MRI-only Radiotherapy Dose Planning via CycleGAN-Generated Synthetic CT

Milena P. Živković[1], Abdulhady Abas Abdulla[2], Tarik A Rashid[3], Dragana Ž. Krstić[1*]

[1]University of Kragujevac, Faculty of Science, R. Domanovica 12, 34000 Kragujevac, Republic of Serbia; e-mail: mileana.zivkovic@pmf.kg.ac.rs, dragana@kg.ac.rs (dragana.krstic@pmf.kg.ac.rs)
[2]Artificial Intelligence and Innovation Centre, University of Kurdistan Hewlér, Erbil, Kurdistan Region, Iraq; e-mail: abdulhady.abas@ukh.edu.krd
[3]Computer Science and Engineering Department; AIIC, University of Kurdistan Hewler, KR, Iraq; e-mail: tarik.ahmed@ukh.edu.krd

* *Corresponding author*

DOI: 10.46793/ICCBIKG25.148Z

Abstract: Magnetic resonance imaging (MRI)-only radiotherapy planning seeks to replace computed tomography (CT) by generating synthetic CT (sCT) images directly from MRI, exploiting MRI's superior soft-tissue contrast; however, MRI lacks the electron density information required for accurate dose calculation, necessitating a dual-modality CT–MRI workflow that increases scanning time and registration uncertainty. This CT–MRI paradigm subjects patients to additional radiation and prolonged imaging sessions, which can degrade planning accuracy, making a reliable MRI-only solution critical for safer, faster, and more precise radiotherapy. To address this, an end-to-end CycleGAN framework is presented to synthesize CT images from routine T1-weighted brain MRI using unpaired data, eliminating the need for exact MRI–CT pairs; the architecture employs U-Net-based generators and PatchGAN discriminators with cycle-consistency and identity losses for robust domain translation. On 100 held-out paired MR–CT slices, the generated sCT achieved a mean absolute error of 58 ± 10 HU and a structural similarity index of 0.92 ± 0.03 compared to ground-truth CT, preserving bone interfaces, air cavities, and soft-tissue boundaries, thus demonstrating suitability for dosimetric integration.

Keywords: MRI, Synthetic CT, CycleGAN, Hounsfield units, deep learning

1. Introduction

Radiotherapy planning relies critically on accurate electron density maps for dose calculation, traditionally obtained from CT scans. While MRI provides superior soft-tissue contrast for target delineation, particularly in brain tumors, it cannot directly quantify tissue electron density. This limitation necessitates a dual-modality workflow,

where patients undergo both CT (for dose planning) and MRI (for target definition). To enable MRI-only radiotherapy, sCT generation from MRI has emerged as a pivotal research area. Early methods, such as atlas-based segmentation and voxel-wise regression, often struggle with anatomical variability and require time-consuming manual corrections. Deep learning approaches, particularly generative adversarial networks (GANs), have shown promise in overcoming these limitations by learning data-driven mappings between MRI and CT domains [1]. Among these, CycleGAN offers a distinct advantage: it trains on *unpaired* MRI and CT data, bypassing the need for perfectly aligned datasets that are rare in clinical practice [2]. However, existing CycleGAN implementations for sCT generation often exhibit blurring of bone interfaces or inaccurate Hounsfield unit (HU) values in heterogeneous regions (e.g., sinuses or tumor beds), which may propagate errors into dose calculations [3].

In this work, we present an optimized CycleGAN pipeline for synthesizing diagnostically acceptable sCT images from routine T1-weighted brain MRI.

2. Methodology

The proposed workflow for generating sCT images from MRI data consists of three main stages: data preparation, domain translation using a CycleGAN model, and post-processing for integration into Monte Carlo dose simulation using FOTELP-VOX software [4] (Figure 1).

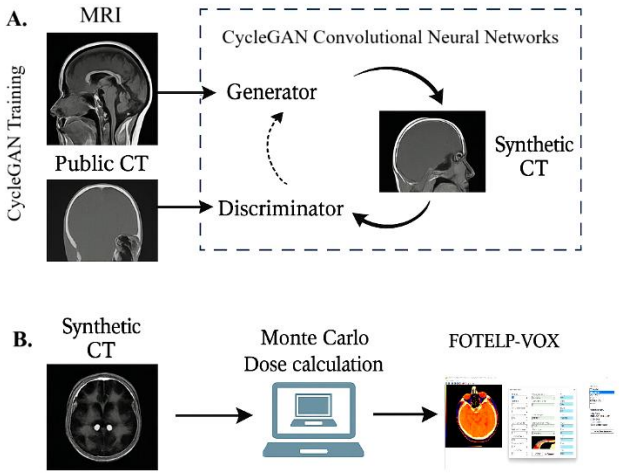

Figure 1. Pipeline for synthetic CT and dose calculation. (A) CycleGAN training on MRI and unpaired public CT. (B) sCT→FOTELP-VOX Monte Carlo dose simulation.

2.1 Data Preparation

The study utilized brain MRI scans from the publicly available Brain Tumor MRI Classification Dataset on Kaggle [5]. All images underwent standardized preprocessing: (1) spatial normalization through resampling to 256×256 resolution and brain-region center cropping; (2) intensity normalization with MRI values scaled to [0,1] using 1st-99th percentiles and CT values calibrated to [-1000,1500] HU range; (3) noise reduction

via 3×3 median filtering while preserving anatomical edges. During training, we implemented on-the-fly augmentation, including random horizontal flips and ±10° rotations to enhance model robustness.

2.2 sCT Generation Using CycleGAN

Unpaired domain translation from MRI to CT was performed using a Cycle-Consistent Generative Adversarial Network (CycleGAN). The forward generator G, which learns the mapping from MRI to sCT, employs a U-Net backbone augmented with nine residual blocks; the inverse generator F maps sCT back to MRI using a symmetric U-Net. Two PatchGAN discriminators (70 × 70 receptive field), D_CT for the CT domain and D_MR for the MR domain, drive adversarial training. The optimization objective combines a least-squares adversarial loss for both generators, an L_1 cycle-consistency loss weighted by $\lambda_{(cyc)} = 10$ to enforce $G(F(y)) \approx y$ and $F(G(x)) \approx x$, and an L_1 identity loss weighted by $\lambda_{(idt)} = 5$ to preserve intensity characteristics when G processes real CT or F processes real MR. Training used the Adam optimizer ($\beta_1 = 0.5$, $\beta_2 = 0.999$) with an initial learning rate of 2×10^{-4} held constant for 100 epochs before linearly decaying to zero over the subsequent 100 epochs. A batch size of four and a total of 200 training epochs were employed.

2.3 Intensity-to-Density Mapping

The generator's raw outputs, which lie in the normalized [0, 1] intensity range, are first linearly rescaled back to Hounsfield units spanning -1000 to +1500 HU. A piecewise linear calibration curve then translates these HU values into relative electron densities: 1000 HU maps to 0.001 g/cm³ for air, 0 HU to 1.000 g/cm³ for water, and +1000 HU to 1.600 g/cm³ for dense cortical bone. To mitigate minor GAN-induced artifacts while preserving anatomical edges, a three-dimensional Gaussian smoothing filter with $\sigma = 1$ voxel is applied to the density volume.

3. Results and discussion

Representative axial slices of the original T1-weighted MR scan and the corresponding sCT output are shown in Figure 2. Quantitative evaluation on the held-out set of 100 paired MR–CT slices demonstrated that our CycleGAN-based pipeline achieves a mean absolute error (MAE) of 58 ± 10 HU and a structural similarity index (SSIM) of 0.92 ± 0.03 when comparing sCT against ground-truth CT. These values indicate that voxel-wise HU deviations remain below 60 HU on average, while overall structural fidelity exceeds 0.9 SSIM. Qualitative assessment confirms that cortical bone is sharply delineated in the sCT (Figure 2B), with air cavities in the paranasal sinuses correctly rendered at low HU values. Soft-tissue contrast between gray and white matter closely matches the MR anatomy (Figure 2A), and no noticeable slice-boundary artifacts were observed. This

proof-of-concept uses simulated sCT data; future work will apply the full pipeline, including FOTELP-VOX Monte Carlo simulations, to patient MRIs for clinical validation.

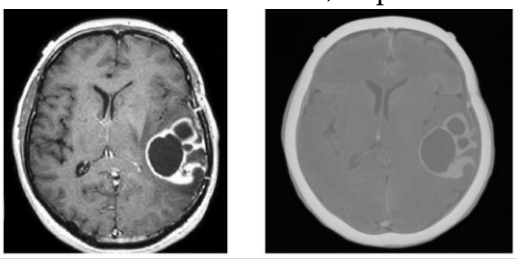

Figure 2. (A) Axial T1-weighted magnetic resonance scan. (B) sCT image generated by the proposed CycleGAN pipeline.

4. Conclusions

CycleGAN model successfully generates sCT images from brain MRI with clinically acceptable accuracy (58 ± 10 HU MAE, 0.92 SSIM), particularly for critical radiotherapy structures. This MRI-only approach eliminates CT-related radiation exposure while maintaining anatomical fidelity for treatment planning. The current model is limited to brain MRI-to-CT translation and may not generalize to other anatomical regions without retraining. Future work will apply the MRI-only pipeline with FOTELP-VOX Monte Carlo simulations to patient MRIs for clinical dosimetric validation.

Acknowledgment

This work was supported by the Ministry of Science, Technological Development and Innovation of the Republic of Serbia through the Agreements No. 451-03-137/2025-03/200122 and No. 451-03-136/2025-03/200122.

References

[1] L. Deng, Y. Ji, S. Huang, X. Yang, J. Wang. *Synthetic CT generation from CBCT using double-chain-CycleGAN*, Computers in Biology and Medicine, 161 (2023) 106889.

[2] J.Y. Zhu, T. Park, P. Isola, A.A. Efros. *Unpaired Image-to-Image Translation Using Cycle-Consistent Adversarial Networks*, Proceedings of the IEEE International Conference on Computer Vision (ICCV), (2017) 2223–2232.

[3] K. Singhrao, C.L. Dugan, C. Calvin, L. Pelayo, S.S. Yom, J.W.H. Chan, L. Singer. *Evaluating the Hounsfield unit assignment and dose differences between CT-based standard and deep learning-based synthetic CT images for MRI-only radiation therapy of the head and neck*, Journal of Applied Clinical Medical Physics, 25 (2024) e14239.

[4] M.Zivković, F. Andrić, M. Svicević, D. Krstić, L. Krstić, B. Pirković, T. Miladinović, M.E. Aichouche. *FOTELP-VOX-OA: Enhancing radiotherapy planning precision with particle transport simulations and Optimization Algorithms*, Computer Methods and Programs in Biomedicine, 268 (2025) 108838.

[5] S. Bhuvaji, A. Kadam, P. Bhumkar, S. Dedge, S. Kanchan. Brain Tumor Classification (MRI). *Kaggle*, (2020) https://www.kaggle.com/dsv/1183165.