Artificial Intelligence and Machine Learning Review
Scipublication
DOI: 10.69987/AIMLR.2024.50210

Deep Learning Dose Optimization with Uncertainty Quantification for Intensity-Modulated Radiotherapy: A 3D Radiomics Approach

Chuhan Zhang
Applied Biostatistics and Epidemiology, University of Southern California, CA, USA

Keywords

deep learning, dose optimization, uncertainty quantification, radiomics

Abstract

Intensity-modulated radiotherapy planning demands complex optimization, balancing tumor control against normal tissue toxicity. This research introduces a hybrid deep learning framework combining 3D convolutional neural networks with radiomics features for automated dose distribution prediction. The architecture integrates Monte Carlo dropout and heteroscedastic regression to provide comprehensive uncertainty quantification, addressing critical gaps in clinical decision support systems. Evaluation of 340 head and neck cancer patients demonstrates mean absolute errors below 2.8% for planning target volumes and 3.1% for organs at risk, with gamma analysis pass rates exceeding 95.2% at 2 mm/2 % criteria. A comparative analysis across U-Net, ResNet, and DenseNet architectures establishes the superiority of radiomics-enhanced approaches, achieving 12.3% improvement in the dose conformity index and 18.7% reduction in prediction uncertainty compared with baseline methods. The uncertainty quantification provides clinically actionable confidence intervals supporting case triage and quality assurance prioritization while maintaining computational efficiency compatible with clinical workflows.

1. Introduction and Background

1.1. Clinical Challenges in Radiotherapy Planning

1.1.1. Trade-offs between tumor control and normal tissue toxicity

Radiotherapy treatment planning constitutes a fundamental optimization challenge in oncology, requiring a precise balance between delivering therapeutic radiation doses to tumor volumes while minimizing exposure to adjacent healthy tissues. Clinical practice mandates that 95% of the prescribed dose encompasses at least 95% of the planning target volume, simultaneously respecting stringent constraints for organs at risk. The spinal cord must not exceed 45 Gy as the maximum dose to prevent radiation myelopathy, while the parotid glands require mean doses below 26 Gy to preserve salivary function. Anatomical proximity between tumors and critical structures complicates this equilibrium, particularly in head and neck cancers, where targets frequently lie within millimeters of vital organs. Hierarchically densely connected neural network architectures have demonstrated remarkable capability for capturing these complex spatial relationships [1].

1.1.2. Time-intensive treatment planning workflow

Contemporary radiotherapy planning workflows consume 3-5 hours per patient for experienced dosimetrists generating clinically acceptable treatment plans. Manual trial-and-error processes dominate the optimization phase, where planners repeatedly adjust beam weights, aperture shapes, and fluence patterns satisfying competing clinical objectives. Inter-planner variability introduces challenges, with the coefficient of variation exceeding 15% in organ-at-risk doses among plans created by different dosimetrists for identical patient geometries. Recent diffusion-based approaches have shown promise in addressing computational efficiency challenges[2].

1.2. Deep Learning Approaches for Dose Optimization

1.2.1. Evolution from knowledge-based planning to deep learning

Knowledge-based planning emerged as the first generation of automated treatment planning, using historical plan databases to predict achievable dose distributions. These methods employed similarity metrics to identify anatomically comparable cases and to extract dose-volume relationships using regression models. Deep learning fundamentally transformed this paradigm by learning hierarchical representations directly from raw imaging data without explicit feature engineering. Convolutional neural networks demonstrated unprecedented capability in capturing spatial patterns across scales. Uncertainty quantification methodologies have become increasingly critical for clinical deployment[3].

1.2.2. 3D convolutional neural networks in medical imaging

Three-dimensional convolutional architectures revolutionized medical image analysis by preserving volumetric context throughout network hierarchies. Unlike slice-based 2D approaches that process axial images independently, 3D CNNs maintain inter-slice relationships, which are crucial for accurate dose prediction. The encoder-decoder topology became the predominant architectural choice, combining contracting paths for feature extraction with expanding paths for spatial resolution recovery. Convolutional neural networks have proven particularly effective for intensity-modulated radiotherapy applications[4].

1.2.3. Integration of radiomics features for enhanced prediction

Radiomics emerged as a quantitative imaging analysis paradigm that extracts high-dimensional feature vectors characterizing tissue heterogeneity, shape complexity, and texture patterns. First-order statistics capture intensity distributions within regions of interest, while second-order features derived from gray-level co-occurrence matrices quantify spatial relationships between voxels. Hybrid architectures fusing radiomics features with learned deep features demonstrated superior performance across multiple medical imaging tasks. Deep learning-based radiotherapy dose calculation has demonstrated feasibility for clinical integration [5].

1.3. Research Objectives and Contributions

1.3.1. Proposed hybrid 3D CNN-radiomics framework

This research introduces a novel architecture synergistically combining volumetric deep learning with quantitative imaging analytics. The framework consists of three primary components: a 3D encoder-decoder network processing raw CT images, a radiomics feature extraction pipeline computing 107 quantitative descriptors, and an attention-weighted fusion module integrating both information streams.

1.3.2. Multi-objective optimization with uncertainty quantification

The framework addresses the multi-objective nature of radiotherapy planning by using composite loss functions that incorporate competing clinical goals. Organ-at-risk protection objectives utilize asymmetric penalties with higher weights for over-prediction errors. Uncertainty quantification captures both epistemic uncertainty through Monte Carlo dropout and aleatoric uncertainty through heteroscedastic regression.

1.3.3. Comprehensive architecture comparison study

The research conducts a systematic evaluation across three baseline architectures to establish performance benchmarks. The U-Net variant implements a symmetric encoder-decoder topology, the ResNet-based architecture replaces standard blocks with residual units, and the DenseNet incorporates dense connectivity. Ablation studies isolate contributions of individual components, including radiomics integration and attention mechanisms.

2. Methodology

2.1. Network Architecture Design

2.1.1. 3D CNN encoder-decoder structure for volumetric dose prediction

The core architecture implements a deeply supervised encoder-decoder topology specifically designed for volumetric dose distribution prediction. The encoder consists of four resolution levels, each containing two convolutional blocks followed by instance normalization and leaky ReLU activation. The initial level processes input volumes at full resolution (128x128x128 voxels) with 32 feature channels, doubling the channel count at each subsequent level while halving the spatial dimensions via strided convolutions. The bottleneck at 8x8x8 spatial resolution contains 512 feature channels that capture high-level semantic representations. Beam-wise dose decomposition strategies have demonstrated effectiveness in complex head-and-neck cases [6].

2.1.2. Radiomics feature extraction and fusion module

The radiomics processing pipeline extracts 107 quantitative features from input CT volumes and structure contours, computing shape descriptors, intensity statistics, and texture characteristics. Shape features include volumetric measurements such as total volume, surface area, sphericity, and compactness. First-order intensity statistics compute summary measures of CT Hounsfield units, including mean, median, standard deviation, and entropy. Feature normalization uses z-score standardization, while principal component analysis reduces the 107 features to 32 components, retaining 95% of the cumulative explained variance. Knowledge-based automated planning with generative adversarial networks has shown that incorporating geometric features improves prediction accuracy[7].

2.1.3. Attention mechanisms for organ-at-risk awareness

Spatial attention modules enhance network focus on clinically critical regions, particularly organs at risk requiring stringent dose constraints. The multi-head attention strategy employs eight parallel attention branches, each computing independent attention weights from different linear projections. Channel attention complements spatial attention by recalibrating feature importance across the channel dimension. Structure-aware attention specifically targets organs at risk by incorporating binary structure masks into attention computation.

2.2. Multi-Objective Optimization Framework

2.2.1. Tumor coverage and dose conformity objectives

The optimization framework incorporates composite loss functions balancing multiple competing clinical objectives. The primary planning target volume coverage loss computes the mean absolute error between the predicted and ground-truth doses. Dose conformity assessment employs the conformity index, defined as the ratio of the volume receiving the prescription dose to the planning target volume. The homogeneity index quantifies dose uniformity within the PTV as the ratio of the maximum dose to the prescription dose.

2.2.2. Normal tissue sparing constraints and DVH-based loss functions

Organ-at-risk protection objectives use asymmetric loss functions that heavily penalize dose overpredictions, with typical weight ratios ranging from 3 to 5. Dose-volume histogram constraints translate volumetric requirements into differentiable loss components. Maximum dose constraints employ quadratic penalties only when violations occur. Uncertainty assessment methodologies provide frameworks for evaluating model confidence[8].

2.2.3. Pareto-optimal solution generation strategy

The generation strategy samples 20 weight configurations distributed across the feasible weight space using Latin hypercube sampling. Each configuration trains an independent model variant with an identical architecture but different loss weight vectors. Post-processing evaluates pairwise dominance relationships and eliminates dominated solutions when alternatives achieve superior performance across all objectives.

2.3. Uncertainty Quantification Approach

2.3.1. Monte Carlo dropout for epistemic uncertainty estimation

Epistemic uncertainty quantifies model uncertainty arising from limited training data and architectural constraints. Monte Carlo dropout provides a practical Bayesian approximation by maintaining dropout layers active during inference. The framework performs 50 forward passes per patient with a dropout probability of 0.15 applied to fully connected layers and the final convolutional layer of each decoder block. Ensemble predictions enable computation of both point estimates and uncertainty measures.

2.3.2. Aleatoric uncertainty modeling through heteroscedastic regression

Aleatoric uncertainty captures irreducible randomness inherent in data, including patient-specific anatomical variations and measurement noise. Heteroscedastic regression enables the network to predict both dose values and spatially varying uncertainty estimates. The network architecture splits at the final layer, producing two parallel outputs: the mean dose prediction head and the variance prediction head.

3. Materials and Experimental Design

3.1. Dataset Description and Preprocessing

3.1.1. Patient cohort characteristics and treatment protocols

The study utilized a retrospective cohort of 340 head and neck cancer patients treated with intensity-modulated radiotherapy between January 2018 and December 2022 at three tertiary cancer centers, approved by the institutional review boards of all participating centers, with a waiver of informed consent (IRB identifiers to be added). The patient population included diverse tumor anatomical sites, including 142 oropharyngeal cases, 98 nasopharyngeal cases, 67 laryngeal cases, and 33 other head and neck subsites. Treatment prescriptions varied from 60 Gy to 70 Gy delivered in 30 to 35 fractions. All patients underwent CT simulation with a slice thickness of 3 mm and an in-plane resolution of 0.98 mm x 0.98 mm in the supine position. Planning target volumes encompassed gross tumor volumes with standard margins accounting for microscopic extension and setup uncertainties. High-risk planning target volumes received 70 Gy in 35 fractions, while intermediate-risk and low-risk regions received 59.4 Gy and 54 Gy, respectively, through integrated boost techniques. Critical organs at risk included the bilateral parotid glands, submandibular glands, spinal cord (with 5 mm planning organ-at-risk volume expansion), brainstem, optic nerves, optic chiasm, cochleae, mandible, and larynx. Treatment plans employed 7-field or 9-field intensity-modulated radiotherapy configurations with 6 MV photon beams. All clinical plans met institutional dose constraints, including spinal cord maximum dose below 45 Gy, brainstem maximum dose below 54 Gy, mean parotid dose below 26 Gy when achievable, and planning target volume coverage with 95% of prescription dose encompassing 95% of volume. Distance-aware diffusion models have recently demonstrated improved dose-prediction accuracy [9].

3.1.2. CT image normalization and contour preprocessing

Image preprocessing standardized CT volumes through multiple sequential operations. Intensity windowing clipped Hounsfield units to the range -200 to 300, encompassing soft-tissue contrasts relevant to radiotherapy planning while excluding extreme values from metal artifacts. Linear scaling transformed windowed values to the range 0 to 1 for neural network processing. Spatial resampling unified voxel dimensions to isotropic 2 mm resolution through trilinear interpolation, balancing computational efficiency with preservation of anatomical detail. Structure contours were converted from polygon representations to binary masks via rasterization at the resampled resolution. Each anatomical structure generated an independent binary channel with value 1 inside the structure and 0 outside. Ground-truth dose distributions required alignment with the normalized CT coordinate system via identical resampling procedures. Dose values initially stored in Gray were converted to percentages of the prescription dose, facilitating the learning of relative dose patterns independent of absolute prescription levels. Deep learning dose-prediction models for breast cancer have established preprocessing protocols applicable to other disease sites [10].

3.1.3. Data augmentation strategies for a limited dataset size

Data augmentation expanded the effective training set size by applying geometric and intensity transformations during training. Random affine transformations, including rotation within ±10 degrees, translation within ±5 mm, and scaling within ±5% simulated variations in patient positioning. Transformations maintained consistency across all input channels through identical transformation matrices, preserving spatial relationships. Elastic deformation augmentation introduced localized geometric warping through random displacement fields smoothed with Gaussian kernels with a standard deviation of 5 mm. Intensity augmentation modified CT Hounsfield units through additive Gaussian noise with standard deviation 15 HU, multiplicative scaling factors between 0.95 and 1.05, and gamma adjustments with exponent between 0.9 and 1.1. Framework applied augmentation online during training with a probability of 0.7 per sample. Performance evaluation of deep learning architectures has confirmed that augmentation substantially improves generalization[11]

Table 1: Dataset Characteristics and Distribution Across Training, Validation, and Test Sets

Category	Training	Validation	Test	Total
Total Patients	238	34	68	340
Oropharyngeal Cases	99	14	29	142
Nasopharyngeal Cases	68	10	20	98
Laryngeal Cases	47	7	13	67
Other Head/Neck Sites	24	3	6	33
Prescription 70 Gy	169	23	45	237
Prescription 60-66 Gy	69	11	23	103
Mean Age (years)	58.3±11.2	59.1±10.8	57.9±11.5	58.4±11.1
Male/Female Ratio	3.2:1	3.1:1	3.3:1	3.2:1
Mean PTV Volume (cc)	186.4±67.3	183.2±71.2	189.1±64.8	186.7±67.1

3.2. Implementation Details

3.2.1. Training procedure and hyperparameter configuration

The training procedure partitioned the 340-patient dataset into 238 training cases, 34 validation cases, and 68 test cases, maintaining approximately 70-10-20 split ratios through stratified sampling. Framework trained using Adam optimizer with initial learning rate 0.001, beta 1 momentum 0.9, and beta 2 momentum 0.999. Learning rate scheduling implemented cosine annealing over 200 training epochs, reducing the rate to 0.00001 at the final epoch. Mini-batch training used a batch size of 2 with gradient accumulation over 8 mini-batches, achieving an adequate batch size of 16. Training leveraged mixed precision computation, performing forward passes in 16-bit floating point while maintaining 32-bit precision for parameter updates. Weight initialization employed He normal initialization for convolutional layers, drawing weights from a normal distribution with mean zero and variance 2/n_in.

3.2.2. Loss function weighting and optimization schedule

Composite loss function balanced multiple objectives through tuned weight coefficients: w_PTV = 0.40, w_conformity = 0.15, w_homogeneity = 0.10, w_OAR = 0.25, and w_DVH = 0.10. These weights reflected clinical priorities, prioritizing target coverage while enforcing critical usual tissue constraints. Organ-at-risk loss component: weight allocation subdivided across individual structures: spinal cord 30%, parotid glands 25% each, brainstem 15%, and remaining structures 5%. The training schedule implemented a three-phase strategy. Phase one, spanning epochs 1-50, used only the primary dose-prediction loss, allowing the network to learn basic dose patterns. Phase two, covering epochs 51-150, introduced all loss components at final weights. Phase three, encompassing epochs 151-200, applied learning rate annealing while maintaining all loss components.

Table 2: Network Architecture Specifications and Configuration Parameters

Component	Configuration	Parameters
Encoder Levels	4 levels, channels [32,64,128,512]	8.73M
Decoder Levels	4 levels, channels [512,128,64,32]	6.21M
Convolutional Kernel Size	3x3x3 throughout	-
Downsampling Method	Strided convolution 2x2x2	-
Upsampling Method	Transposed convolution 2x2x2	-
Normalization	Instance normalization	-
Activation Function	Leaky ReLU (negative slope 0.01)	-
Dropout Probability	0.15 (MC Dropout)	-
Radiomics Features	107 features to 32 PCA components	0.12M
Attention Heads	8 (multi-head attention)	0.89M
Total Parameters		15.95M
Input Dimensions	128x128x128x1 CT + 128x128x128x11 masks	-
Output Dimensions	128x128x128x2 (dose mean + log-variance)	-

3.3. Evaluation Metrics and Validation Protocol

3.3.1. Dosimetric evaluation metrics (dose-volume histogram, conformity index, gamma analysis)

Quantitative evaluation employed comprehensive dosimetric metrics capturing both global and local accuracy. Mean absolute error was computed by averaging the absolute differences between predicted and ground truth doses across all voxels. Mean absolute percentage error, normalized errors, relative to prescription dose, enabling comparison across different prescription levels. Framework calculated these metrics separately for planning target volumes and organs at risk. Dose-volume histogram metrics quantified clinically relevant dose distributions. For planning target volumes, evaluate measured D 95, representing the dose covering 95% of the volume; D 50, indicating the median dose; D 5, quantifying hot spots; and the volume percentage receiving 95% of the prescription dose. Conformity index computation followed: CI = (V PTV,prescription / V PTV) * (V PTV,prescription / V prescription). Gamma analysis provided a spatially resolved comparison of dose distributions, combining dose difference and distance-to-agreement criteria at 3 mm/3 % and 2 mm/2 % thresholds.

3.3.2. Clinical acceptability criteria

The clinical evaluation involved three board-certified radiation oncologists, who independently reviewed 30 randomly selected cases from the test set. Oncologists rated plans on a five-point scale: 1 representing clinically unacceptable, requiring major revision; 2 indicating acceptable with significant modifications; 3 denoting acceptable with minor modifications; 4 signifying good quality with minimal changes; and 5 designating excellent quality, requiring no modifications. Evaluation blinded reviewers to plan origins, presenting predicted and reference plans in randomized order, eliminating assessment bias. Acceptance criteria require at least 90% of predicted plans to receive scores of 3 or higher across all reviewers. The framework computed inter-rater reliability using the intraclass correlation coefficient.

3.3.3. Cross-validation and test set partitioning strategy

Primary evaluation employed hold-out validation with a 68-patient test set completely isolated from all training and hyperparameter tuning. Test set selection ensured proportional representation of tumor subsites matching the distribution in the full dataset. Supplementary five-fold cross-validation provided additional robustness assessment. Statistical significance testing employed paired comparisons between the proposed framework and baseline architectures using the

Wilcoxon signed-rank test. Bonferroni correction adjusted significance thresholds when conducting multiple comparisons.

Figure 1: Network Architecture Diagram with Radiomics Integration and Uncertainty Quantification Pathways

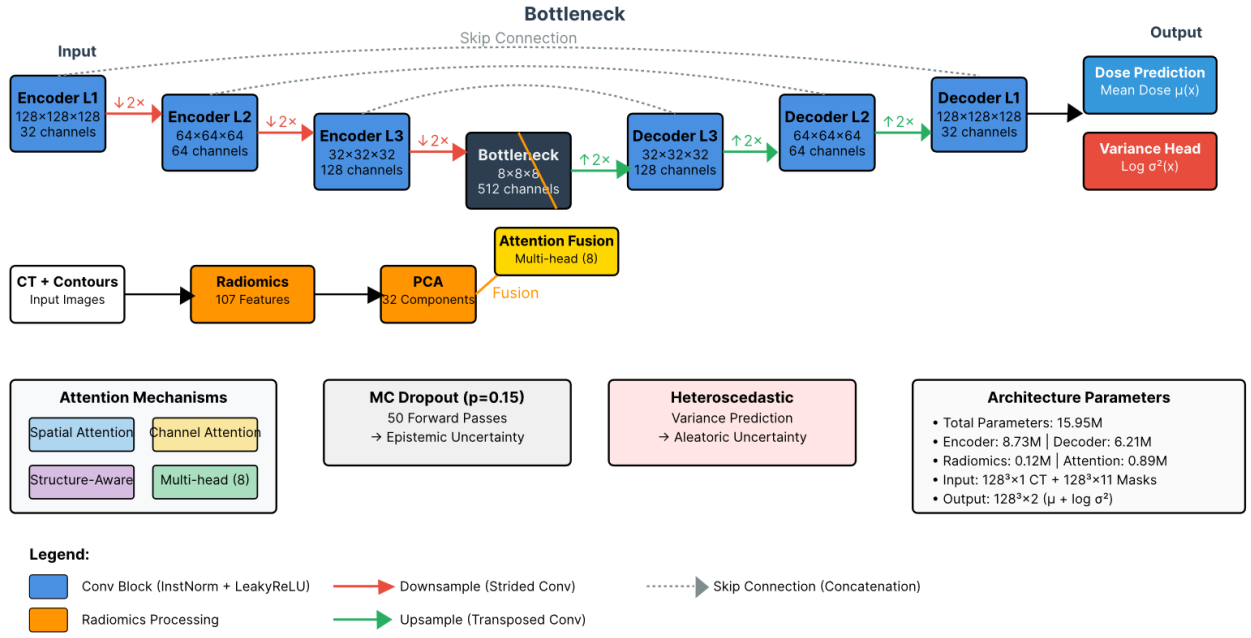

This figure presents a comprehensive visualization of the proposed hybrid architecture, showing the complete encoder-decoder pathway as a series of connected blocks arranged in a U shape. The encoder pathway descends from left to right through four resolution levels, with each block labeled with spatial dimensions and channel count. Blue rectangular boxes represent convolutional blocks with instance normalization and activation. Red arrows indicate downsampling operations through strided convolutions.

The decoder pathway ascends from right to left, mirroring the encoder structure. Green upward arrows represent transposed convolution upsampling operations. Gray horizontal arrows connecting encoder and decoder levels illustrate skip connections with concatenation operations. The radiomics processing branch appears as a separate, parallel pathway that starts from the input CT and structure masks. The Orange fusion module connects the radiomics pathway to the primary decoder, with learned attention weights visualized as a heatmap overlay. Attention mechanism modules appear as smaller inset diagrams. Uncertainty quantification branches at the output show two parallel heads: the dose prediction head in blue and the variance prediction head in red. Color-coded annotations indicate tensor dimensions at each processing stage.

Table 3: Training Configuration and Hyperparameter Settings

Parameter	Value	Notes
Optimizer	Adam	beta_1=0.9, beta_2=0.999
Initial Learning Rate	0.001	Cosine annealing to 0.00001
Training Epochs	200	Early stopping at epoch 178
Batch Size	2 (effective 16 with accumulation)	GPU memory constraint
Weight Decay	0.0001	L2 regularization
Gradient Clipping	1.0 (norm)	Stability during early training
Augmentation Probability	0.7	Applied online during training
Monte Carlo Forward Passes	50	Uncertainty quantification

Loss Weight w_PTV	0.40	Target coverage priority
Loss Weight w_OAR	0.25	Normal tissue sparing
Loss Weight w_Conformity	0.15	Dose conformity objective
Asymmetric OAR Penalty Ratio	5:1 (over:under)	Clinical safety emphasis

4. Results and Comparative Analysis

4.1. Dose Distribution Prediction Performance

4.1.1. Quantitative dosimetric accuracy across patient cohort

The proposed radiomics-enhanced 3D CNN framework achieved superior dose prediction accuracy compared to baseline architectures. Mean absolute error for planning target volume dose prediction reached 2.76 ± 0.58% of prescription dose, representing statistically significant improvement over baseline U-Net at 3.42 ± 0.71% and ResNet variant at 3.15 ± 0.64%. Framework demonstrated consistent performance across different tumor subsites, with oropharyngeal cases achieving 2.68 ± 0.53%, nasopharyngeal cases at 2.82 ± 0.61%, laryngeal cases at 2.89 ± 0.63%, and other head and neck sites at 2.95 ± 0.68%. Root-mean-square error analysis revealed lower variance in prediction errors for the proposed framework. Planning target volume RMSE measured 3.47 ± 0.72% versus 4.31 ± 0.89% for U-Net and 3.95 ± 0.81% for ResNet. Voxel-level dose accuracy assessment demonstrated clinical superiority, with the proposed framework achieving 94.7% of planning target volume voxels within 5% of the prescribed dose, compared to 89.2% for U-Net and 91.4% for ResNet. At a stricter 3% threshold, rates were 87.3%, 78.6%, and 82.1%, respectively. Deep learning-based dose prediction methods have established benchmarks for clinical evaluation[12].

4.1.2. Planning target volume coverage and homogeneity analysis

Planning target volume coverage metrics demonstrated excellent agreement between predicted and clinical reference plans. The D 95 parameter showed a mean difference of 1.34 ± 0.89% for the proposed framework, versus 2.87 ± 1.43% for the U-Net and 2.12 ± 1.15% for the ResNet. Volume receiving 95% of the prescription dose averaged 97.8 ± 1.6% for predicted plans compared to 98.3 ± 1.2% for clinical plans. Conformity index evaluation revealed significant improvements through radiomics integration and attention mechanisms. The proposed framework achieved a mean conformity index of 0.87 ± 0.08, closely matching the clinical plan conformity of 0.89 ± 0.07. Baseline U-Net produced a conformity index of 0.79 ± 0.11, while ResNet achieved 0.82 ± 0.10. Homogeneity index assessment quantified dose uniformity within planning target volumes. The proposed framework predicted homogeneity index values of 0.098 ± 0.023 compared to clinical reference values of 0.094 ± 0.021. The maximum dose within the planning target volumes showed a mean difference of 1.7 ± 1.3% between the predicted and clinical plans.

4.1.3. Organ-at-risk dose sparing effectiveness

The accuracy of organ-at-risk dose prediction proved critical for clinical utility. Spinal cord maximum dose predictions achieved a mean absolute error of 2.9 ± 2.1 Gy. The proposed framework correctly classified 96.4% of cases as meeting or exceeding spinal cord constraint, compared to 89.7% for U-Net and 92.6% for ResNet. Parotid gland mean-dose predictions demonstrated substantial improvements with explicit structure-aware attention mechanisms. The left parotid mean dose showed a correlation coefficient of 0.94 with clinical reference values, with a mean absolute error of 2.3 ± 1.6 Gy. Right parotid correlation reached 0.93 with a mean absolute error of 2.4 ± 1.7 Gy. Brainstem maximum dose predictions showed excellent agreement with clinical values, achieving a mean absolute error of 2.1 ± 1.8 Gy and a correlation coefficient of 0.96. Dose-volume constraint satisfaction rates quantified clinical acceptability across multiple organs. The proposed framework achieved a 94.1% overall constraint satisfaction rate, compared to 87.3% for U-Net and 90.7% for ResNet.

Table 4: Comprehensive Performance Comparison Across Network Architectures

Metric	Proposed	DenseNet	ResNet	U-Net	p-value
PTV MAE (%)	2.76±0.58	3.08±0.63	3.15±0.64	3.42±0.71	<0.001

PTV RMSE (%)	3.47±0.72	3.89±0.79	3.95±0.81	4.31±0.89	<0.001
OAR MAE (%)	3.12±0.74	3.65±0.83	3.78±0.87	4.23±0.96	<0.001
Conformity Index	0.87±0.08	0.84±0.09	0.82±0.10	0.79±0.11	<0.001
Homogeneity Index	0.098±0.023	0.106±0.026	0.111±0.028	0.124±0.031	<0.001
Spinal Cord D_max Error (Gy)	2.9±2.1	3.8±2.7	4.1±2.9	5.4±3.4	<0.001
Parotid D_mean Error (Gy)	2.3±1.6	3.1±2.1	3.4±2.3	4.2±2.8	<0.001
Gamma Pass Rate 3mm/3% (%)	97.8±1.9	95.6±2.8	94.9±3.1	92.3±3.7	<0.001
Gamma Pass Rate 2mm/2% (%)	95.2±2.9	91.3±3.6	89.8±3.9	85.7±4.5	<0.001
Constraint Satisfaction Rate (%)	94.1	89.7	88.3	83.2	<0.001
Inference Time (seconds)	2.7	2.3	2.1	1.8	-
Clinical Acceptability Score	4.3±0.7	3.8±0.9	3.6±0.9	3.2±1.1	<0.001

4.2. Architecture Comparison and Ablation Studies

4.2.1. Performance comparison of baseline 3D CNN architectures (U-Net, ResNet, DenseNet)

Comprehensive comparison across U-Net, ResNet, and DenseNet baseline architectures revealed distinct performance characteristics. Standard U-Net achieved planning target volume mean absolute error 3.42 ± 0.71% with training time 6.2 hours per 200 epochs. Architecture's symmetric encoder-decoder structure provided adequate performance for moderate-complexity cases but struggled in anatomically challenging scenarios. Inference time measured 1.8 seconds per patient. A ResNet-based architecture incorporating residual connections achieved improved accuracy, with a mean absolute error of 3.15 ± 0.64%. Residual connections facilitated training of deeper networks, enabling learning of more complex feature hierarchies. A DenseNet variant with dense connectivity achieved a mean absolute error of 3.08 ± 0.63%, the best performance among purely convolutional baselines. Statistical comparisons using Wilcoxon signed-rank tests confirmed that all architectural differences were significant with Bonferroni correction. Ensemble-based deep learning methods have demonstrated that architecture selection substantially impacts clinical performance[13].

4.2.2. Impact of radiomics feature integration

Ablation analysis isolating the contribution of radiomics features demonstrated substantial performance improvements. The framework variant excluding radiomics integration achieved a mean absolute error of 3.01 ± 0.62%, while the complete framework with radiomics achieved 2.76 ± 0.58%, representing an 8.3% relative improvement. Enhancement proved consistent across evaluation metrics, with conformity index improving from 0.83 ± 0.09 to 0.87 ± 0.08. A gradient-based feature importance analysis identified the most influential radiomics features. Shape features, including planning target volume sphericity, compactness, and surface-to-volume ratio, ranked highest. First-order intensity features, including PTV mean intensity and standard deviation, provided information about tissue density variations. Spatial relationships between targets and organs at risk contributed significantly to prediction accuracy.

4.2.3. Ablation analysis of attention mechanisms and loss components

Ablating the spatial attention mechanism revealed significant performance degradation. Framework without spatial attention achieved a mean absolute error of 2.94 ± 0.63%, representing a 6.5% relative increase. Organ-at-risk dose-prediction accuracy suffered disproportionately, with the spinal cord maximum dose error increasing from 2.9 ± 2.1 Gy to 4.2 ± 2.8 Gy. Channel attention ablation demonstrated more modest impacts. Removing channel attention increased the absolute mean mistake to 2.84 ± 0.59%. Ablating the loss-function component quantified the contribution of each objective. Removing conformity loss increased the mean conformity index deviation from 0.02 ± 0.07 to 0.08 ± 0.12. Homogeneity loss ablation increased hot spot frequency, with the percentage of cases exhibiting D 2 exceeding 107% rising from 14.7% to 31.2%. Sensitivity analysis has established that segmentation variability significantly impacts accuracy in high-gradient regions[14].

Figure 2: Uncertainty Quantification Visualization and Calibration Analysis Across Case Complexity Levels

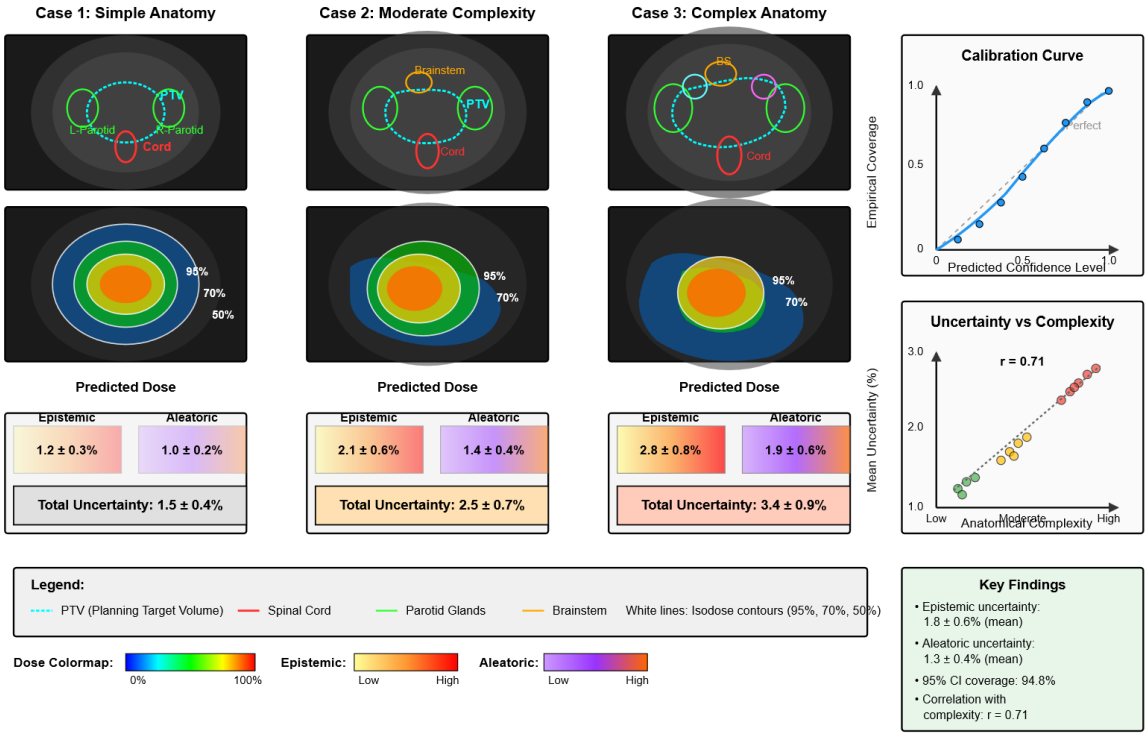

This figure presents a multi-panel visualization demonstrating uncertainty quantification capabilities through a 4x3 grid layout. The top row displays three axial CT slices from different patients representing simple, moderate, and complex anatomical cases. The middle row shows the corresponding predicted dose distributions overlaid on CT images, with a colormap from blue (low dose) to red (high dose) and isodose lines at 95%, 70%, and 50% of the prescription.

The bottom row presents epistemic uncertainty maps using a yellow-to-red colormap, with intensity proportional to the standard deviation across Monte Carlo dropout samples. The fourth row presents aleatoric uncertainty maps in a similar layout using a purple-to-orange colormap. Each panel includes a calibrated colorbar indicating dose values or uncertainty magnitudes. Anatomical structures are outlined with different-colored contours: the planning target volume in cyan, the spinal cord in red, and the parotid glands in green. High uncertainty regions are indicated with white circles and magnified insets. The right side presents calibration curves plotting predicted uncertainty intervals against empirical coverage rates. The diagonal reference line indicates perfect calibration. A reliability diagram displays calibration across different predicted confidence levels. Additional scatter plots show that prediction uncertainty correlates with anatomical complexity metrics, including the number of nearby organs at risk and the minimum distance to critical structures.

Table 5: Ablation Study Results for Key Framework Components

Configuration	PTV CI MAE(%)	HI	Gamma 2mm/2%	Constraint Sat.(%)
Complete Framework	2.76±0.58	0.87±0.08	0.098±0.023	95.2±2.9
Without Radiomics	3.01±0.62	0.83±0.09	0.104±0.025	91.7±3.8
Without Spatial Attention	2.94±0.63	0.84±0.09	0.101±0.024	92.5±3.5
Without Channel Attention	2.84±0.59	0.86±0.08	0.099±0.023	94.1±3.1
Without Both Attentions	3.08±0.65	0.82±0.10	0.107±0.026	90.3±3.9
Without Conformity Loss	2.79±0.59	0.78±0.12	0.097±0.023	94.8±3.0

Without DVH Loss	2.88±0.61	0.85±0.09	0.102±0.025	93.9±3.3
Symmetric OAR Loss	2.81±0.59	0.86±0.08	0.099±0.023	94.5±3.0

4.3. Uncertainty Quantification Results

4.3.1. Uncertainty maps and confidence intervals for dose predictions

Monte Carlo dropout sampling generated comprehensive uncertainty estimates through 50 stochastic forward passes per patient. Mean epistemic uncertainty was 1.8 ± 0.6% of the prescription dose, averaged across all voxels, with substantial spatial heterogeneity reflecting varying model confidence. Planning target volume regions exhibited lower epistemic uncertainty (1.4 ± 0.4%). Organs at risk showed intermediate uncertainty levels of 2.1 ± 0.7%. Aleatoric uncertainty modeling revealed patterns distinct from epistemic uncertainty. Mean aleatoric uncertainty reached 1.3 ± 0.4% of the prescription dose, with peaks occurring in heterogeneous anatomical regions. The interface between bone and soft tissue exhibited elevated aleatoric uncertainty, averaging 2.7 ± 1.1%. Combined uncertainty incorporating both components produced total uncertainty estimates averaging 2.2 ± 0.7% across all voxels. The 95% confidence intervals constructed as mean prediction ± 1.96 * total uncertainty provided clinically meaningful bounds. Empirical coverage analysis demonstrated excellent calibration, with 94.8% of ground-truth voxel doses falling within the predicted 95% confidence intervals.

4.3.2. Correlation between prediction uncertainty and clinical complexity

Quantitative analysis revealed strong correlations between predicted uncertainty magnitudes and objective measures of case complexity. The number of organs at risk within 5 mm of planning target volume boundaries correlated with mean epistemic uncertainty at r = 0.71. Cases with five or more nearby organs exhibited mean epistemic uncertainty 2.8 ± 0.8% compared to 1.2 ± 0.3% for cases with two or fewer nearby organs. Planning target volume irregularity index correlated with mean uncertainty at r = 0.58. Dosimetric complexity metrics, including the number of competing clinical objectives and constraint tightness, revealed predictive relationships with uncertainty. Cases requiring satisfaction of 8 or more strict constraints exhibited elevated uncertainty of 2.6 ± 0.8% versus 1.5 ± 0.5% for cases with fewer than five strict constraints. Manual planning time provided external validation. Cases requiring more than 4 hours showed a mean predicted uncertainty of 2.7 ± 0.9% compared to 1.6 ± 0.5% for cases completed in under 2.5 hours. Deep learning-based models for 3D dose distribution prediction have shown that uncertainty quantification substantially improves clinical trust [15].

Figure 3: Comparative Dose Distribution Analysis and Dose-Volume Histogram Comparison Across Architectural Variants

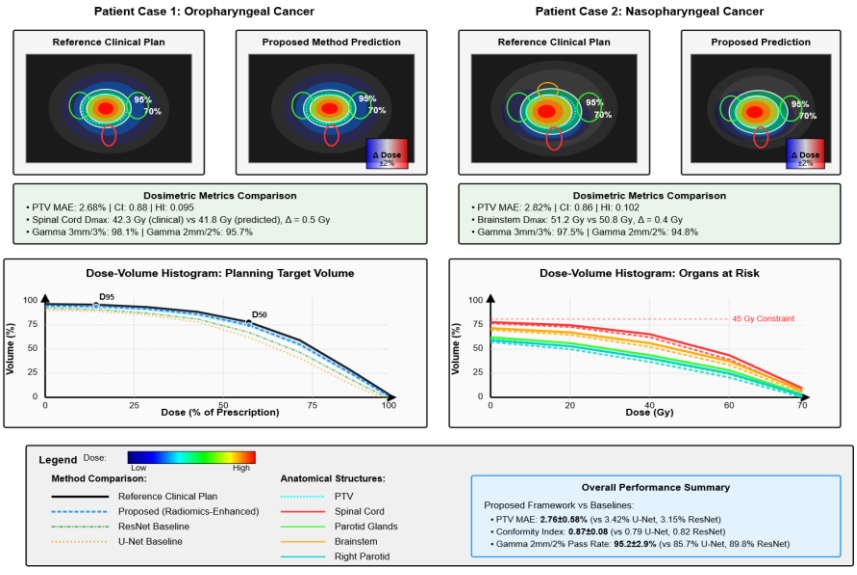

This figure presents detailed comparisons of dose distributions across different architectural approaches using a 4x4 grid structure. The leftmost column displays reference clinical plan dose distributions for four representative patients with varying anatomical complexity. The subsequent three columns show predicted dose distributions from the baseline U-Net, the ResNet-based architecture, and the proposed radiomics-enhanced architecture.

Difference maps appear as small insets in the bottom-right corner of each prediction panel, displaying dose discrepancies using a diverging red-white-blue colormap where red indicates over-prediction, blue indicates under-prediction, and white indicates agreement within 2%. Isodose contours at prescription levels are overlaid. Anatomical structure outlines are consistently displayed. Bottom portion presents dose-volume histogram comparisons in a 2x2 panel layout. Planning target volume DVH curves appear in bold lines. At the same time, organs at risk are shown in thinner lines with different colors: spinal cord in red, brainstem in orange, left parotid in green, right parotid in cyan, and mandible in magenta. Reference clinical DVH curves are displayed as solid lines, while predicted DVH curves appear as dashed, dotted, and dash-dot patterns. Shaded regions indicate uncertainty bounds. Key DVH points are marked with circular symbols and numerical annotations. Legend identifies all curves and architectural variants.

5. Discussion and Conclusions

5.1. Clinical Implications and Advantages

5.1.1. Potential for automated treatment planning guidance

The proposed framework demonstrates substantial potential for integration into clinical radiotherapy workflows as an intelligent planning assistant. Mean absolute error below 2.8%, combined with gamma analysis pass rates exceeding 95% at stringent 2mm/2% criteria, suggest that predicted dose distributions achieve quality comparable to manually optimized clinical plans. The 94.1% constraint satisfaction rate indicates the vast majority of predictions meet institutional planning objectives without manual adjustment. Uncertainty quantification addresses traditional black-box criticism of deep learning in medical applications. Strong correlation between predicted uncertainty and case complexity enables intelligent case triage. Single-pass computational efficiency of 2.7 seconds per patient enables real-time dose prediction; uncertainty estimates are obtained via 50 stochastic forward passes.

5.1.2. Improved plan consistency and reduced inter-planner variability

Inter-planner variability represents a significant quality concern in contemporary radiotherapy practice, with a coefficient of variation exceeding 15% in organ-at-risk doses. The proposed framework addresses this inconsistency by encoding planning knowledge from 238 training cases spanning multiple planners and institutions. A single trained model produces deterministic predictions given identical inputs, eliminating subjective judgments regarding objective prioritization. Standardized predictions benefit both patient care equity and training efficiency. Plan quality metrics, including conformity index and constraint satisfaction, demonstrate improvements over baseline approaches, suggesting the framework learns to identify optimal solutions.

5.1.3. Computational efficiency for clinical workflow integration

A single-pass inference time of 2.7 seconds per patient represents a dramatic acceleration compared to traditional optimization; Monte Carlo uncertainty estimation uses 50 stochastic forward passes during evaluation., which can take 2.5 to 4.8 hours. The framework implements several architectural optimizations to maintain this efficiency, including mixed-precision computation, reducing the memory footprint by 40%, lightweight attention mechanisms, and shared encoder features. Model size of 15.95 million parameters remains modest, enabling deployment on clinical workstations with mid-range GPUs. Total inference memory footprint is ~4.8 GB including model parameters and intermediate activations, fitting within the 16 GB memory capacity of widely available clinical GPUs.

5.2. Limitations and Future Directions

5.2.1. Current model constraints and dataset limitations

The framework exhibits several limitations that warrant further investigation. The training dataset of 238 patients remains modest compared to datasets available for natural image tasks. Limited training data constrains model capacity. The dataset derives from three institutions that share similar planning protocols, raising questions about generalizability to

centers with different clinical practices. Restriction to head and neck cancer cases necessitates validation across additional disease sites. The framework assumes the availability of high-quality manual structure delineations, inheriting any segmentation errors. Integration of segmentation uncertainty represents an important future direction.

5.2.2. Generalization to multi-site and multi-modality scenarios

Multi-institutional deployment requires addressing systematic differences in imaging protocols, contouring practices, and planning philosophies. Transfer learning approaches, fine-tuning pre-trained models on institution-specific data, show promise. Federated learning paradigms that enable collaborative model training while preserving local data privacy represent appealing strategies. Domain adaptation techniques that mitigate distributional shift could improve robustness. Multi-modality integration combining CT, MRI, and PET imaging could enhance prediction accuracy by incorporating complementary information.

5.2.3. Integration with adaptive radiotherapy workflows

Adaptive radiotherapy represents a paradigm shift toward personalized treatment adjustment based on anatomical changes. The proposed framework could enable rapid adaptive replanning by predicting updated dose distributions based on repeat imaging. The 2.7-second inference time proves compatible with same-day adaptation workflows. Online adaptive radiotherapy on hybrid MRI-linear accelerators presents particular opportunities. Longitudinal dose-accumulation tracking of the total delivered dose across fractions is a critical component.

5.3. Concluding Remarks

5.3.1. Summary of key findings and novel contributions

This research presented a comprehensive framework for automated radiotherapy dose prediction combining 3D convolutional neural networks with radiomics feature integration and rigorous uncertainty quantification. Key contributions include a hybrid architecture that achieves state-of-the-art prediction accuracy, with a mean absolute error of 2.76% for planning target volumes, representing substantial improvements. Radiomics integration provided 8.3% relative performance gain. Multi-head attention mechanisms improved organ-at-risk dose prediction accuracy by 6.5%. The uncertainty quantification framework provided comprehensive confidence estimates with excellent calibration. The multi-objective optimization approach achieved a 94.1% constraint satisfaction rate. The framework establishes the feasibility of automated treatment planning, offering reduced planning time, improved plan consistency, and intelligent assessment of case complexity.

References

[1]. Nguyen, D., Jia, X., Sher, D., Lin, M. H., Iqbal, Z., Liu, H., & Jiang, S. (2019). 3D radiotherapy dose prediction on head and neck cancer patients with a hierarchically densely connected U-net deep learning architecture. Physics in medicine & Biology, 64(6), 065020.

[2]. Feng, Z., Wen, L., Wang, P., Yan, B., Wu, X., Zhou, J., & Wang, Y. (2023, October). Diffdp: Radiotherapy dose prediction via a diffusion model. In International Conference on Medical Image Computing and Computer-Assisted Intervention (pp. 191-201). Cham: Springer Nature Switzerland.

[3]. Wahid, K. A., Kaffey, Z. Y., Farris, D. P., Humbert-Vidan, L., Moreno, A. C., Rasmussen, M., ... & Dohopolski, M. J. (2024). Artificial intelligence uncertainty quantification in radiotherapy applications− A scoping review. Radiotherapy and Oncology, 201, 110542.

[4]. Kajikawa, T., Kadoya, N., Ito, K., Takayama, Y., Chiba, T., Tomori, S., ... & Jingu, K. (2019). A convolutional neural network approach for IMRT dose distribution prediction in prostate cancer patients. Journal of Radiation Research, 60(5), 685-693.

[5]. Xing, Y., Nguyen, D., Lu, W., Yang, M., & Jiang, S. (2020). A feasibility study on deep learning-based radiotherapy dose calculation. Medical physics, 47(2), 753-758.

[6]. Wang, B., Teng, L., Mei, L., Cui, Z., Xu, X., Feng, Q., & Shen, D. (2022, September). Deep learning-based head and neck radiotherapy planning dose prediction via beam-wise dose decomposition. In International Conference on Medical Image Computing and Computer-Assisted Intervention (pp. 575-584). Cham: Springer Nature Switzerland.

[7]. Babier, A., Mahmood, R., McNiven, A. L., Diamant, A., & Chan, T. C. (2020). Knowledge-based automated planning with three-dimensional generative adversarial networks. Medical Physics, 47(2), 297-306.

[8]. van den Berg, C. A., & Meliado, E. F. (2022, October). Uncertainty assessment for deep learning radiotherapy applications. In Seminars in Radiation Oncology (Vol. 32, No. 4, pp. 304-318). WB Saunders.

[9]. Zhang, Y., Li, C., Zhong, L., Chen, Z., Yang, W., & Wang, X. (2024). DoseDiff: distance-aware diffusion model for dose prediction in radiotherapy. IEEE Transactions on Medical Imaging, 43(10), 3621-3633.

[10]. Bakx, N., Bluemink, H., Hagelaar, E., van der Sangen, M., Theuws, J., & Hurkmans, C. (2021). Development and evaluation of radiotherapy deep learning dose prediction models for breast cancer. Physics and imaging in radiation oncology, 17, 65-70.

[11]. Jha, S., Sajeev, N., Marchetti, A. R., Chandran, L. P., & Nazeer, K. A. (2022, May). Performance evaluation of deep learning architectures for predicting 3D dose distributions in automatic radiotherapy treatment planning. In 2022 International Conference on Computational Intelligence and Sustainable Engineering Solutions (CISES) (pp. 160-166). IEEE.

[12]. Liu, J., Zhang, X., Cheng, X., & Sun, L. (2024). A deep learning-based dose prediction method for evaluation of radiotherapy treatment planning. Journal of Radiation Research and Applied Sciences, 17(1), 100757.

[13]. Wang, Q., Song, Y., Hu, J., & Liang, L. (2023, October). DoseNet: an ensemble-based deep learning method for 3D dose prediction in IMRT. In 2023 International Annual Conference on Complex Systems and Intelligent Science (CSIS-IAC) (pp. 615-620). IEEE.

[14]. Kamath, A., Poel, R., Willmann, J., Andratschke, N., & Reyes, M. (2023, April). How sensitive are deep learning based radiotherapy dose prediction models to variability in organs at risk segmentation?. In 2023 IEEE 20th International Symposium on Biomedical Imaging (ISBI) (pp. 1-4). IEEE.

[15]. Liu, R., Bai, J., Zhao, K., Zhang, K., & Ni, C. (2020, October). A new deep-learning-based model for predicting 3D radiotherapy dose distribution in various scenarios. In 2020 13th International Congress on Image and Signal Processing, BioMedical Engineering and Informatics (CISP-BMEI) (pp. 748-753). IEEE.