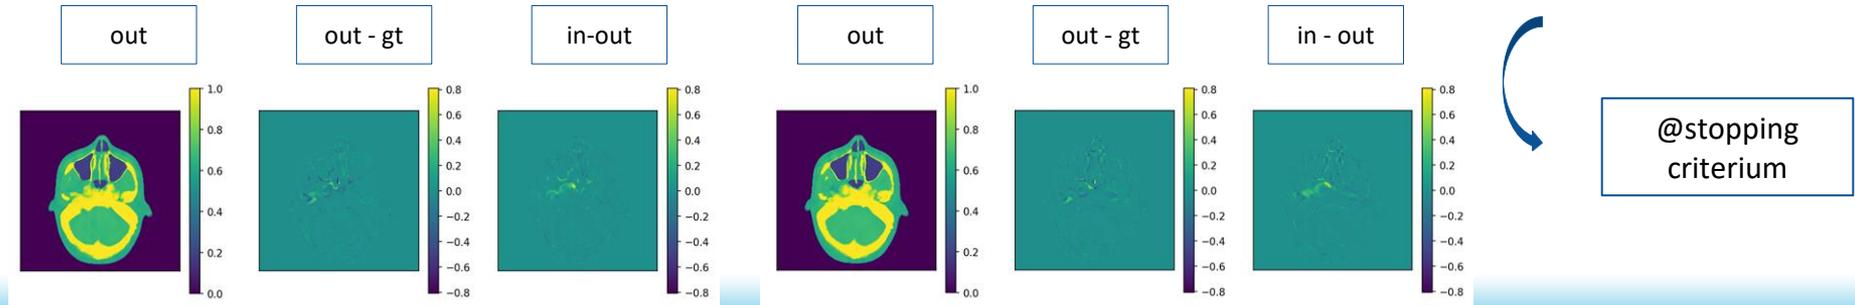
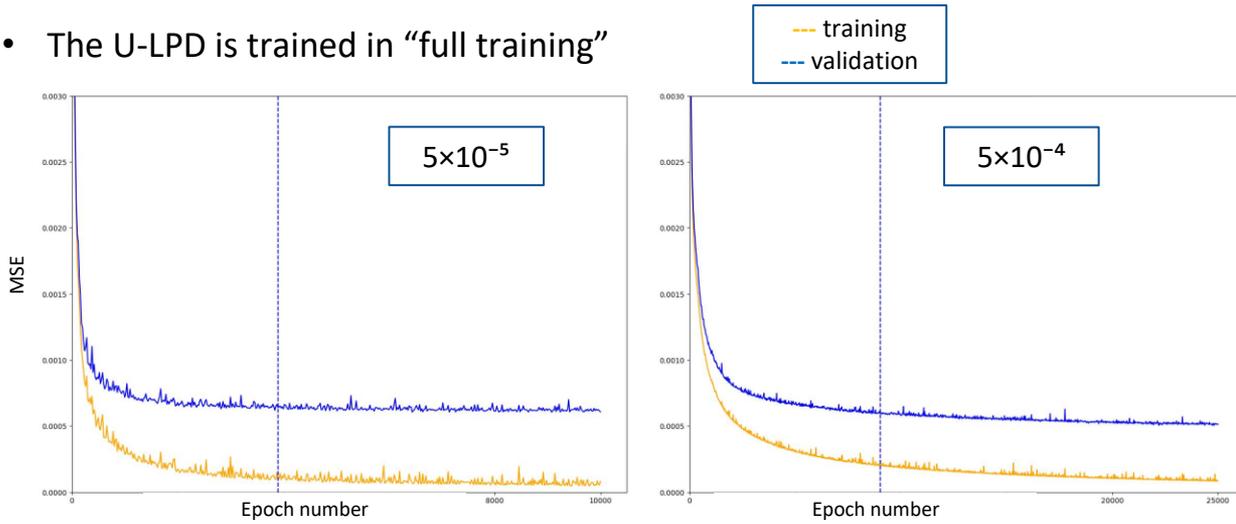


@stopping
criterium

“Full training” U-LPD



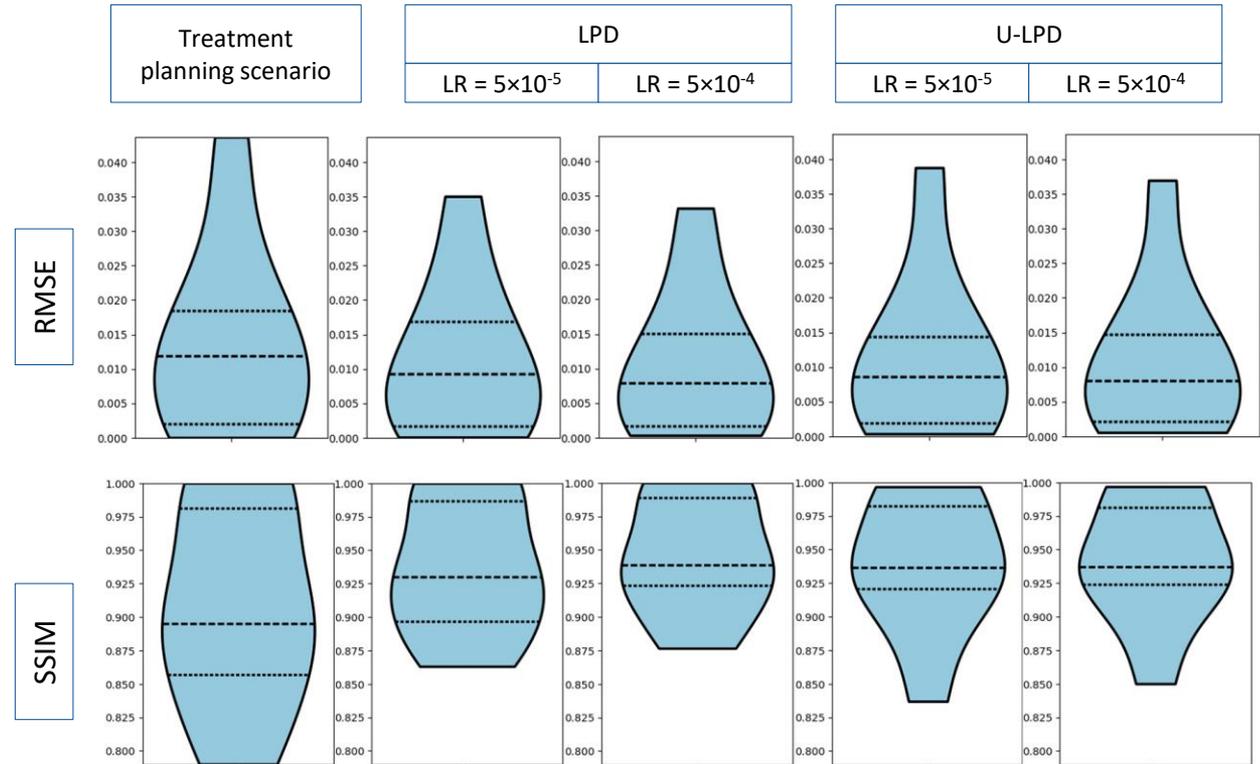
- The U-LPD is trained in “full training”



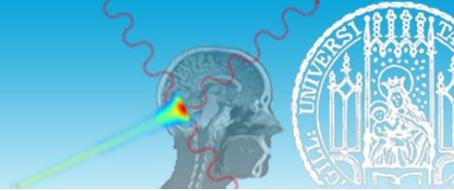
Task-specific evaluation metrics validation dataset



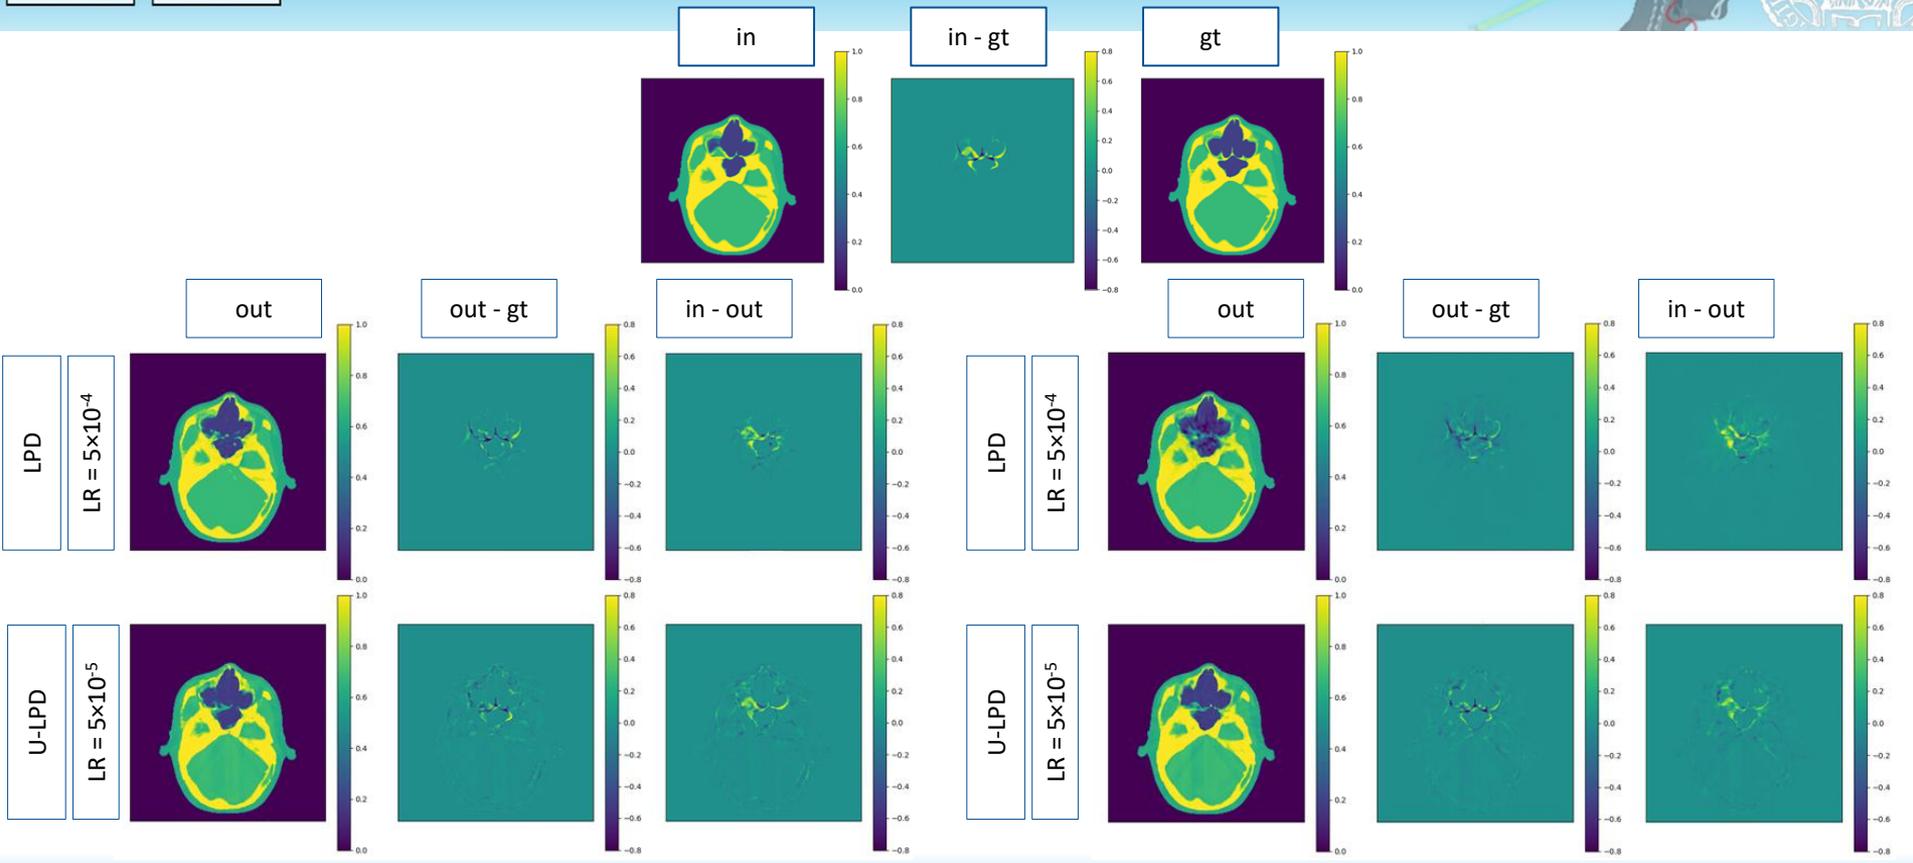
- The metrics suggest that both the LPD and U-LPD networks are capable of capturing the **transformation**
- The U-LPD network performs **comparably** to the LPD network while demonstrating more **stable** training and validation curves



Task-specific evaluation metrics

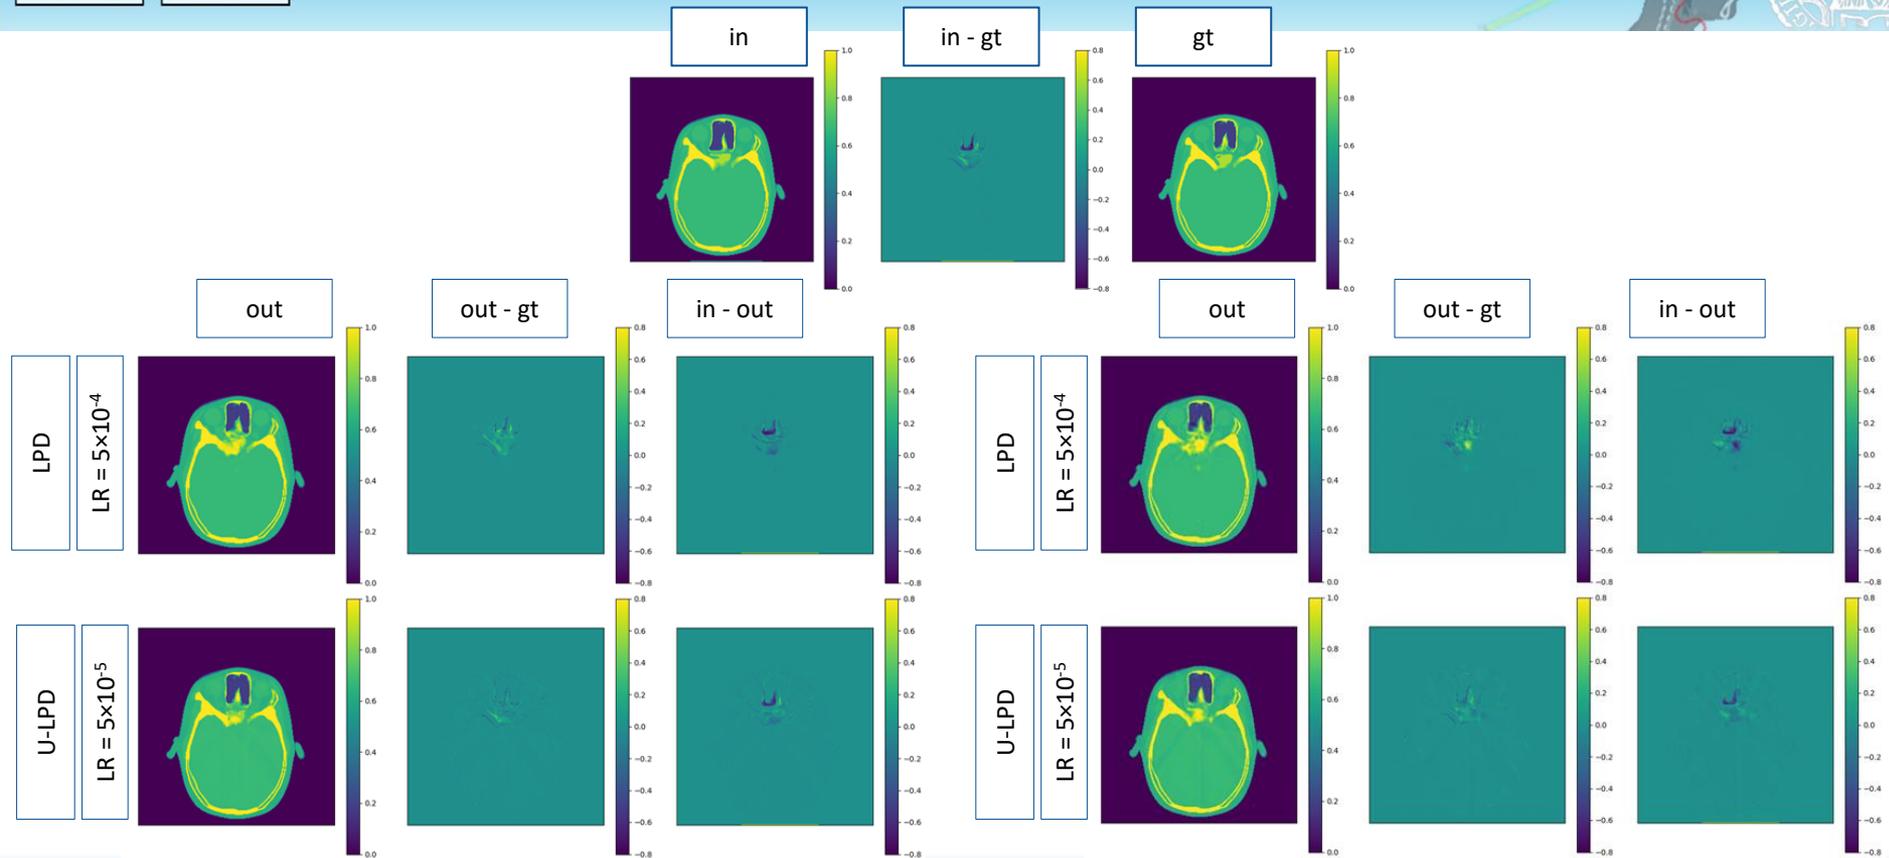


test dataset



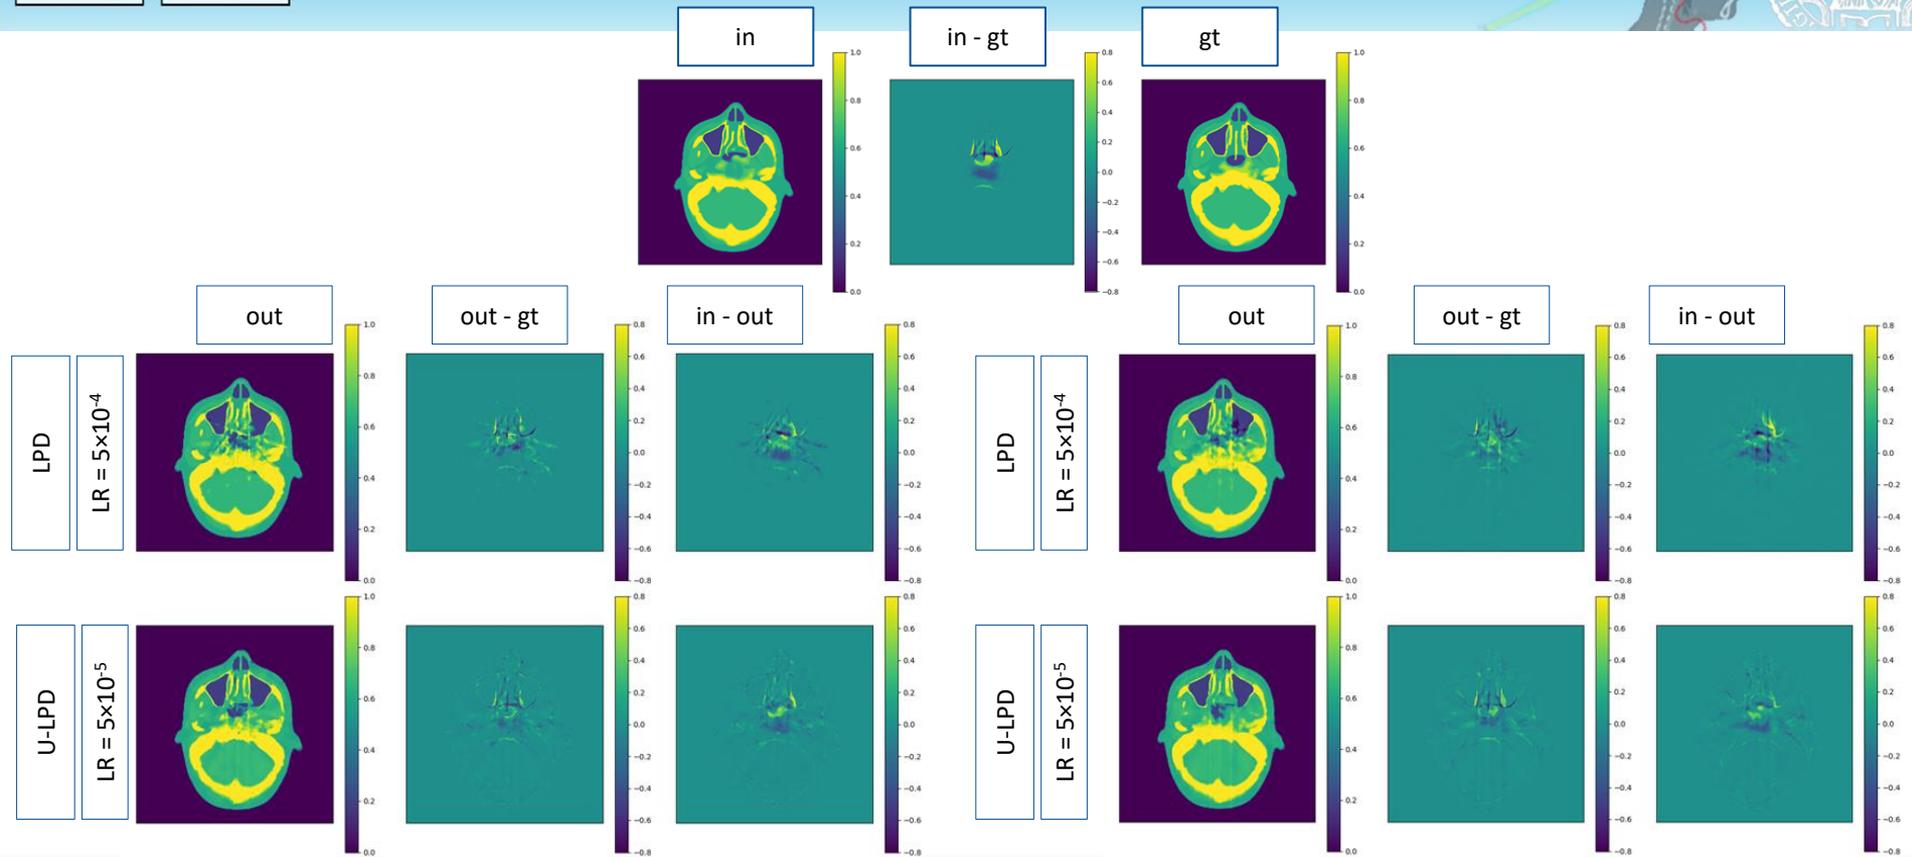
Task-specific evaluation metrics

test dataset



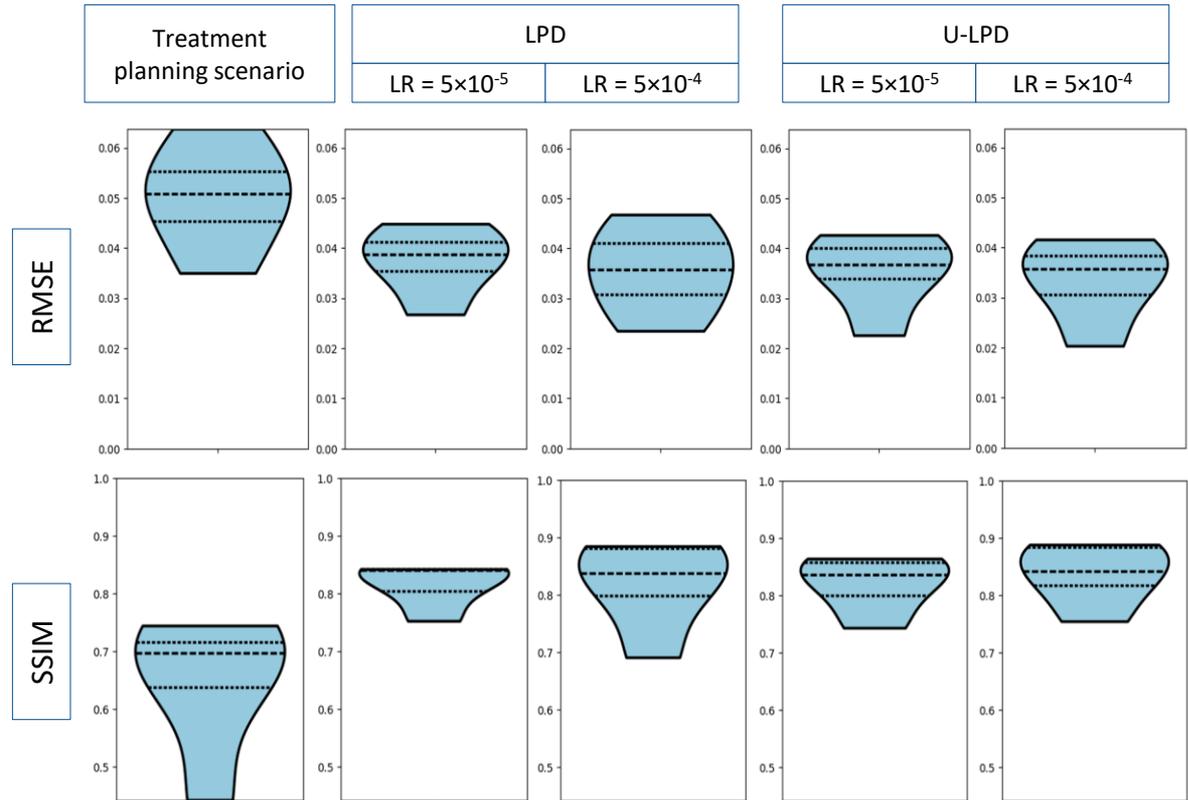
Task-specific evaluation metrics

test dataset





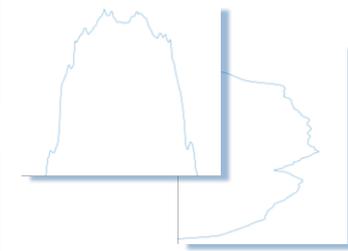
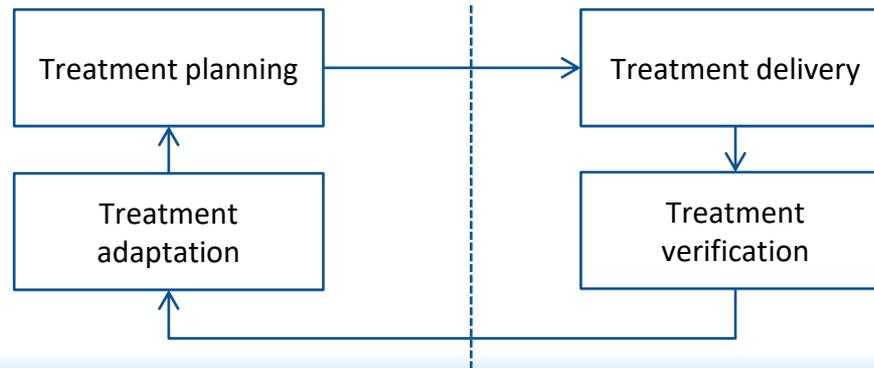
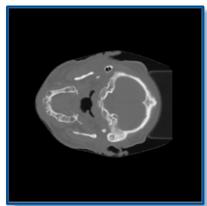
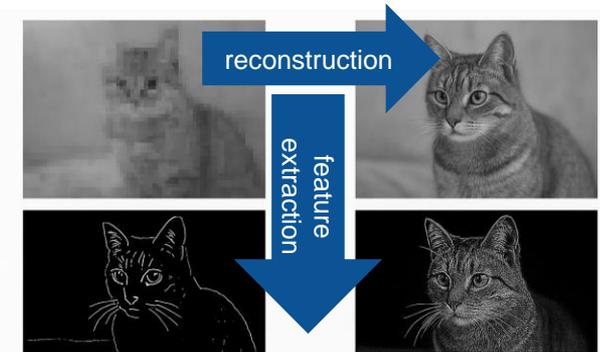
- The deformations applied to the phantom in the **test dataset** result in greater **intensity differences** relative to the ground truth, and this is reflected in the evaluation metrics
- The metrics confirm the results observed on the validation dataset, highlighting both the ability of the two architectures to capture the **deformation**



Treatment adaptation based on sparse ion radiographies



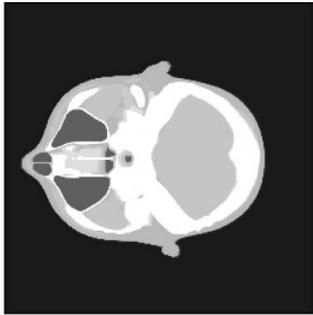
- Further investigation of the architecture → Explore optimal placement and depth of U-Nets
- Early U-Nets: may capture features on coarse reconstruction
- Late U-Nets: may overcomplicate an already refined reconstruction
- Combination with the calibration inaccuracies → Explore joint inference of “image and motion”



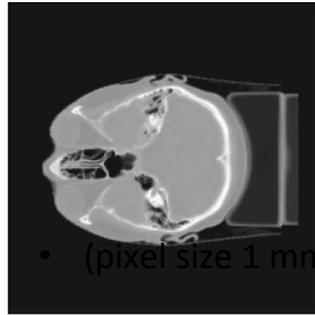
Treatment adaptation based on sparse ion radiographies



- Monte Carlo simulations on patient data → Explore patient-specific and detector-specific realistic data

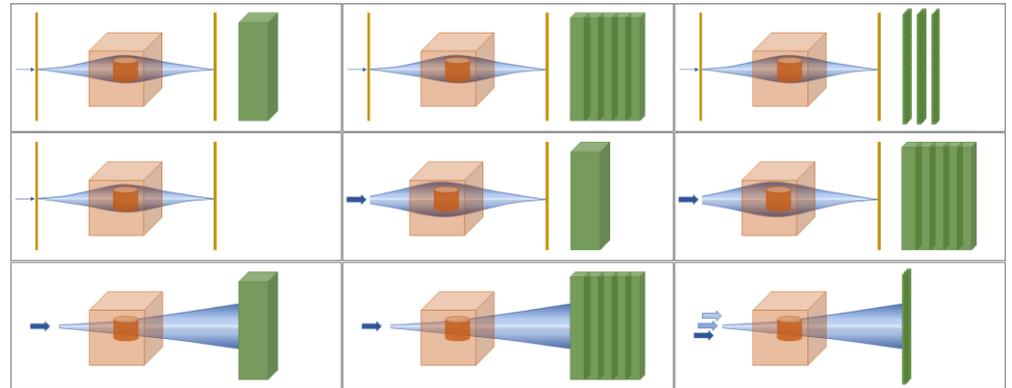
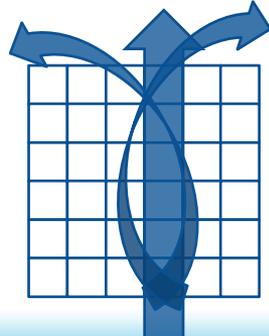
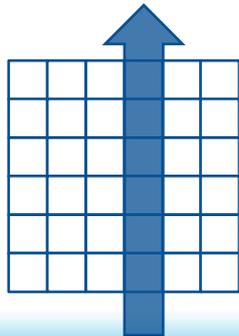


phantom data



(pixel size 1 mm)

clinical data



Paganelli, C., Gianoli, C., & Knopf, A. (2024). Imaging in Particle Therapy: Current practice and future trends. IOP Publishing.



Thank you for your attention

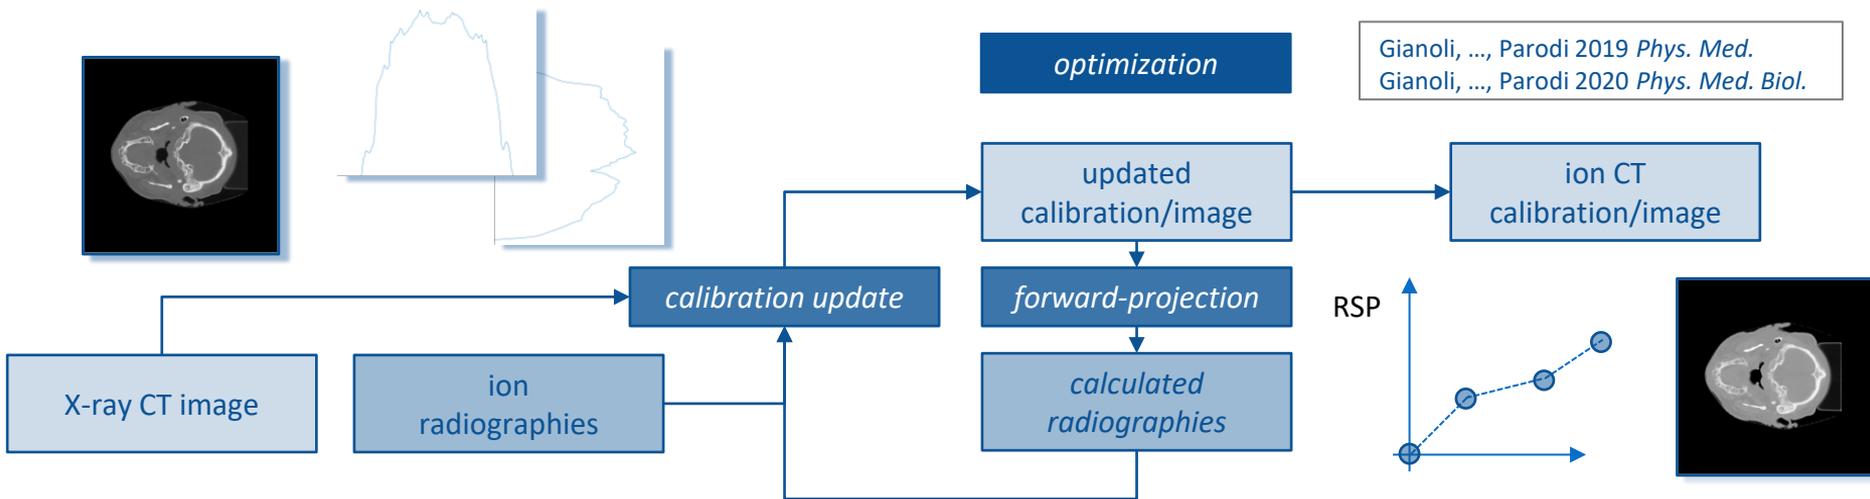
Additional slides



The role of ion imaging in treatment planning

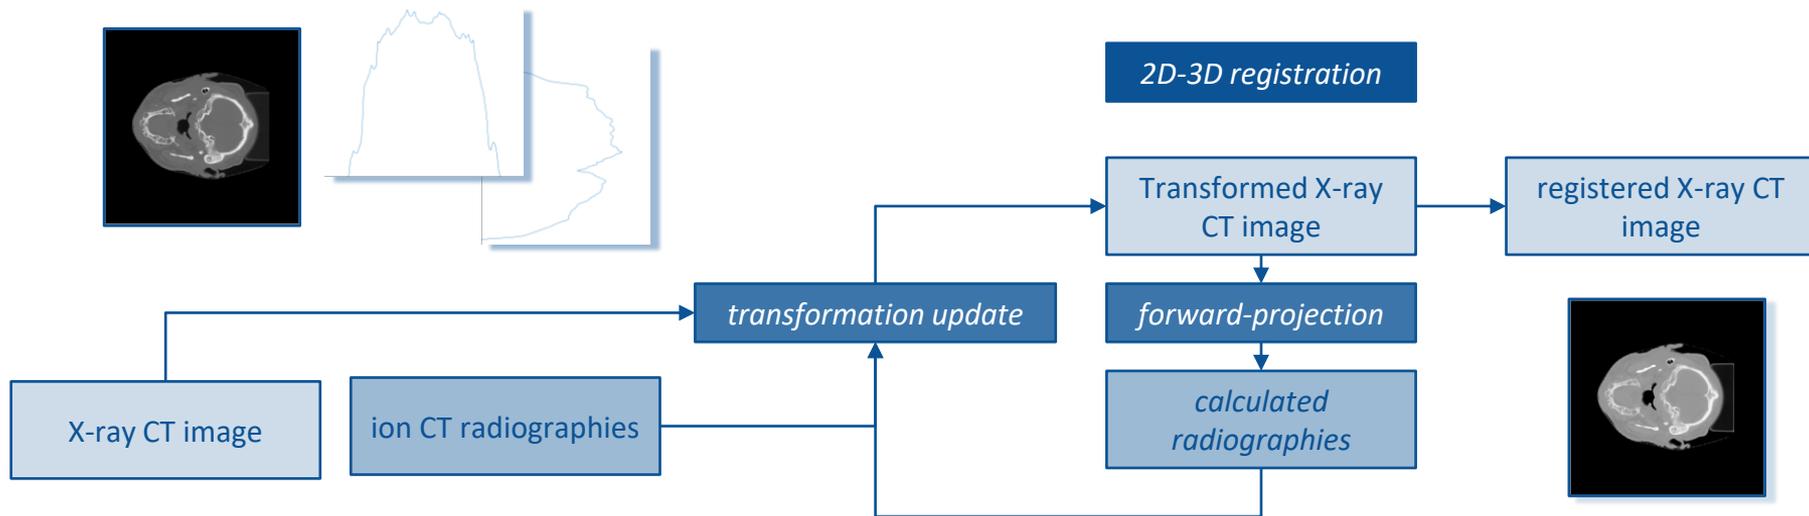


- The optimization of the X-ray CT image **calibration** using sparse ion radiographies simplifies the tomographic image reconstruction problem, with the aim of reducing **calibration inaccuracies**
- The prior information of the calibrated X-ray CT image and the reduced number of unknowns (i.e., the points of the calibration curve) are hypothesized to serve the solution of the ill-posed inverse problem





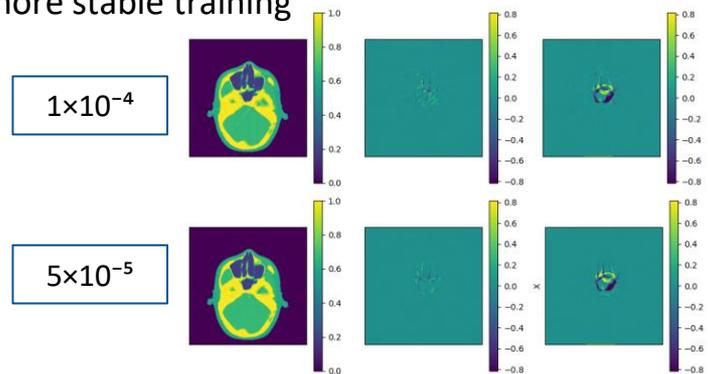
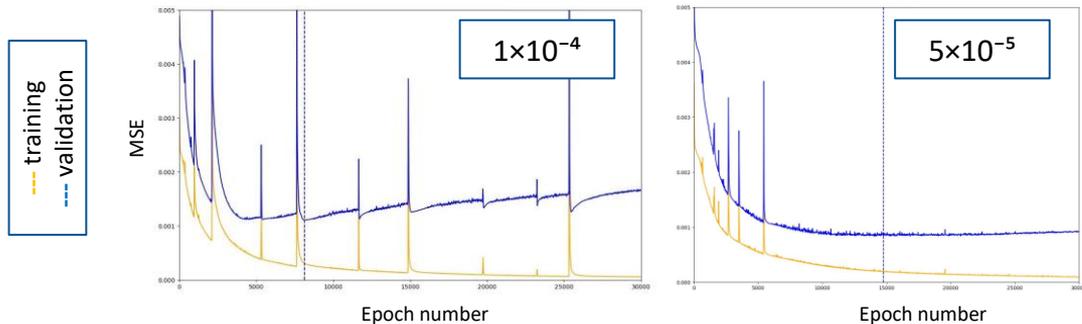
- A formulation of the **image registration** problem based on sparse ion radiographies has been firstly proposed, aiming to minimize anatomical changes in adaptive radiation therapy
- The forward-projection operator is exploited to transform the X-ray CT image according to the sparse ion radiographies



“Subset training” learning rate



- LPD: The learning rate is reduced (from 1×10^{-4} to 5×10^{-5}) to achieve more stable training



- U-LPD: The learning rate is increased (from 5×10^{-5} to 5×10^{-4}) to achieve more stable training

