Research Square

Preprints are preliminary reports that have not undergone peer review.
They should not be considered conclusive, used to inform clinical practice,
or referenced by the media as validated information.

# Enhancing the Spatial Resolution of Dose Measurements Using a Super-Resolution Geometry-Informed Neural Network

**Kamila Kalecińska**

kamila.kalecinska@agh.edu.pl

AGH University of Krakow

**Bartłomiej Rachwał**

AGH University of Krakow

**Jakub Hajduga**

AGH University of Krakow

**Jan Kapusta**

AGH University of Krakow

**Daria Kokot**

AGH University of Krakow

**Joanna Dudała**

AGH University of Krakow

**Tomasz Fiutowski**

AGH University of Krakow

**Damian Kabat**

The Maria Sklodowska-Curie National Research Institute of Oncology

**Bartłomiej Łach**

AGH University of Krakow

**Bartosz Mindur**

AGH University of Krakow

**Jakbub Moroń**

AGH University of Krakow

**Piotr Wiącek**

AGH University of Krakow

**Tomasz Szumlak**

AGH University of Krakow

---

**Article**

**Keywords:**

**Posted Date:** July 3rd, 2026

**DOI:** https://doi.org/10.21203/rs.3.rs-10091556/v1

**License:** This work is licensed under a Creative Commons Attribution 4.0 International License.
Read Full License

**Additional Declarations:** No competing interests reported.

# Enhancing the Spatial Resolution of Dose Measurements Using a Super-Resolution Geometry-Informed Neural Network

**Kamila Kalecińska[1,*], Bartłomiej Rachwał[1], Jakub Hajduga[1], Jan Kapusta[1], Daria Kokot[1], Joanna Dudała[1], Tomasz Fiutowski[1], Damian Kabat[2], Bartłomiej Łach[1], Bartosz Mindur[1], Jakub Moroń[1], Piotr Wiącek[1], and Tomasz Szumlak[1]**

[1]AGH University of Krakow, Faculty of Physics and Applied Computer Science, al. A. Mickiewicza 30, 30-059 Krakow, Poland
[2]Maria Skłodowska-Curie National Research Institute of Oncology Krakow Branch, ul. Garncarska 11, 31-115 Krakow, Poland
[*]kamila.kalecinska@agh.edu.pl

## ABSTRACT

Accurate verification of photon radiotherapy treatment plans requires dose measurements with spatial resolution sufficient to resolve steep dose gradients. However, for hardware-based three-dimensional phantom detectors, the achievable measurement resolution is fundamentally limited by detector size, resulting in dose maps that are substantially coarser than the voxel grids used by treatment planning systems. Here, we present a DICOM RT-Plan-Guided Super-Resolution (SR) framework that reconstructs voxel-scale ($1\,mm^3$) dose distributions from cell-scale ($1\,cm^3$) measurements acquired with the D3DF reconfigurable phantom. The framework incorporates a Geometric Aperture Moment (GAM), a treatment-plan-derived descriptor of multileaf collimator (MLC) geometry, together with a depth-weighted beam mask as auxiliary inputs to a UNet++ model, enabling recovery of spatial details that are not directly observable in the detector measurements. The model is trained on paired cell- and voxel-resolution datasets generated using a Geant4-based Monte Carlo platform and evaluated using both peak signal-to-noise ratio (PSNR) and the clinically established Gamma Passing Rate (GPR). In a dedicated test where identical cell-level dose measurements corresponded to different underlying MLC configurations, the framework successfully reconstructed the correct voxel-scale dose geometry, demonstrating that RT-Plan-derived information provides complementary spatial information beyond the measured dose alone. These results indicate that routinely available DICOM RT-Plan information can substantially improve dose reconstruction fidelity and may enable high-resolution treatment-plan verification without the need for either denser detector hardware or high-resolution voxel scoring in Monte Carlo simulations.

## 1 Introduction

The development of a real-time dose deposition monitoring system within a defined 3D region can offer numerous benefits in dosimetry and quality assurance of treatment plans. Each treatment plan presents a spatial dose distribution, often with highly complex shapes and significant local dose gradients. Moreover, respiratory-gated radiotherapy is a commonly used technique, and it represents another potential application area for a 3D phantom capable of real-time measurement, particularly for quality assurance in such treatments.

To address these needs, our team, Dose3D Future (D3DF), has developed a scalable detection system for evaluating dose distributions in a reconfigurable 3D phantom intended for use in photon radiotherapy treatment planning[1–4]. A schematic representation of the system is shown in Fig. 1. Throughout the development process, attention was also given to ensuring compliance with current Digital Imaging and Communications in Medicine (DICOM) standards. The prototype versions of the phantom's active component currently consist of cubic cells with a 10 mm side length. The use of Machine Learning (ML) methods aims to increase the spatial resolution of measured dose distributions to match those produced by treatment planning systems.

### 1.1 Image Resolution Within the DICOM Standard

Medical images, such as Computed Tomography (CT) scans and DICOM Radiotherapy Dose Data (RT-Dose), are stored in the standardised DICOM format. In addition, anatomical structures, including the tumor, are defined in files DICOM Radiotherapy Structure Set (RT-Struct), which provide essential segmentation data for treatment planning. While the CT image resolution in axial plane (slice thickness) typically has a value from 1 mm to 5 mm

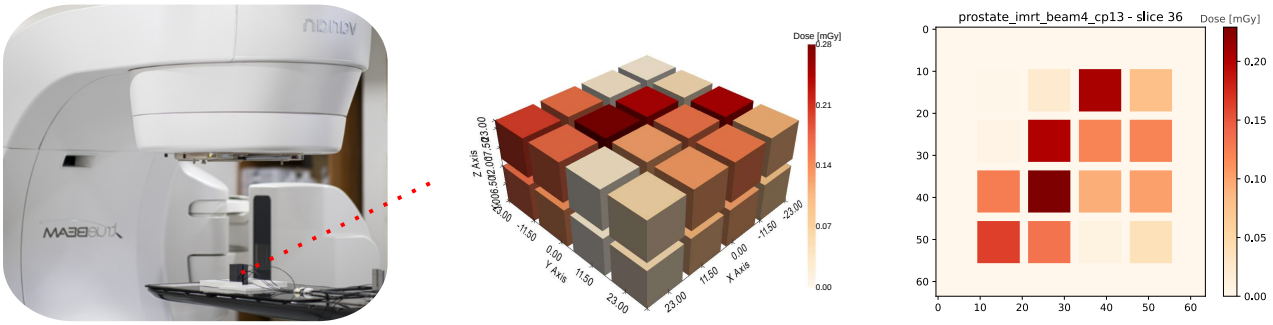



**Figure 1.** Schematic representation of the linac gantry and the spatial dose detector phantom from the D3DF project, used in Monte Carlo simulations in this study. The phantom consists of a $1\,\mathrm{mm}^3$ active volume discretized into $4 \times 4 \times 2$ scoring cells for dose estimation. The right panel shows a single CT slice of the detector mapped onto a $64 \times 64$ CT grid derived from treatment planning data — the beam configuration and control point define the corresponding mask.

and in-plane resolution (pixel size) is about $1\,\mathrm{mm}^2$. The resolution (voxel size) used in Treatment Planing System (TPS) is commonly valued around 2.5 mm in every direction.

The reported resolution values of the information provided by the TPS indicate that dose calculations are performed at relatively coarse spatial resolutions. Consequently, several techniques - such as interpolation, resampling, and subdivision - are employed to enhance the resolution of dose distribution representation. These refinements, automated within TPS, are crucial because RT-Struct data, used alongside RT-Dose data, serve as the basis for computing essential metrics such as the Dose Volume Histogram (DVH), a fundamental tool for assessing dose distribution across different tissues and volumes[5–7]. Improving the quality of imaging and dose calculation methods directly enhances the accuracy of dose delivery to the Gross Tumor Volume (GTV), which corresponds to the tumor identified in imaging studies. Precise targeting of this volume, along with careful consideration of other patient volumes, is essential for effective and personalized treatment planning.

However, given the current spatial resolution of our prototype measurement system, where each detector element has a large active volume it is not feasible to apply the aforementioned algorithms to accurately resolve dose variations at the scale comparable to TPS voxel grids.

### 1.2 Objective of the Study

The novel contributions of this work are: (i) a UNet++ architecture applied for volumetric dose super-resolution from $1\,\mathrm{cm}^3$ to $1\,\mathrm{mm}^3$; (ii) the Geometric Aperture Moment (GAM), a per-voxel geometric encoding of the MLC configuration extracted from the RT-Plan, used as additional input to the network aimed at enhancing its performance; (iii) a controlled experimental design in which cell-level dose is held constant across different MLC configurations, isolating the contribution of plan-guided features at the sub-cell scale; and (iv) a dual evaluation combining standard image-similarity metrics with the clinically established Gamma Passing Rate.

## 2 Methodology

### 2.1 Plan Guided Neural Network

Incorporating physical knowledge directly into the learning process, Physics-Informed Neural Networks (PINNs) have emerged as a prominent approach[8,9]. The hallmark of these networks is their loss function, which typically consists of two components: a data loss term, representing the standard error metric of neural networks, and a physical loss term, derived from the governing physical laws of the system under study. However, our research introduces a fundamentally different approach. Rather than embedding physical principles directly into the neural network, we adopt a standard ML framework enhanced by a novel geometric rescaling of the feature distribution. This rescaling operates independently of the underlying physical processes and leverages information extracted from the DICOM Radiotherapy Plan Data (RT-Plan), a standard format in radiotherapy planning. This methodology allows us to retain the flexibility of standard ML techniques while incorporating domain-relevant geometric insights, bridging the gap between physics-based modeling and data-driven approaches.

## 2.2 Dataset production

All experiments and analyses presented in this work are based on datasets generated using Monte Carlo (MC) simulations, ensuring physically accurate modeling of radiotherapy dose deposition. An in-house Geant4 Radiotherapy (G4RT) simulation platform, based on the Geant4 Software Toolkit (Geant4)[10,11] has been utilized for training and evaluation data production. The Monte Carlo (MC) simulation serves as the foundation for accurately modelling dose deposition within an artificial patient of the structured grid-cell system (4x4x2).

The dataset utilized in this work comprise paired input (sparse-grid; cell-level) and target (high-resolution-grid, voxel-level) samples which are generated for a set of Multileaf Collimator (MLC) configurations extracted from DICOM Radiotherapy Plan (DICOM RT-Plan) files obtained from clinical practice at the Maria Skłodowska-Curie National Research Institute of Oncology, Krakow Branch. Subsequently, a Control Point (CP) is defined by a single MLC configuration together with the corresponding gantry angle orientation. In the present study, only simulations with a gantry angle of 0° were considered, corresponding to the Z-axis direction in the gantry coordinate system.

As part of the data preprocessing pipeline, both input and target samples are mapped onto a common grid (CT-like data) corresponding to the CT-tube geometry (in this work, it's corresponds to the 6.4 $cm^3$ volume of the 64x64x64 resolution). Combining CT-grid and detector-level scoring, the simulation bridges the gap between experimental D3DF imaging data and the accuracy of dose reconstruction from the TPS. The complete data flow used in this work is presented in Fig. 2, where patient data from the RT-Plan are used in MC simulations to generate both input (cell-level) and target (voxel-level) datasets. After preprocessing, the data are mapped onto the CT grid, and the Geometric Aperture Moment (GAM) is evaluated for each voxel. The resulting MC-generated cell-level and voxel-level datasets provide geometry-informed input and ground-truth target data for model training. During inference, the model takes cell-level data together with the control point configuration as input and reconstructs super-resolution dose distributions aligned with the CT grid.

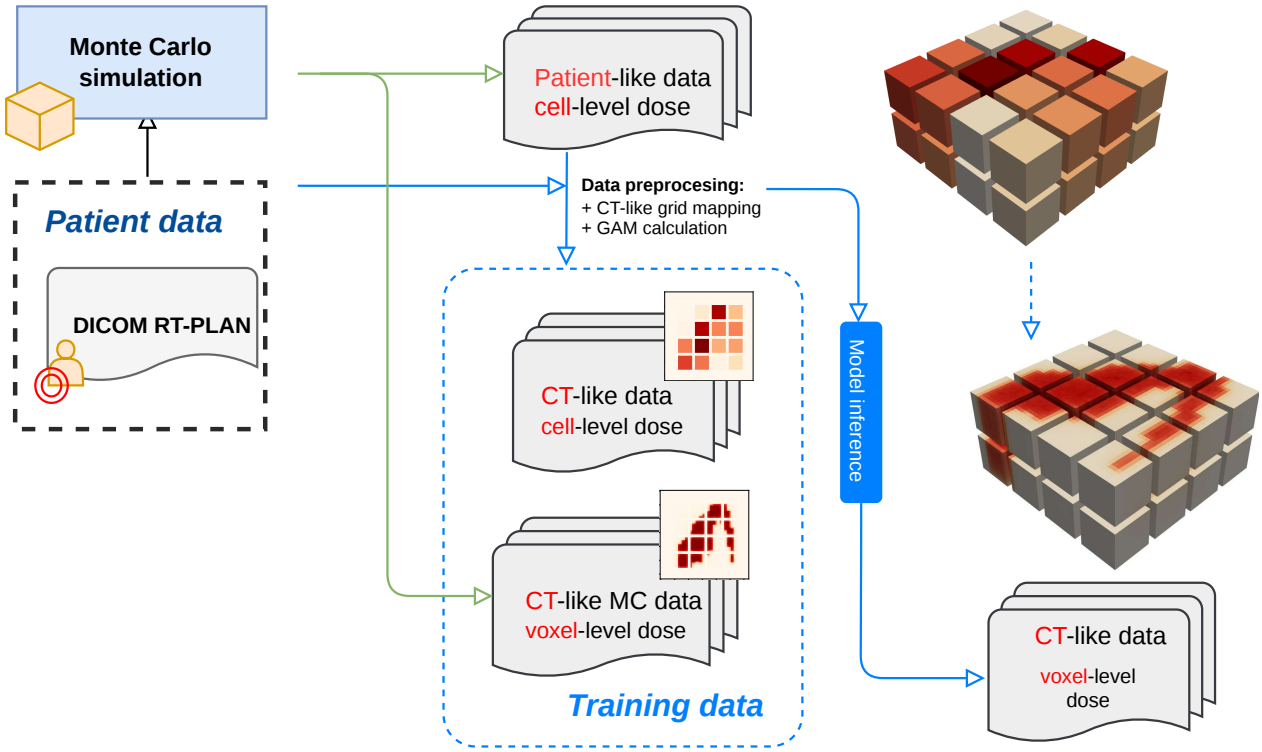



**Figure 2.** Schematic representation of the data-processing pipeline for preprocessing, model training and inference.

## 2.3 Deep Learning (DL) Model fo Super-Resolution (SR)

The Super-Resolution (SR) has been a well-known issue in image processing and ML. Due to spatial dependency in our data, they can be treated as 3D images, so the natural choice of technique was Convolutional Neural Networks

(CNNs). We utilize the UNet++[12], which is built on the core of U-Net[13]. Since it was invented for medical purposes, we decided to use a medical framework - *MONAI*[14], which contains convenient implementations of it.

The UNet++ is a denser and more advanced version of basic U-Net. The largest innovation is the fact that in this architecture, the encoder and decoder subnetworks are connected through a series of nested, dense skip pathways instead of the original skip connections.

### 2.4 Training Process Details

For a single beam direction (i.e., gantry angle), the dataset comprising 512 elements was split into training, validation, and test subsets in a 7:2:1 ratio.

The *CombinedLossFunction* was constructed by combining two core elements: Mean Squared Error (MSE) and Structural Similarity Index Measure (SSIM)[15] Loss (Equation 1), with equal weighting (50/50) between them. Combining different loss functions is a common practice in the field of ML, and our goal was to balance low-level fidelity (MSE) and high-level structural similarity (SSIM) )in our results.

$$\mathcal{L}_{SSIM}(x,y) = 1 - SSIM(x,y) \tag{1}$$

A deep supervision mechanism was employed to supervise each layer of the feature reconstruction process, incorporating weighted loss values from each resolution level into the final loss function. The weighting strategy was experimentally selected.

### 2.5 The TPS Plan-Guided Rescaling Descriptor

The TPS is a specialized software used in radiation therapy to design and optimize treatment plans. It provides detailed geometrical information, including the beam shape (see schematic view in Fig. 3a) and its orientation (i.e., gantry angle) relative to the patient geometry and segmentation. The system employs individual CP to capture the dynamic modulation of MLC leaf positions. These variations reflect the implementation of Intensity-Modulated Radiotherapy (IMRT), in which each beam is delivered as a sequence of CP, each defined by a specific MLC leaf configuration. All of this information is stored in the DICOM RT-Plan. From it, we derive a per-voxel geometric descriptor, the Geometry Aperture Moment (GAM), that links the plan-level beam geometry to the spatial dose distribution and provides conditioning information for the SR model. The GAM is constructed as a second moment of the MLC leaf-tip positions about the local beam axis, computed independently at each voxel. By definition, the GAM can be determined at each point within the spatial domain of interest, which, in our study, corresponds to the centers of individual voxels. The formula is defined in Equation 2.

$$GAM = \sum_{\text{leaf}=1}^{64} ((\lambda_{\text{leaf}}^{left})^2 + (\lambda_{\text{leaf}}^{right})^2), \tag{2}$$

where $\lambda_{\text{leaf}}^{\text{left}}$ and $\lambda_{\text{leaf}}^{\text{right}}$ represent the leaf bearing angle for the left and right sides of the MLC, respectively. The $i$-th component of the sum refers to the relation of the given *point-of-interest* ($P$) and the $i$-th MLC leave position as defined in the Equation 3. The formula has the structure of a second moment analogous to a moment of inertia of the leaf-tip distribution about the voxel-to-source axis with squaring weighting leaves by their angular distance from that axis.

$$\lambda_{leaf_i} = \angle(\mathbf{r}_{MLC}, \mathbf{r}_{leaf_i}), \tag{3}$$

which corresponds to the angle between vectors $\mathbf{r}_{MLC} = MLC_{centre} - P$ and $\mathbf{r}_{leaf_i} = MLC_{leaf_i} - P$. These vectors are constructed from given *point-of-interest* to the centre of the MLC and $i$-th leaf position, respectively (see also Fig. 3b). The $MLC_{centre}$ is a constant point of the setup which refers to the real Varian apparatus and $MLC_{centre} = (0, 0, -625.25)$.

The spatial distributions of the GAM for a representative patient and two MLC configurations are shown in Fig.4. Since the two distributions are visually difficult to distinguish, the difference between them (b-a) is presented in Fig.4e to highlight their discrepancies.

The GAM descriptor is being calculated for each CP and the target grid. To improve the numerical stability the min-max rescaling is being implied for all the features, including the dose distribution itself (eq. 4), where the min,

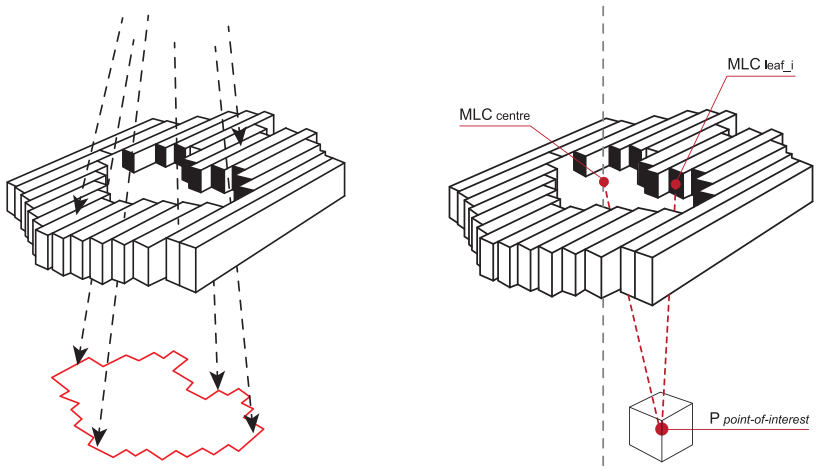



**Figure 3.** Schematic view of MLC configuration from a radiotherapy RT plan illustrating beam shaping into mask geometry (a). Reference points used in the analytical expression for the GAM spatial descriptor (b).

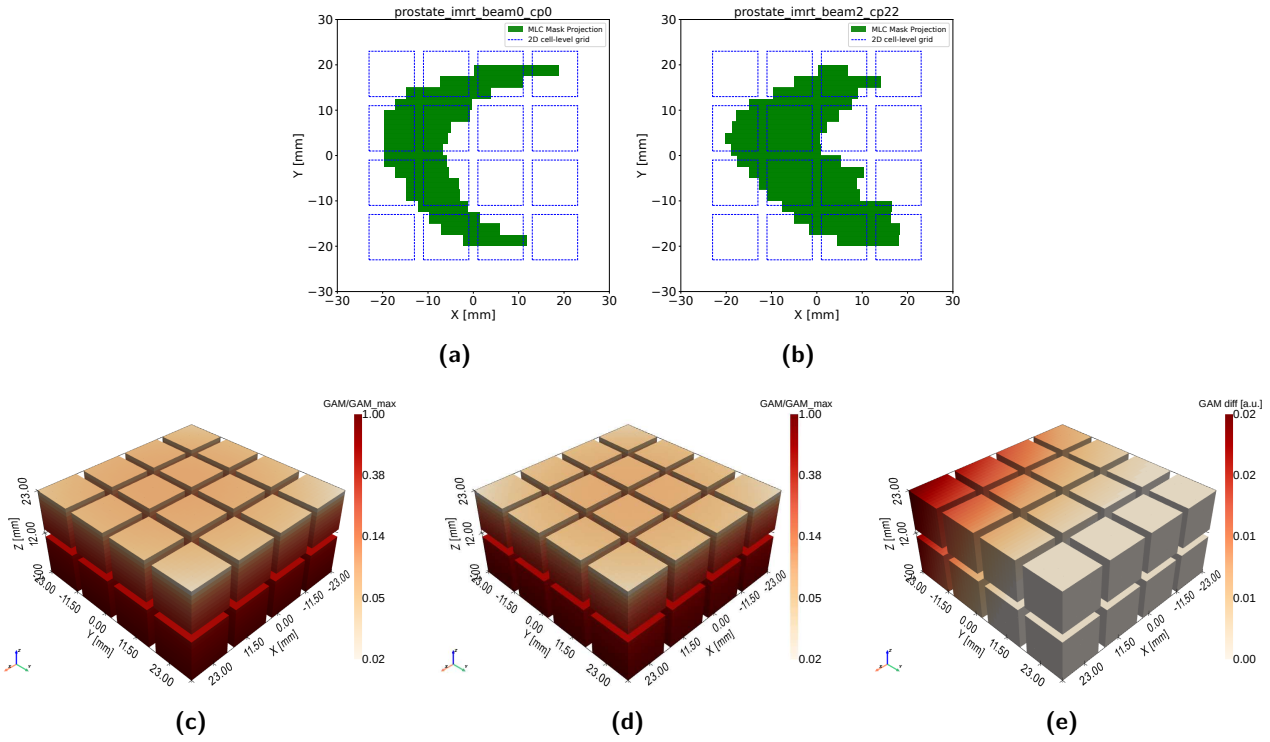



**Figure 4.** Examples of beam masks from individual control points (a,b) show dynamic MLC leaf modulation. The image canvas is scaled to match the actual mask dimensions (a fraction of the total MLC field). Corresponding GAM spatial distributions are shown in (c,d) with logarithmic scaling. The difference map (e = d–c) highlights changes between the two divisions.

max values are extracted from all values of the given feature (i.e., *GAM* or *dose*) in the given configuration (i.e. a given CP), while $max_{new} = 0.02$ and $max_{new} = 0.98$ rather than [0,1] in order to avoid saturation effects at the boundaries of the activation space, which may negatively affect gradient propagation and training stability.

$$Feature_{scaled} = \frac{Feature - min}{max - min} * (max_{new} - min_{new}) + min_{new} \tag{4}$$

### 2.6 Evaluation of Model Stability in Mask Orientation Recognition

Analysis of the training loss function and PSNR metrics suggests that the GAM factor does not substantially affect global model performance. However, its impact becomes apparent during inference of cell-scale dose distributions: the cellular distribution remains consistent, while voxel-level discrepancies arise from variations in beam configurations. To validate this scenario, two specially designed artificial beam masks were prepared to produce consistent dose distributions at the cell-level, despite differing voxel-level patterns (see Fig. 5). This setup enables a controlled assessment of the GAM factor's influence under conditions where voxel-level differences do not alter the aggregated dose at the cellular scale.

### 2.7 Evaluation Metrics

Standard ML metrics are initially employed to assess model performance. To evaluate accuracy in the context of field-specific criteria, a domain-specific evaluation method, the Gamma Passing Rate (GPR) is subsequently applied, as it is essential for determining the clinical relevance and precision of the model's predictions.

#### *2.7.1 ML evaluation*

The Peak Signal-to-Noise Ratio (PSNR), next to SSIM (as defined in section 2.4) is a commonly used metric to evaluate image similarity, especially for ML tasks like SR or image denoising. It measures the ratio between the maximum possible signal power and the power of the reconstruction error (represented by the mean squared error, MSE), expressed in decibels (dB). PSNR focuses solely on pixel-wise differences, without considering spatial context or clinical tolerances. We mostly use PSNR to evaluate the model during training.

$$\mathrm{PSNR} = 10\log_{10}\left(\frac{\mathrm{MAX}^2}{\mathrm{MSE}}\right) = 20\log_{10}\left(\frac{\mathrm{MAX}}{\sqrt{\mathrm{MSE}}}\right) \tag{5}$$

where: MAX is the maximum possible value of the signal,

#### *2.7.2 Domain-specific evaluation*

In this work, we use the **Gamma Passing Rate**, GPR, metric as a secondary evaluation measure for model performance on the test dataset. The GPR represents the percentage of points where $\gamma(p) \leq 1$ based on the **Gamma Index**, $\gamma$, a widely used metric in radiotherapy Quality Assurance (QA) for quantitatively comparing measured and calculated dose distributions.

The Gamma Index, introduced by Low et al. (1998)[16], simultaneously evaluates dose difference (DD) and distance-to-agreement (DTA) criteria. It combines these metrics into a single unitless value by normalizing each by its acceptance limit (e.g., 3% for $\Delta$D and 3 mm for DTA)[17]. Typical acceptance criteria require a Gamma Passing Rate (GPR) of 90 −95 % for common tolerances like 3%/3 mm or stricter standards like 3%/2 mm[18][19]. This metric is widely reported for dose verification in radiation therapy; however, acceptance criteria such as 3%/3 mm are not absolute and may be replaced by stricter tolerances depending on treatment complexity and clinical requirements.

## 3 Computational Setup

### 3.1 Data Preprocessing

In contrast to the original U-Net and UNet++ architectures, which are primarily designed for segmentation tasks where the input and output tensors naturally share the same spatial dimensions, the super-resolution problem requires reconstruction of a higher-resolution output from lower-resolution input data. To address this within the same architectural framework, the low-resolution input data were first voxelized and resampled to match the spatial grid of the target high-resolution output. The modified configuration preserves the structural symmetry of UNet++ and skip connections enable efficient transfer of multi-scale contextual information between encoder and decoder stages as shown in Fig. 6.

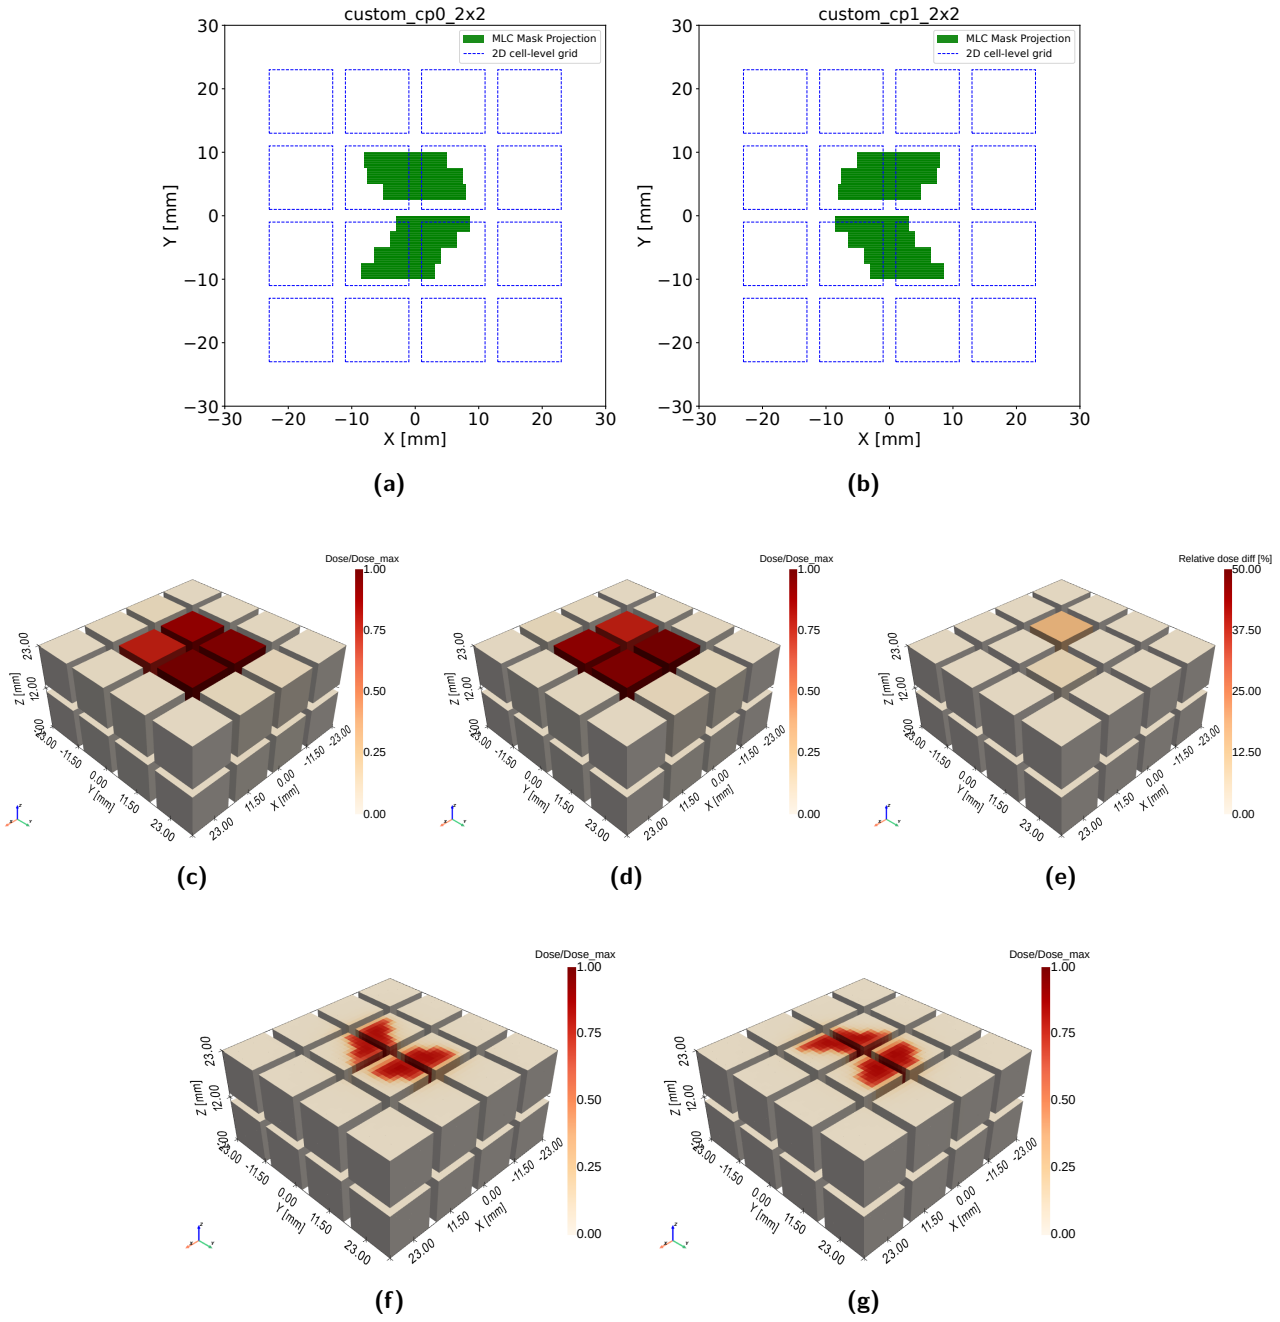



**Figure 5.** Artificial beam masks designed to produce consistent dose distributions at the cell level (a, b). Corresponding dose distributions at the cell level, showing high consistency between the two masks (c, d). Spatial map of the relative dose difference (c & d), expressed with respect to the cell-wise mean of the compared dose values (e). Corresponding dose distributions at the voxel level, illustrating clear differences resulting from the beam configurations (f, g).

During preprocessing, two three-dimensional input channels are constructed. Both volumes are then treated as separate input channels. The first channel consists of the voxel-wise input dose distribution rescaled by the corresponding target values of the three-dimensional GAM. Since both dose and GAM are represented as 3D volumes with the same dimensions, the scaling is performed element-wise, preserving full spatial correspondence. The second channel is derived from a binary 2D mask describing the field shape. The mask is replicated along the Z axis of the patient to form a 3D volume and subsequently weighted as a function of depth using GAM. Specifically, GAM values are summed for each axial slice, normalized to the [0,1] range, and used as slice-wise scaling factors for the mask. As a result, the field geometry remains unchanged, while its intensity varies along the depth (Z-axis) according to the GAM profile (see example Z-axis profile of the slice-wise aggregated GAM in Fig. 7).

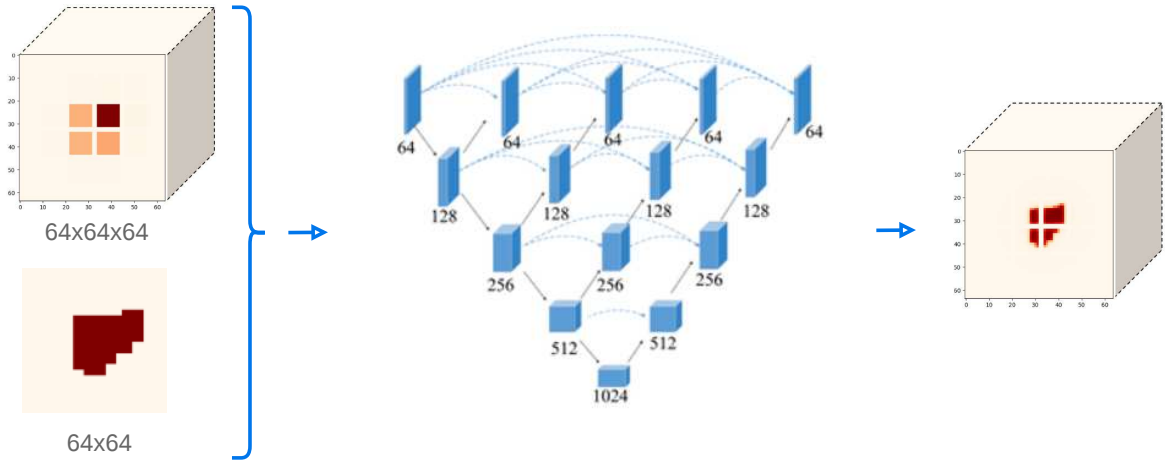



**Figure 6.** Schematic representation of the UNet++ dataflow used for SR. The low-resolution input data are voxelized and resampled to match the spatial grid of the high-resolution target output. The network preserves the characteristic symmetric encoder–decoder structure of UNet++, including dense skip connections that transfer multi-scale contextual features between corresponding stages.

### 3.2 Model Inference

Model inference in our pipeline relies on high-resolution voxelisation for accurate dose prediction. To enable this, the cell-level data, which is initially defined on a coarser grid (e.g., 4x4x2 grid-cells), is mapped to a finer voxel-level grid. Each voxel is assigned a computed GAM during this mapping process. This rescaling ensures that each voxel reflects the appropriate dose contribution based on the field size and shape. This high-resolution mapping and the assignment of GAM descriptors enable model inference, allowing the ML model to generate precise, localized dose predictions across the entire voxel grid.

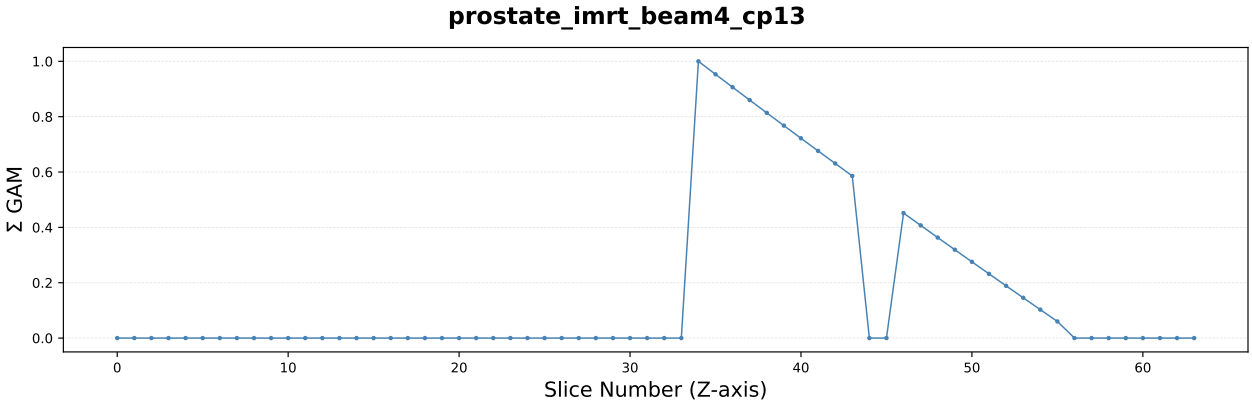



**Figure 7.** Example of the slice-wise aggregated GAM profile as a function of depth along the Z-axis. Nonzero Σ GAM values clearly identify the active detector volume, revealing the presence of two detector cells along the Z-axis.

# 4 Results

## 4.1 Performance on Test Subset

The best results 94.52 % GPR and 48.11 dB PSNR for the test data were obtained at 1740th epoch with a learning rate 0.0001 and Adam optimizer, deep-supervision weights: $[0.1, 0.2, 0.5, 1.0]$ and *CombinedLossFunction* with 50/50 input of the mentioned elements.

## 4.2 Accuracy of SR Dose Measurement

During the inference process, the real low-resolution data from measurement (afterward proper preprocessing — as described in section 3.1) can be passed to the model. Using the test dataset, the model performance is evaluated on CP that mimic real TPS data as shown in Fig. 8. All results are presented for a representative central slice of the volume. In addition, the projection of the MLC mask is overlaid on the cell-level dose distribution (leftmost panel) to facilitate visual comparison between the irradiated region and the detector response.

An additional evaluation was conducted on a dedicated test dataset to assess the influence of the GAM factor on cell-to-voxel shape reconstruction by comparing the performance of two models. The baseline scenario consisted of training the network using only the dose distribution, without the GAM distribution, serving as a reference for the more advanced GAM-integrated configuration. The obtained results are presented in Fig. 9 and Fig. 10, respectively. The results highlight the impact of the GAM on the model's ability to accurately reproduce high-resolution dose distributions and geometric shapes.

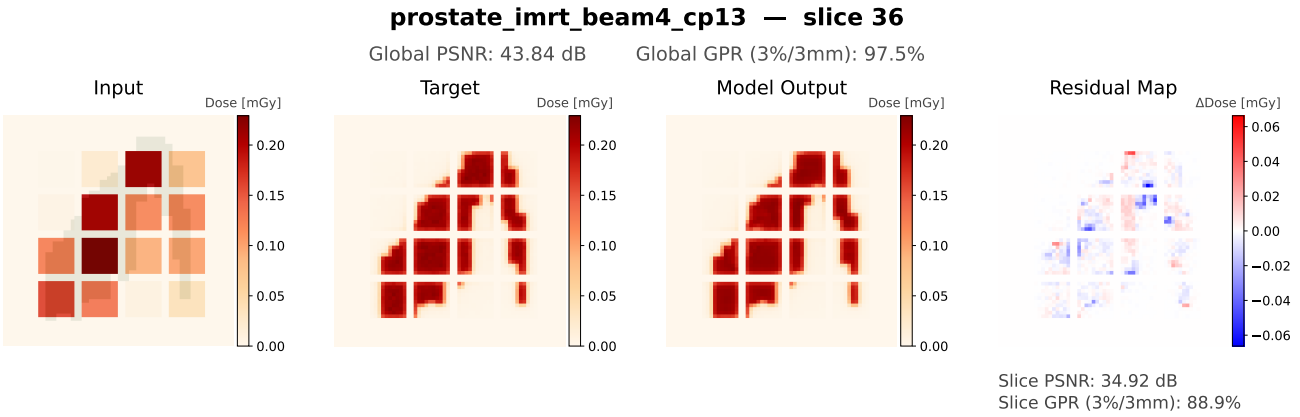



**Figure 8.** Results obtained for clinically representative irradiation fields indicate strong reconstruction performance, as reflected by consistently high PSNR and GPR values. The projected MLC mask is overlaid on the cell-level dose distribution shown in the leftmost panel. In the displayed slice, the predicted and target dose distributions are visually nearly indistinguishable. The corresponding difference map shows that the largest reconstruction errors are concentrated near the distribution boundaries, suggesting that edge regions constitute the most challenging areas for the model. Nevertheless, these discrepancies remain localized and do not significantly affect the overall agreement between the reconstructed and reference dose distributions.

# 5 Discussion

Traditional measurement techniques, while reliable, are constrained by hardware limitations that restrict the achievable spatial resolution. In contrast, DL-driven methods, particularly CNN-based approaches, provide an alternative by learning complex spatial patterns and enabling enhanced reconstruction of dose distributions from lower-resolution measurements. As demonstrated in this work, augmenting the CNN with geometric information extracted from RT-Plan further improves its ability to capture relevant spatial relationships, which is essential for accurate dose reconstruction. The CNN model demonstrates the capability to reconstruct high-resolution dose maps from coarse inputs, suggesting potential utility in improving radiotherapy workflows. After appropriate adaptation and validation, the SR model may serve as a computationally efficient alternative to MC simulations, which are highly accurate but remain computationally expensive and time-consuming. Although training the DL model requires significant computational resources, this cost is limited to the offline training phase. Once trained, the model enables inference that is orders of magnitude faster than MC-based approaches, potentially supporting faster

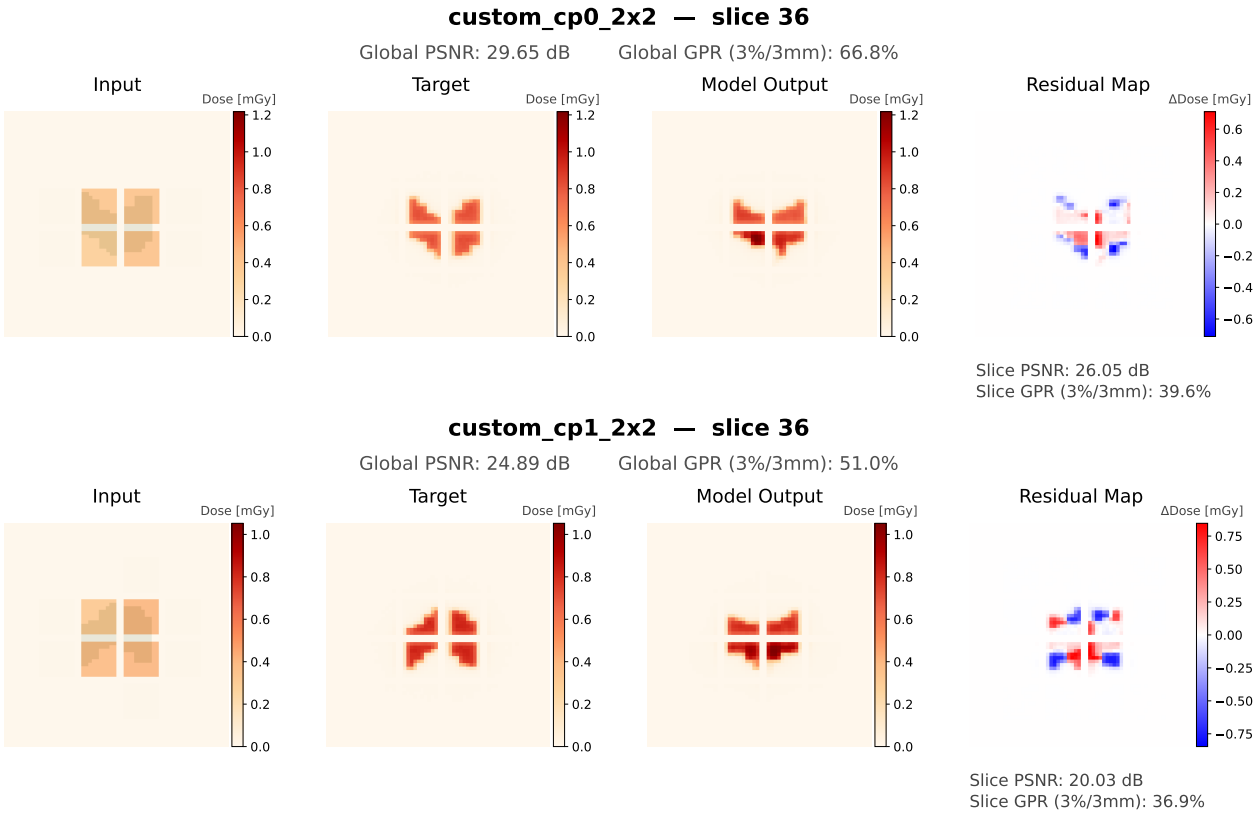



**Figure 9.** Results obtained using the baseline model trained without GAM-based weighting of the field mask for the special test case. The projected MLC masks are overlaid on the cell-level dose distribution shown in the leftmost panels. The low-resolution inputs correspond to similar dose distributions, while the high-resolution targets are mirrored versions of each other. The model struggles to reconstruct the high-resolution dose distributions.

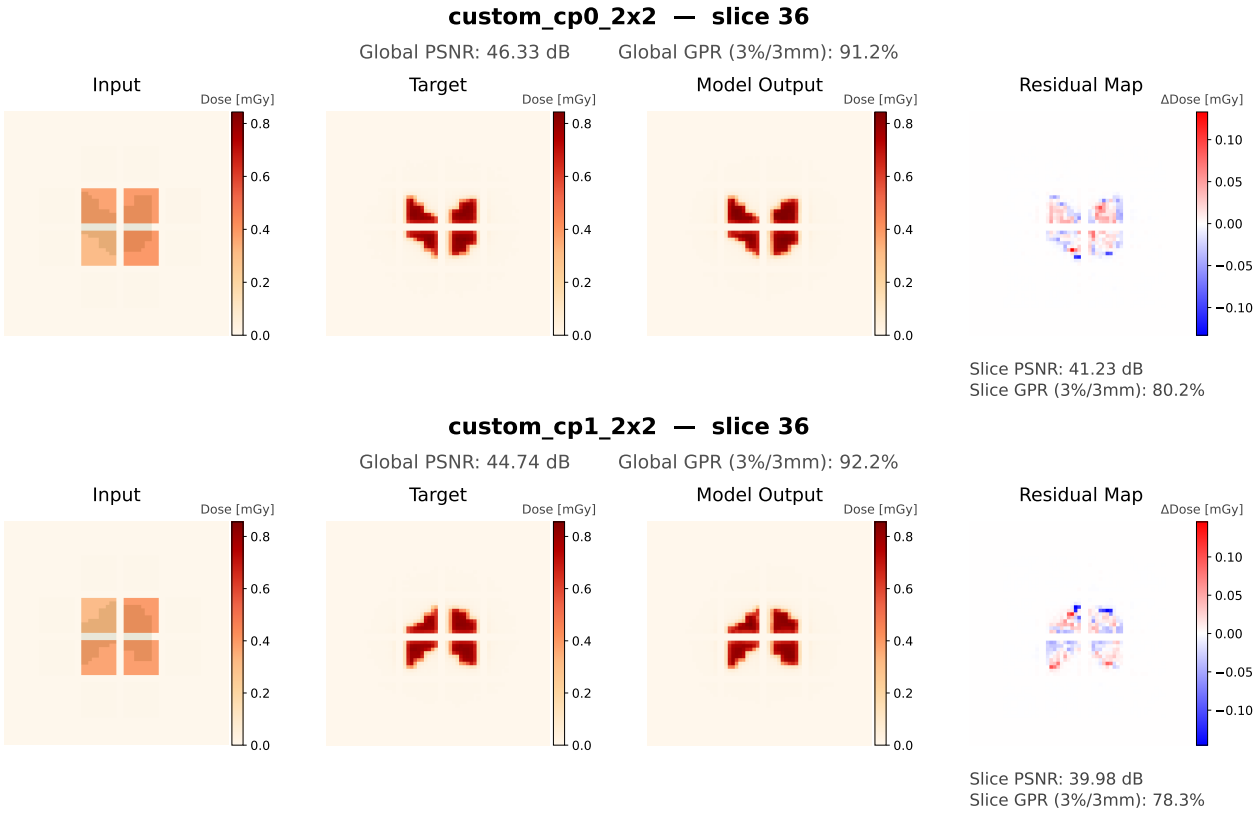



**Figure 10.** Results from a dedicated test designed to evaluate the influence of the GAM descriptor in scenarios where voxel-level variations do not alter the aggregated dose at the cellular scale. The projected MLC masks are overlaid on the cell-level dose distribution shown in the leftmost panels. The model accurately reconstructs the high-resolution targets, preserving both the spatial dose distribution and the geometric characteristics of the irradiated region.

treatment planning and more adaptive radiotherapy strategies. However, these advantages should be interpreted in the context of further validation requirements and clinical constraints.

The current results achieve an average GPR of 94.52 % at 3%/3 mm, whereas the clinically accepted threshold is 90 %[20]. It is important to note that this metric is strongly dependent on the specific CP configuration — for some CP, the performance is satisfactory (over 90 %), while for others it decreases. Moreover, the mean GPR observed across all CP in the training dataset does not differ substantially from that of the test dataset, which likely indicates a limited diversity of dose distribution patterns within the training data. While MC simulations represent the main bottleneck in data preparation due to their high CPU computation time requirements, future work will focus on generating a more extensive and diverse dataset to improve model generalizability.

## References


1. Kopec, M. *et al.* A reconfigurable detector for measuring the spatial distribution of radiation dose for applications in the preparation of individual patient treatment plans. *Nucl. Instruments Methods Phys. Res. Sect. A: Accel. Spectrometers, Detect. Assoc. Equip.* **1048**, 167937, DOI: 10.1016/j.nima.2022.167937 (2023).
2. Jurgielewicz, P. *et al.* A modular data acquisition system for reconstruction of radiation dose spatial distribution in radiotherapy treatment planning. *Nucl. Instruments Methods Phys. Res. Sect. A: Accel. Spectrometers, Detect. Assoc. Equip.* **1045**, 167607, DOI: 10.1016/j.nima.2022.167607 (2023).
3. Mindur, B. *et al.* System for radiation dose distribution monitoring in radiotherapy treatment planning. *Nucl. Instruments Methods Phys. Res. Sect. A: Accel. Spectrometers, Detect. Assoc. Equip.* **1069**, 169834, DOI: 10.1016/j.nima.2024.169834 (2024).
4. Mindur, B. *et al.* Data acquisition system for radiation dose distribution monitoring in photon radiotherapy. *J. Instrumentation* **20**, P06040, DOI: 10.1088/1748-0221/20/06/P06040 (2025).
5. Khan, F. M., Gibbons, J. P. & Sperduto, P. M. *Khan's Treatment Planning in Radiation Oncology* (Lippincott Williams and Wilkins, 2020), 4 edn.
6. De Martino, F., Clemente, S., Graeff, C., Palma, G. & Cella, L. Dose calculation algorithms for external radiation therapy: An overview for practitioners. *Appl. Sci.* **11**, 6806, DOI: 10.3390/app11156806 (2021).
7. Podgorsak, E. B. Review of radiation oncology physics: a handbook for teachers and students. *Vienna, Austria: IAEA* **19**, 133 (2003).
8. Farea, A., Yli-Harja, O. & Emmert-Streib, F. Understanding physics-informed neural networks: Techniques, applications, trends, and challenges. *AI* **5**, 1534–1557, DOI: 10.3390/ai5030074 (2024).
9. Wang, S., Sankaran, S., Wang, H. & Perdikaris, P. An expert's guide to training physics-informed neural networks. *arXiv preprint* (2023). 2308.08468.
10. Agostinelli, S. *et al.* Geant4—a simulation toolkit. *Nucl. Instruments Methods Phys. Res. A* DOI: 10.1016/S0168-9002(03)01368-8 (2003).
11. Hajduga, J. *et al.* Monte carlo simulation platform and software stack in dose-3d project. *Acta Phys. Pol. B Proc. Suppl. 15, 4-A8* DOI: 10.5506/APhysPolBSupp.15.4-A8 (2022).
12. Zhou, Z., Siddiquee, M. M. R., Tajbakhsh, N. & Liang, J. Unet++: A nested u-net architecture for medical image segmentation. *CoRR* **abs/1807.10165** (2018). 1807.10165.
13. Ronneberger, O., Fischer, P. & Brox, T. U-net: Convolutional networks for biomedical image segmentation. *CoRR* **abs/1505.04597** (2015). 1505.04597.
14. Cardoso, M. J. *et al.* Monai: An open-source framework for deep learning in healthcare (2022). 2211.02701.
15. Wang, Z., Bovik, A. C., Sheikh, H. R. & Simoncelli, E. P. Image quality assessment: from error visibility to structural similarity. *IEEE Transactions on Image Process.* **13**, 600–612, DOI: 10.1109/TIP.2003.819861 (2004).
16. Low, D. A., Harms, W. B., Mutic, S. & Purdy, J. A. A technique for the quantitative evaluation of dose distributions. *Med. physics* **25**, 656–661 (1998).
17. Diamantopoulos, S. *et al.* Treatment plan verification: A review on the comparison of dose distributions. *Phys. Medica* **67**, 107–115, DOI: 10.1016/j.ejmp.2019.10.029 (2019).
18. Deng, J., Liu, S., Huang, Y., Li, X. & Wu, X. Evaluating aapm-tg-218 recommendations: Gamma index tolerance and action limits in imrt and vmat quality assurance using suncheck. *J. Appl. Clin. Med. Phys.* **25**, e14277 (2024).

**19.** Atiq, M. *et al.* Interpretation of gamma index for quality assurance of simultaneously integrated boost (sib) imrt plans for head and neck carcinoma. *Pol. J. Med. Phys. Eng.* **23**, 93 (2017).

**20.** Das, S. *et al.* Gamma index analysis as a patient-specific quality assurance tool for high-precision radiotherapy: A clinical perspective of single institute experience. *Cureus* **14**, e30885, DOI: 10.7759/cureus.30885 (2022).

## Funding

This work was supported by the Foundation for Polish Science under the project POIR.04.04.00-00-15E5/18.

## Acknowledgements

We gratefully acknowledge Poland's high-performance Infrastructure PLGrid ACC Cyfronet AGH for providing computer facilities and support within computational grant no PLG/2025/018314.

The authors acknowledge the Maria Skłodowska-Curie National Research Institute of Oncology, Kraków Branch, for providing access to the medical data used in this study under the framework agreement on cooperation. The data were anonymised before being made available for research purposes.

## Author contributions statement

K.K. and B.R. contributed to conceptualization, formal analysis, and writing of the original draft. J.H. contributed to software development and visualization. J.K. and D.Ko. contributed to formal analysis. D.Ka. contributed to methodology. J.D. contributed to review and editing of the manuscript. T.F., B.Ł., B.M., J.M., and P.W. contributed resources. T.S. contributed to conceptualization and review and editing of the manuscript. All authors reviewed and approved the final manuscript.

## Additional information

**Competing interests:** The authors declare no competing interests.

## Data availability

The datasets generated and analysed during the current study are not publicly available due to institutional restrictions related to patient data but are available from the corresponding author on reasonable request.