Demo: Generative AI helps Radiotherapy Planning with User Preference

Riqiang Gao[1], Simon Arberet[1], Martin Kraus[1], Han Liu[1], Wilko F.A.R. Verbakel[2], Dorin Comaniciu[1], Florin C. Ghesu[1], Ali Kamen[1]

[1]Digital Technology and Innovation, Siemens Healthineers [2]Varian Medical Systems, Siemens Healthineers

NEURAL INFORMATION PROCESSING SYSTEMS 2025

Abstract

Background: Radiotherapy (RT) planning is highly complex and varies significantly across institutions. Most deep learning dose predictors do not enable user interface, potentially biasing models toward specific planning styles.

Method: We introduce **Flexible Dose Proposer (FDP)**, a novel generative model that predicts 3D dose distributions based on **user-defined preferences via interactive sliders**. These customizable preferences enable planners to prioritize specific trade-offs between organs-at-risk (OARs) and planning target volumes (PTVs), offering greater flexibility and personalization.

Result: FDP demonstrates superior DVH estimation accuracy and plan quality compared to Varian RapidPlan™ in some scenarios.

Introduction & Motivation

Radiotherapy Planning Background:

- RT is critical: ∼50% cancer patients receive RT treatment
- RT planning is complex, time-consuming, and subjective
- Involves multidisciplinary team with varying preferences

Current Limitations:

- Highly variable planning styles across institutions and planners
- Deep learning models trained on reference plans inherit specific biases
- Limited ability to interactively customize PTV/OAR trade-offs
- **RapidPlan limitations:** DVH-only predictions (no spatial dose), small training sets (∼50 plans), institution-specific models

Our Contributions:

- **Novel two-stage training framework** with foundational dose decoder for physically plausible outputs
- **First dose prediction model with interactive sliders** for real-time customization of PTV/OAR trade-offs
- **Clinical integration** with Eclipse™ treatment planning system
- **Superior performance:** better DVH estimation accuracy and plan quality

Method: Flexible Dose Proposer (FDP)

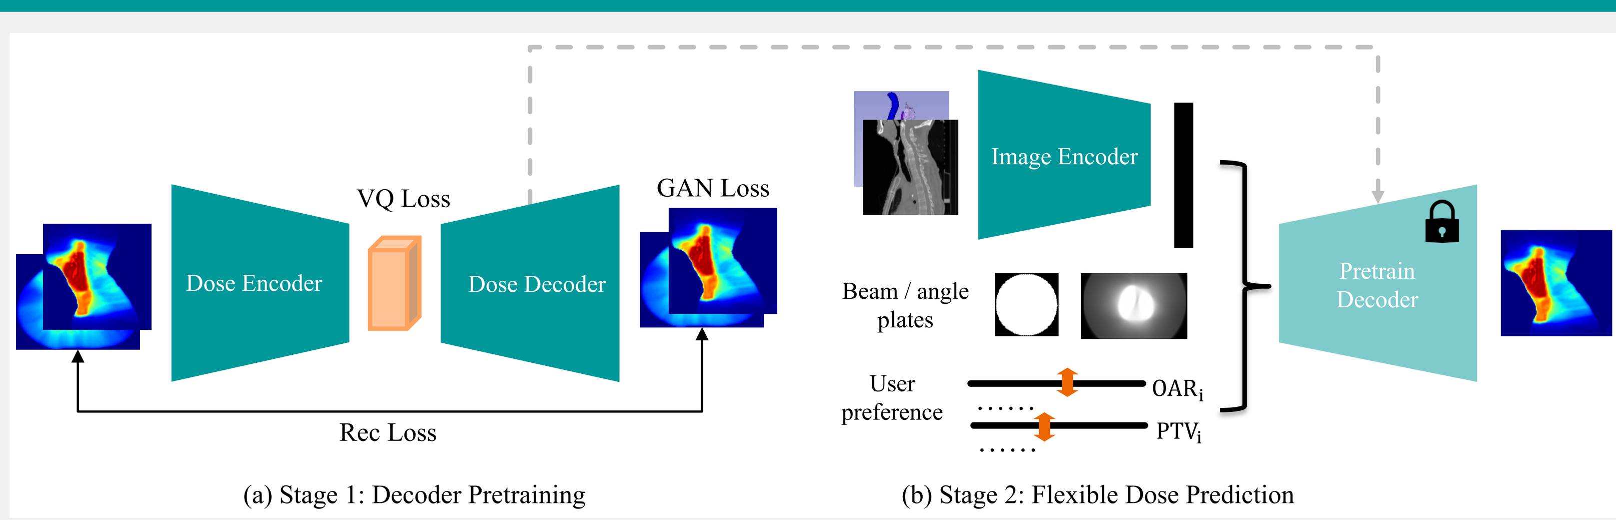

Figure: Two-stage training pipeline: Stage I learns realistic dose distributions via VQ-VAE pre-training; Stage II encodes user preferences for flexible prediction.

Stage I: Foundational Dose Decoder

- VQ-VAE architecture pre-trained on 31K doses from diverse sources
- Generates physically plausible dose distributions
- Stabilizes Stage II training and prevents unrealistic artifacts

Stage II: Flexible Prediction with User Preferences

- Multi-conditional encoder: CT, structures, beams, **user preference sliders**
- Adaptive Instance Normalization (AdaIN) modulates generation based on slider values
- Random sampling of preferences during training enables continuous control space
- One-step GAN generation for fast real-time inference

Experiments & Dataset

Dataset:

- 820 head-and-neck cancer cases across 6 cohorts
- Stage I: 31K doses for foundational decoder pre-training
- Stage II: 820 training, 103 validation, 113 test cases
- All test plans follow RapidPlan requirements for fair comparison

Evaluation Metrics:

- **Intra-patient:** Expected vs. achieved DVH differences (prediction accuracy)
- **Inter-patient:** DVH variability across patients (generalization robustness)
- Lower std → better estimation reliability
- **Plan Quality:** OAR mean dose, PTV Homogeneity Index (HI), Conformity Index (CI)

Baseline: Varian RapidPlan™ (widely used commercial model)

Key Results

DVH Prediction Accuracy (Superior Generalization):

- **Intra-patient (std):** FDP outperforms RapidPlan on **15/15 OARs**
- **Inter-patient:** FDP shows lower variability on **12/15 OARs**
- Better generalization to unseen patients
- Generalizable to different treatment mode

Plan Quality (After Optimization in Eclipse):

- **OARs:** FDP better on **14/15**, worse on **0/15**
- Top improvements: SubmandR (71%), OCavity (65%), SubmandL (60%)
- **PTVs:** FDP better on **1/6**, worse on **0/6**
- Maintains PTV quality while improving OAR sparing

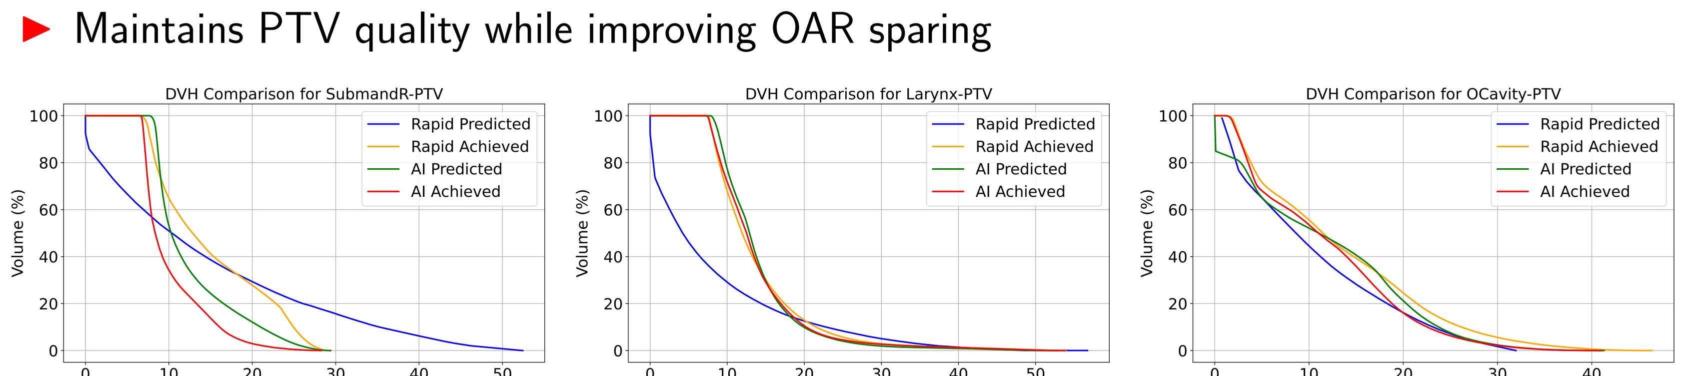

Figure: DVH comparison: FDP (blue/orange close) vs RapidPlan (larger gaps). Expected vs. achieved DVHs show FDP's superior prediction accuracy.

Ablation Study: Stage I Pre-training

Quantitative Impact:

MAE slightly reduces from 2.63 to **2.56** with Stage I pre-training

Qualitative Benefits:

- Reduces unrealistic boundary artifacts at PTV/OAR interfaces
- Improves anatomically plausible dose gradients
- Generates physically realistic dose distributions
- Stabilizes training by constraining to plausible dose space

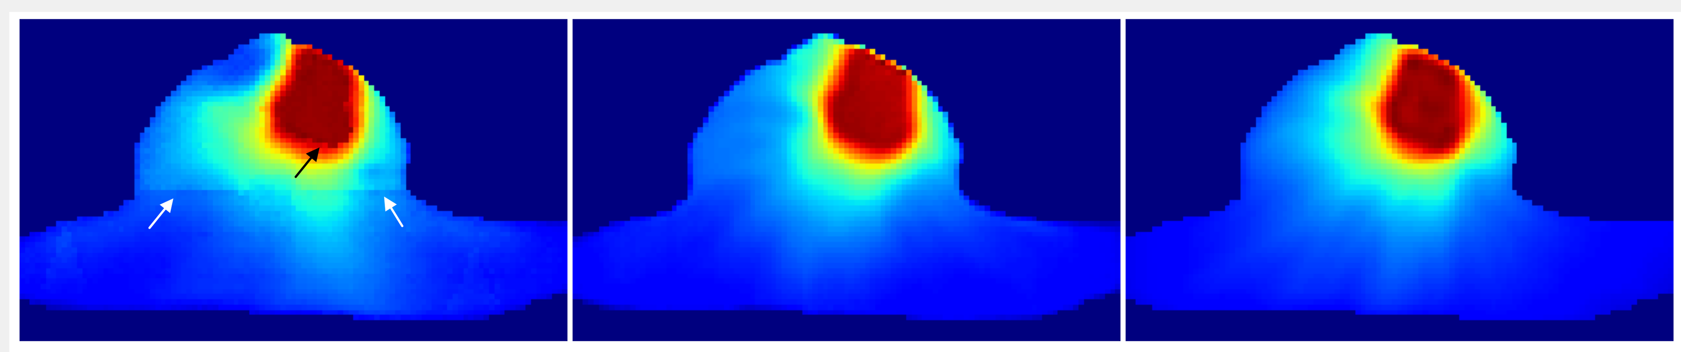

(1) w/o. pretrain decoder (2) w. pretrain decoder (3) reference

Figure: Dose distribution in different scenarios. Black/white arrows show boundary artifacts without Stage I. With Stage I, the predicted dose distribution is more physically plausible.

Interactive User Preference

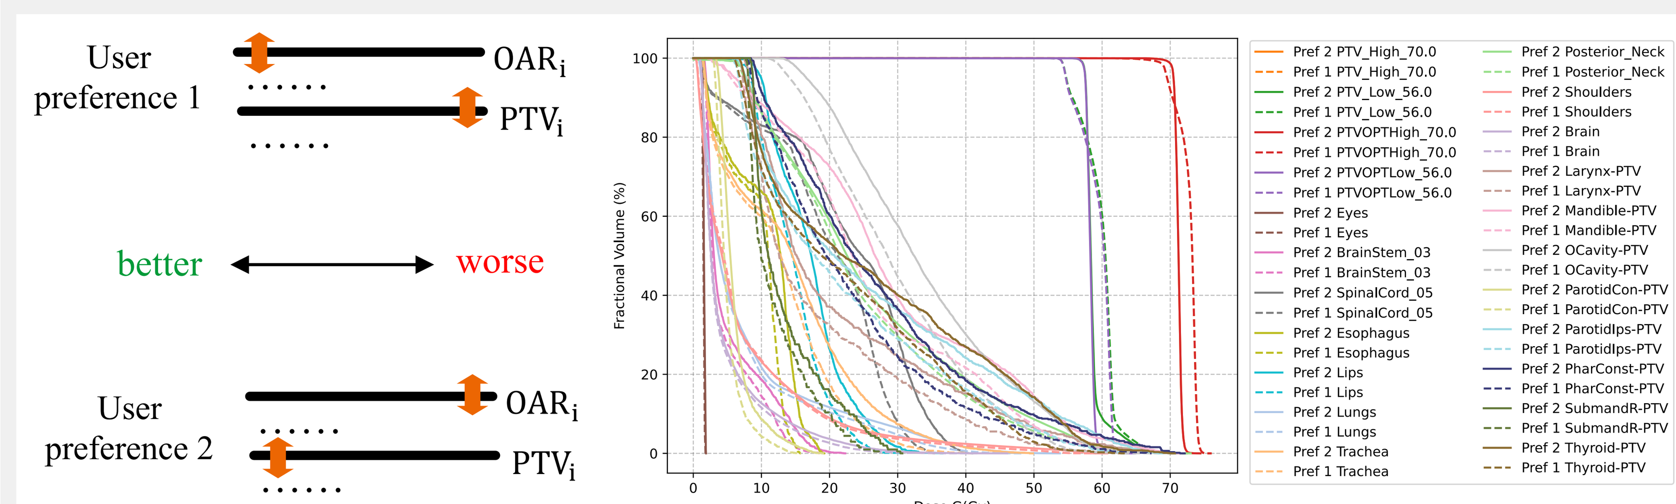

Figure: Interactive sliders control trade-offs: P1 (OAR-sparing focused) vs. P2 (PTV-homogeneity focused). DVH comparison is shown in the right panel.

Real-time Interaction:

- Continuous interpolation between preferences
- Model responds within seconds (full in <5s, DL model only 0.1s)
- Enables exploration of clinical trade-off space

Clinical Integration

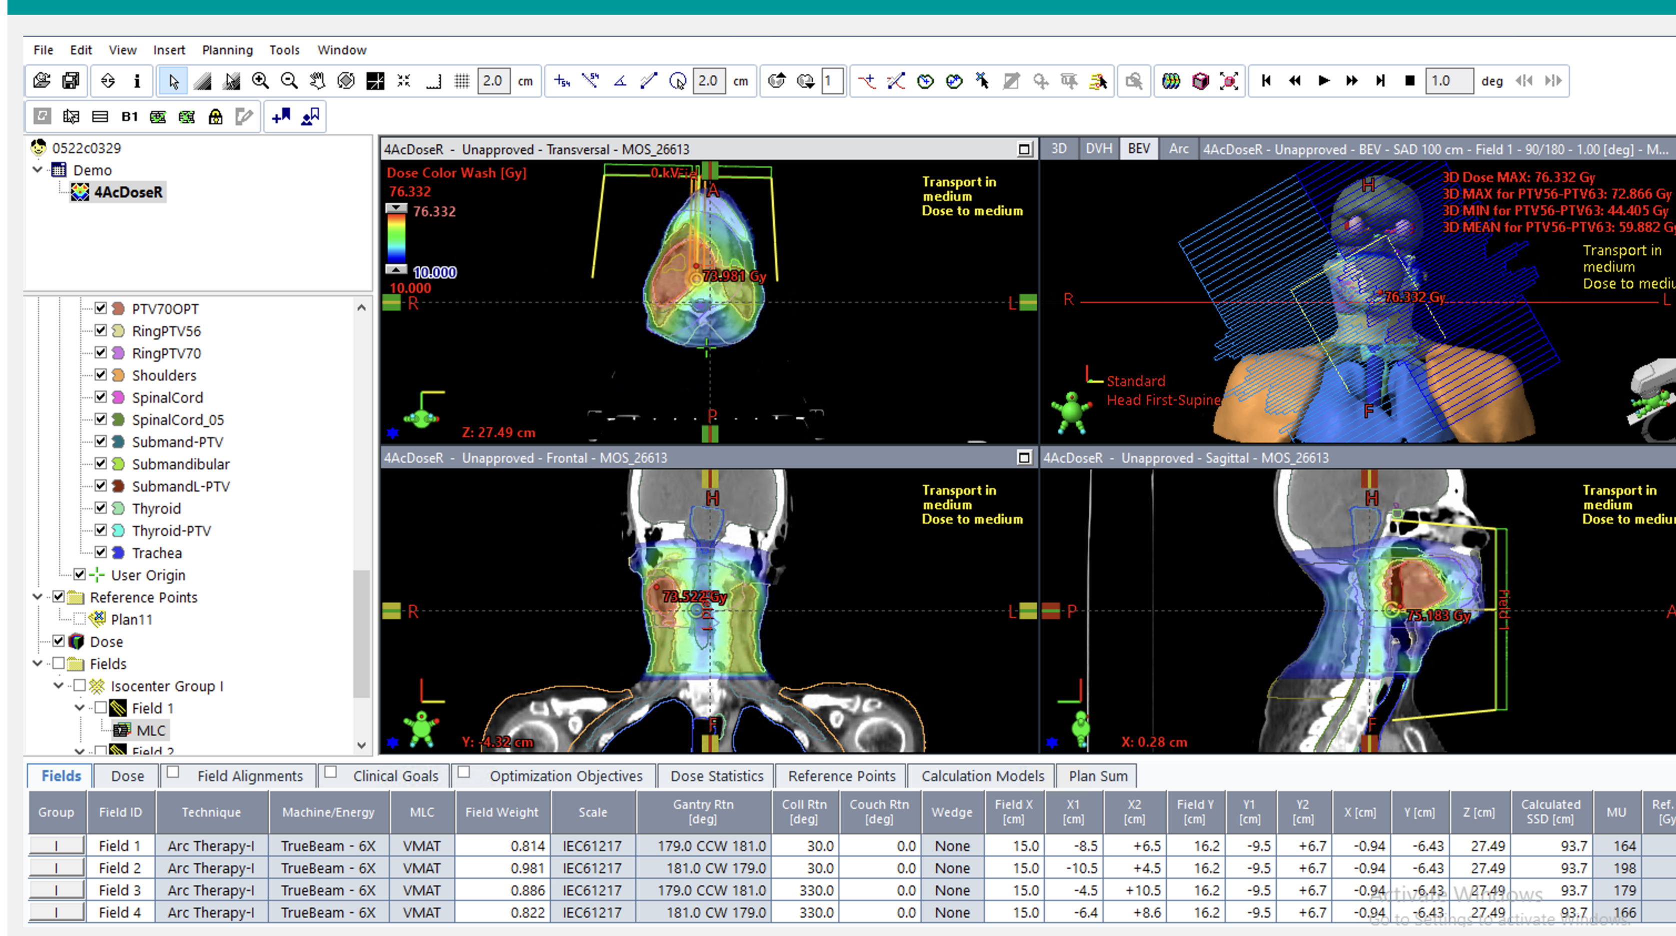

Figure: Integration with Eclipse™ planning system enables deliverable clinical plans.

Clinical Workflow:

- Load patient CT and structures into FDP interface
- Adjust user preference sliders to explore trade-off space
- Export AI dose to Eclipse™ for automatic objective extraction
- Plan optimization produces deliverable clinical plans
- Full quality assurance (QA) and safety verification

Conclusion & Future Work

Key Contributions Summary:

- Novel two-stage framework with foundational decoder ensuring physically plausible dose distributions
- First interactive dose prediction model with real-time slider-based customization
- Superior DVH estimation: Better or similar to RapidPlan on all 15 OARs
- Better plan quality: 14/15 OARs improved, 0/15 worse; PTV quality maintained
- Integration with Eclipse™ treatment planning system

Future Directions:

- Extension to other treatment sites (prostate, lung, breast, etc)
- Comprehensive multi-center clinical validation

Disclaimer: Research results not commercially available. Future availability cannot be guaranteed.