Attention is Matter for Inclusiveness: Generating Synthetic CT for Patients with Hip Implants

Nico Camillo Zala[1,2*], Mariia Lapaeva[1,2*], Manuel Günther[2], Vittoria Banchieri[6], Jeanette Carmen Deck[2,3,5], Reto Sutter[2,5], Constantin von Deuster[3,4], Nicolaus Andratschke[1,2], Matthias Guckenberger[1,2], Stephanie Tanadini-Lang[1,2], and Riccardo Dal Bello[1,2]

[1] University Hospital Zurich, Zurich, Switzerland
[2] University of Zurich, Zurich, Switzerland
[3] Siemens Healthineers International AG, Zurich, Switzerland
[4] Balgrist Campus AG, Zurich, Switzerland
[5] Balgrist University Hospital, Zurich, Switzerland
[6] ETH Zurich, Zurich, Switzerland

Abstract. MR-based synthetic CT (sCT) generation has enabled MR-only radiotherapy workflows for selected anatomical sites. Despite that, patients with metal implants remain largely excluded from both clinically approved workflows and academic deep learning (DL) studies due to severe MR signal voids and metal-induced artifacts. In this work, we propose a dedicated DL framework for sCT generation in underrepresented patient cohorts with metal implants, with a focus on hip implants. Specifically, we introduce: (1) an automated method to derive and embed MR-based implant attention masks; (2) physically-guided MR augmentation strategy tailored to metal-induced artifacts and CT augmentation; (3) a systematic ablation study of multiple attention mechanisms and augmentation schemes using evaluation metrics targeting implant reconstruction. Experiments on two in-house pelvic datasets and the public SynthRAD dataset highlight clinically relevant benefits of the proposed methods, including for rare implant types. The DL-generated sCTs show strong dosimetric agreement, demonstrating clinical feasibility and supporting further development of inclusive workflows for populations previously considered unsuitable for MR-only radiotherapy. The source code is available at: https://github.com/medical-physics-usz/sct-implants.

Keywords: MR-only radiotherapy · Synthetic CT · Metal implants

1 Introduction

Modern radiotherapy treatment planning relies on multimodal imaging, primarily integrating computed tomography (CT) and magnetic resonance (MR) imaging. MR provides superior soft-tissue contrast for accurate tumor and organ-at-risk (OAR) delineation, while CT supplies electron density information required for dose calculation. Advances in deep learning (DL)-based synthetic CT

* These authors contributed equally. Corresponding author: mariia.lapaeva@uzh.ch

(sCT) generation has enabled clinical adoption of MR-only workflows for selected sites, particularly those less affected by motion, deformation, or air inclusions. MR-only radiotherapy reduces patient radiation exposure and introduces more cost-efficient workflows [24, 27].

Nevertheless, cohorts of patients with metal implants remain largely excluded from MR-only workflows due to severe MR signal voids and metal-induced artifacts [5] that result from strong magnetic susceptibility differences between metal and tissue, causing local field inhomogeneities, signal-loss and geometric distortions [10]. On MR images, metal regions appear as low-intensity "air-like" areas, which DL models often misinterpret during sCT generation [17], leading to incorrect Hounsfield unit (HU) assignments (-1000 HU for air instead of around 8 HU for soft tissue, 1200 HU for bone, or 3000 HU for metal) and consequently compromising radiation dose calculations. To date, sCT generation in the presence of metal implants remains largely unexplored, with prior work limited to bulk-density [12] or atlas-based approaches [31], and cone-beam CT DL studies [2]. Recent MR-only radiotherapy guidance further identifies DL hallucinations in the presence of metal as a major technical challenge [27].

At the same time, the prevalence of metal implants is increasing, especially in older populations [18, 28], with an estimated 10-15 % of prostate radiotherapy patients having implants. However, paired CT-MR datasets required for training sCT generation methods are scarce and subject to strict ethical constraints, leaving implant cases underrepresented in DL studies that depend on large, annotated datasets. Therefore, developing realistic data augmentation strategies for diverse implant types in underrepresented cohorts is of particular interest.

In this work, we propose an automated pipeline to enhance DL-based sCT generation for patients with metal implants, supporting inclusive MR-only radiotherapy. Starting with total hip replacements (THR) and intramedullary nails, the proposed framework is extendable to other implant types. We exploit attention mechanisms previously shown effective in medical imaging tasks [19], investigate implant-aware modules tailored for sCT generation and validate it on an external dataset. A further contribution is the derivation of metal attention masks directly from MR images, ensuring compatibility with MR-only workflows. To mitigate data scarcity, we are among the first to propose a paired data augmentation strategy for MR-to-CT translation involving implants, introducing physics-guided augmentation of MR metal-induced distortions with corresponding CT augmentation to model diverse implant appearances. An ablation study evaluates component impact using metal-related metrics for photon planning. Collectively, the results demonstrate that the proposed framework improves baseline DL approaches, enhances anatomical consistency, and enables clinically meaningful sCT generation for underrepresented cohorts.

2 Method

We combine three components to enable MR-based sCT generation of patients with metal implants. First, data augmentation expands training set. Afterward, we create a metal mask directly from MR voids, which we then utilize to guide sCT generation with various attention mechanisms, using Pix2Pix [9] as baseline, due to its proven robustness in the task [3, 5].

2.1 Implant-Specific Data Augmentation

Donor-Receiver (DR) Augmentation. The first proposed augmentation strategy follows a naïve implant-transfer approach. A segmented hip implant from a donor patient (with implant) is transferred to a receiver patient (without implant) in both CT and MR. The femur bone contour serves as the anatomical reference for alignment due to its rigid structure and low inter-patient variability.

Principal component analysis is applied to donor and receiver femur contours to estimate the primary anatomical axis, where the first eigenvector corresponds to the femoral neck direction. For the femoral centroids, a rigid transformation comprising rotation and translation is computed to align the donor's femur to the receiver's, with rotation derived via Rodrigues' formula [21]. This transformation is then applied to transfer both the implants from the donor CT and the corresponding signal void from the donor MR to the receiver CT-MR pair. The implant mask is extracted from donor CT by thresholding HU>2000, retaining only metal voxels overlapping the femur to suppress unrelated metallic objects. After transformation, the implant is inserted into the receiver CT by overwriting voxels with original metal density. Adjacent bone voxels are reassigned to soft-tissue values to refine the augmentation. The MR void mask is defined in a 2.5 cm neighborhood around the implant CT mask. Within this region, voxels with normalized MR intensity $I_{\mathrm{MR}} \in [0, 1]$ are classified as void when $I_{\mathrm{MR}} < 0.1$. This threshold was empirically determined from internal dataset analysis and applied unchanged to the external dataset for validation. The proposed pipeline supports unilateral or bilateral insertion and cross-donor combinations, enabling a larger number of anatomically plausible augmented samples.

Physics-Guided (Ph) Augmentation of Metal-Induced Distortions. Second augmentation strategy differs only in the MR component and leverages physical properties of metal-induced distortions. Susceptibility-driven intravoxel dephasing in the periprosthetic region leads to pronounced signal voids [6]. Unlike prior metal artifact simulators developed for spin-echo imaging using spectral binning [11], our method is one of the first to operate directly on real GRE/Dixon images while remaining sequence-configurable. To model the metal-related MR signal loss, we employ a sinc-product intravoxel dephasing model driven by local $B0$ field gradients [4, 15]. The off-resonance (off-frequency) map Δf is computed using the implant CT mask and the magnetic susceptibility values of the corresponding materials [1]. Susceptibility values for common implant alloys, cobalt chromium and titanium, are taken from the literature [22] to simulate the implant induced perturbation of the local magnetic field. From the resulting distorted $B0$ field map Δf, the spatial field gradients are computed:

$$g_x = \frac{\partial \Delta f}{\partial x} \text{ [Hz/mm]}, g_y = \frac{\partial \Delta f}{\partial y} \text{ [Hz/mm]}, g_z = \frac{\partial \Delta f}{\partial z} \text{ [Hz/mm]} \tag{1}$$

The signal attenuation is modeled under the assumption of a piecewise linear field variation within each voxel. For a gradient g along a voxel dimension Δ, accumulated phase dispersion at echo time TE yields attenuation: $A = |\text{sinc}(g * \Delta * \text{TE})|$. Thus, the total 3D intravoxel dephasing attenuation is:

$$A_{total} = |\text{sinc}(g_x * \Delta x * \text{TE})| \cdot |\text{sinc}(g_y * \Delta y * \text{TE})| \cdot |\text{sinc}(g_z * \Delta z * \text{TE})| \quad (2)$$

The final attenuation mask is multiplied pixel-wise with the MR image to simulate signal voids. The processing of CT follows the DR strategy, described above.

2.2 Implant-aware Synthetic CT Generation

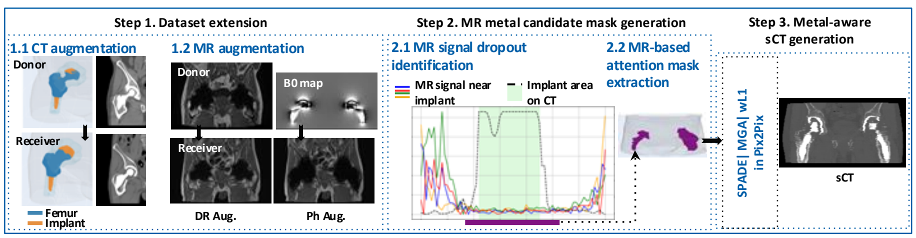

Fig. 1. Overview of proposed framework

MR-only Metal Candidate Mask Extraction. The macroscopic magnetic field inhomogeneity due to the metal causes a dephasing of the spins [4, 15, 32], leading to progressive signal loss and complete dropout near the implant. We exploit this property to derive an implant surrogate directly from MR via thresholding, without CT supervision. A binary MR candidate mask is obtained from the normalized in-phase MR image ($I_{\text{MR}} \in [0, 1]$) by thresholding $I_{\text{MR}} < 0.025$, where the threshold optimized via parameter search maximizing the overlap of MR-derived candidate mask with CT metal masks. Connected components $<$5000 voxels are removed to suppress spurious detections. After morphological smoothing, two geometric constraints are applied. First, regions with a flatness ratio < 0.1 (min/max side length of the 3D bounding box) were excluded, removing thin, sheet-like structures typical for air pockets rather than metal-induced voids. Second, hollow ring-shaped regions around cortical bone are discarded [14], as metal-induced voids form solid 3D structures. The resulting mask M highlights plausible implant locations (Fig. 1).

Attention-Guided Implant-Aware Synthetic CT Generation. Diverse attention-based strategies are explored to incorporate MR-derived metal candidate masks and improve metal-affected sCT generation. As the first method, we explore adding a penalty term α to the conditional L_1 loss of Pix2Pix sCT generator $G(x_{MR}, z)$ [9] to emphasize the voxels in the metal mask M:

$$wL_1(G) = E_{(x_{MR},x_{CT})\sim p_{data}(x_{MR},x_{CT})}\left[\left\|(1+\alpha M)\odot(x_{CT}-G(x_{MR},z))\right\|_1\right] \quad (3)$$

where z is random noise and (x_{MR}, x_{CT}) are MR-CT training pairs. The second approach explores spatially-adaptive denormalization (**SPADE**) layers that are integrated *within* ResNet blocks of the Pix2Pix generator to enable spatial conditioning on the MR-derived metal candidate mask. The design is based on [20], who showed that naïvely adding a segmentation mask as an input often fails to preserve semantic information across the network. The third algorithm, inspired by [23], applies metal-guided attention (**MGA**), which uses the spatial mask to guide per-location gating, to condition the Pix2Pix generator. The MR-derived metal mask is encoded into a key representation and combined with MR features to produce a gating map that modulates feature flow *before* ResNet blocks.

3 Experiments and Results

3.1 Datasets

We evaluate a total of three datasets. **Dataset 1** contains in-house MR-CT pairs from 20 male patients with pelvic tumors and **THR**. Data were divided into training (16 patients, 1682 axial slices) and test (4 patients, 490 slices) sets. For augmentation, image volumes of 64 patients without implants (6874 slices) were included. MR (T1 Dixon in-phase, 1.5T, Magnetom Sola) and CT (120 kVp, Somatom Definition AS) were rigidly registered, resampled to $1.63 \times 1.63\text{mm}^2$, cropped to 256×256 pixels, clipped (for CT to $[-1024, 3000]$ HU, for MR only top 2%), and scaled to $[0, 1]$. Five-fold cross-validation preserving implant laterality is used to evaluate performance stability. **Dataset 2** contains **rare implants.** MR-CT pairs of 5 patients with intramedullary nails were available internally, reflecting the scarcity of such cases in routine clinical cohorts. We use this dataset to assess feasibility under extreme data limitations, with a split into training (3 patients, 189 slices) and test (2 patients, 169 slices) sets. Data were acquired and preprocessed identically to Dataset 1, with augmentation using a subset of 24 implant-free cases from Dataset 1 (2851 slices). Ethics protocol for both datasets is BASEC 2023-01508. **Dataset 3** is used for **external validation**. The publicly available SynthRAD2023 [25] dataset (Center A) was manually curated to identify THR cases and split into training (4 patients, 450 slices) and test (3 patients, 330 slices), with additionally identified implant-free cases (32 patients, 3007 slices) used for augmentation. MR (T1 SPGR, 1.5/3T, Ingenia) and CT (Brilliance Big Bore, 120 kVp) are preprocessed similarly to internal datasets. Tagging can be found in our repository to ensure reproducibility.

3.2 Ablation Study Configuration and Evaluation Metrics

Datasets 1 and 2 are evaluated without and with DR and Ph augmentation. For each setting, we evaluate wL1, SPADE, MGA to disentangle the impact of attention mechanisms, dataset sizes and augmentation strategies. Dataset 3 is used for testing method transferability (training and test on Dataset 3) and cross-dataset generalization (training on Dataset 1, test on Dataset 3, denoted with $^\diamond$). The configuration without augmentation or attention is defined as baseline.

sCT quality is assessed by image similarity and dosimetric accuracy metrics. Image similarity is quantified using the mean absolute error (MAE), previously shown to correlate with dosimetry [7, 16], and computed in 3 regions: (i) MAE_{in}: inside the expanded implant mask T, defined by thresholding CT$>$ 2000 HU with 2 cm margin; (ii) MAE_{out}: outside T but within the body contour; and (iii) MAE_{B}: the entire body contour. We use the true positive rate (TPR), defined as the proportion of reference metal-containing axial slices in which the sCT correctly predicted metal within T, whereas the false positive rate (FPR) denotes the proportion of all slices in which metal was incorrectly predicted outside T. Dosimetric accuracy is evaluated by comparing dose-volume histogram (DVH) parameters of photon plans (5×8 Gy), optimized on sCT and transferred to CT using matRad [30] (pencil beam, 2 mm^3 grid, plan is available in repository). DVH metrics [7, 8, 16] include D95%, D2%, Dmean for prostate as planning target volume (PTV) and D2% and Dmean for OARs (bladder; rectum for Datasets 1-2, colon for Dataset 3), calculated in partial-arc (avoiding implants) and full-arc (covering 360°) scenarios. DVH metrics below 8 Gy are excluded due to limited clinical relevance. Agreement is further quantified using 3D local gamma analysis [29] with 2%/2 mm criteria at 90%/50% dose thresholds.

3.3 Results and Discussion

sCT Generation for Dataset 1 without augmentation or attention are already robust: 4 of 5 cross-validation splits have all mean DVH deviations <0.95% within clinically acceptable range [13] of ±2%, with only one minor outlier for Rectum_{D2} (2.2%, 0.8 Gy), mean MAE_{B} of 73.5 HU, TPR of 91%, FPR of 1.9%, gamma D90 and D50 (2%/2 mm) of 99.9%. In contrast, the 5th split showed multiple dosimetric outliers due to large air inclusions as shown in Fig. 2 (yellow arrow) that caused the baseline model to mispredict high-density regions. As shown in [16], small test sets may mask potential outliers, producing overly optimistic performance estimates that only emerge when more challenging cases or larger datasets are evaluated. This motivates focused analysis of the 5th split as the primary training configuration to capture relevant variability.

Tab. 1 and Fig. 2 demonstrate that attention or augmentation alone is insufficient for challenging cases. The combined wL_1 loss and Ph augmentation improves generation by reducing FPR and eliminating PTV DVH outliers, while mitigating metal–air ambiguity. Residual rectum outliers (D2: 3.9%/1.3 Gy, Dmean: 2.1%/0.3 Gy, full-arc) likely reflect MR-CT registration uncertainty and may benefit from more advanced CT augmentation. DR augmentation can misalign under-pronounced anatomical variability. Preoperative surgical planning tools [26] could further aid rare implant simulation, but are beyond the scope of this study. End-to-end inference per 3D volume is 2 min, supporting clinical feasibility.

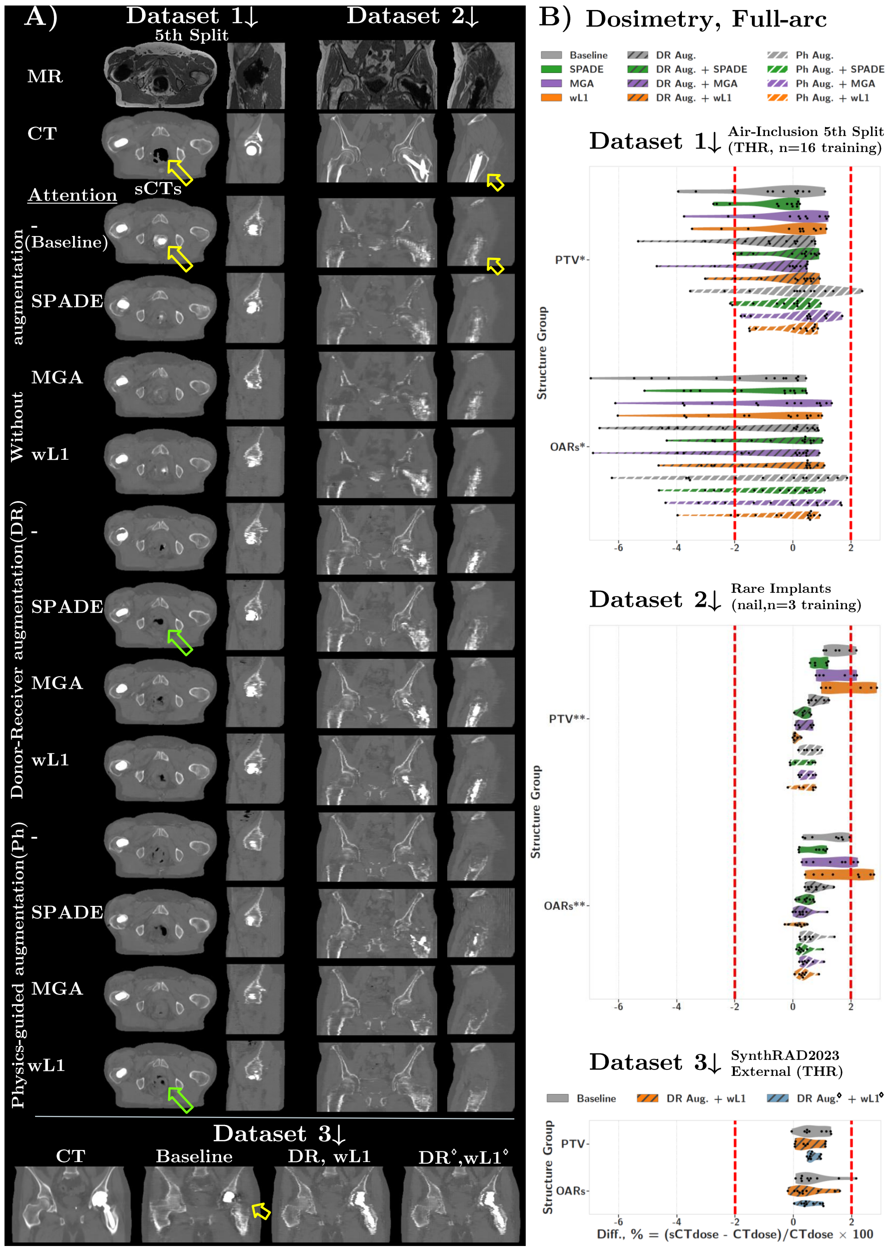

Fig. 2. Representative examples of sCT volumes generated from diverse datasets (A) and differences in DVH parameters between plans calculated on CT and generated sCT (B). Only the combination of augmentation and attention prevents erroneous metal predictions in air regions (A, Dataset 1, green arrows) compared to baseline. (B): Red lines highlight the clinically relevant range of mean DVH deviations within ±2%. PERMANOVA, Dataset 1:$^{*}p < 0.05, R^2 \sim 0.06$; Dataset 2:$^{**}p < 0.01, R^2 \sim 0.70$.

Table 1. Quantitative comparison of sCT quality. Rectum$_{D2}$ is most affected by outliers; #o are DVH outliers with $|\Delta| > 2\%$ across all parameters; DVH diff. and Gamma reports mean ± SD; best per dataset values in bold. For Dataset 3, Colon$_{D2}$ is reported instead of Rectum$_{D2}$ due to contour availability.

Methods		**Image similarity**					**DVH diff.,partial-arc**			**Gamma (2/2)**	
Atten.	**Aug.**	**MAE$_B$** (HU↓)	**MAE$_{in}$** (HU↓)	**MAE$_{out}$** (HU↓)	**TPR** (%↑)	**FPR** (%↓)	**PTV$_{D95}$** (%↓)	**Rectum$_{D2}$** (%↓)	**#o** (↓)	**D90** (%↑)	**D50** (%↑)
Dataset 1: Total hip replacement											
–	–	73.17	334.28	50.20	92.6	13.0	-0.67±1.29	-3.22±3.78	9	91.2±12	97.5±3
SPADE	–	67.29	301.42	47.17	91.9	0.50	-0.76±1.01	-2.58±3.00	8	92.5±9	98.5±2
MGA	–	73.54	348.46	49.73	90.1	4.80	0.19±1.14	-2.08±3.58	6	93.9±9	98.3±2
wL1	–	67.37	287.82	47.94	93.7	4.80	-0.11±0.97	-2.27±3.22	6	93.5±10	98.6±2
–	DR	**64.72**	286.63	**44.40**	91.9	2.50	-0.47±0.96	-2.59±3.75	8	90.8±11	98.1±2
SPADE	DR	69.60	313.63	48.80	91.2	2.50	-0.12±1.00	-1.88±2.72	5	95.3±5	99.3±1
MGA	DR	68.04	322.69	45.99	**95.1**	3.67	-0.31±1.07	-2.50±3.63	6	93.6±9	98.4±2
wL1	DR	66.60	**279.74**	47.99	93.0	1.83	-0.10±1.13	-1.96±2.87	6	94.8±6	99.0±1
–	Ph	70.87	332.43	47.41	87.7	5.70	-0.04±1.25	-2.54±3.78	7	90.3±7	97.6±1
SPADE	Ph	71.75	317.07	49.61	93.7	0.50	-0.49±0.90	-1.85±2.72	6	94.6±8	98.8±1
MGA	Ph	74.30	325.62	51.67	89.1	1.83	**0.05±0.7**	-1.97±3.02	5	94.5±8	99.2±1
wL1	Ph	67.78	315.50	45.88	90.2	**0.46**	-0.07±1.06	**-1.21±2.4**	**3**	**96.5±5**	**99.4±1**
Dataset 2: Rare implants											
–	–	79.87	322.32	67.75	9.0	0	1.42±0.92	1.23±1.10	2	94.7±7	97.3±3
SPADE	–	64.78	306.43	52.55	12.4	0	0.66±0.74	0.52±0.67	0	100±0	100±0
MGA	–	77.46	323.65	65.01	13.5	2.37	1.29±1.26	1.10±1.34	3	94.1±8	98.0±2
wL1	–	84.75	316.68	72.74	18.0	4.73	1.78±1.47	1.64±1.69	7	79.4±29	91±12
–	DR	62.40	281.20	50.99	31.5	0	0.63±0.06	0.52±0.27	0	100±0	100±0
SPADE	DR	**59.82**	279.16	**48.60**	4.5	0	0.23±0.14	0.18±0.21	0	100±0	100±0
MGA	DR	60.18	**269.34**	49.23	**45.0**	0	0.26±0.04	**0.23±0.08**	0	100±0	100±0
wL1	DR	60.68	270.26	49.68	42.7	0	**-0.10±0.2**	-0.26±0.40	0	100±0	100±0
–	Ph	63.36	296.20	51.22	0	0	0.44±0.01	0.36±0.18	0	100±0	100±0
SPADE	Ph	72.47	292.48	61.06	34.8	0	0.16±0.14	0.24±0.12	0	100±0	100±0
MGA	Ph	63.27	298.47	51.33	12.4	0	0.24±0.04	0.26±0.21	0	100±0	100±0
wL1	Ph	61.50	294.83	49.59	1.1	0	0.30±0.05	0.14±0.43	0	100±0	100±0
Dataset 3: External dataset SynthRAD2023											
–	–	68.37	370.26	54.37	69.7	6.06	0.72±0.33	0.29±0.18	0	99.9±0	99.9±0
wL1	DR	**53.66**	**355.69**	**39.26**	87.2	0.61	0.40±0.42	**0.10±0.1**	0	100±0	**100±0**
wL1$^{\diamond}$	DR$^{\diamond}$	61.22	357.63	47.41	**96.7**	0.61	**0.46±0.2**	0.38±0.39	0	100±0	99.9±0

The **Ablation Study on Dataset 2** shows that DR augmentation combined with attention enables reliable sCT generation even for extremely rare implant cohorts (3 training cases), improving metal-region reconstruction by reducing MAE$_{in}$ and increasing TPR, while maintaining DVH deviations <2% for partial- and full-arc plans. Full-arc evaluation is particularly relevant, as it probes robustness in scenarios where the tumor is located adjacent to an implant and the beams traverse the implant region. Importantly, this demonstrates the method's potential for inclusive MR-only workflows, supporting underrepresented patient cohorts. Nevertheless, results should be interpreted cautiously, as validation is limited to a single-scanner, two-patient test set with photon therapy plans only.

The **Evaluation on External Dataset 3** confirms feasibility and transferability of the proposed workflow across MR protocols, with MR-only candidate mask extraction and DR augmentation applied without additional fine-tuning or parameter adjustment, demonstrating DVH deviations within the clinically acceptable range of ± 2%.

4 Conclusion

In this paper, we demonstrated the feasibility of the proposed pipeline to enhance DL-based sCT generation in patients with metal implants, supporting more inclusive MR-only radiotherapy workflows. By combining MR-deduced implant-aware attention with multiple proposed augmentation strategies, we showed improved anatomical consistency in metal-affected regions and robust sCT generation for underrepresented cohorts. Across all evaluated datasets and types of metal implants (THR, intramedullary nails), all mean DVH differences stayed below 2%, indicating clinically relevant dosimetric agreement. While our study only focused on hip implants, the proposed framework is extendable to other implant types, MR sequences and treatment sites.

Disclosure of Interests. The Department of Radiation Oncology, University Hospital Zurich has teaching and research agreements with Siemens Healthineers, which however did not influence this work. Jeanette Carmen Deck and Constantin von Deuster are employed by Siemens Healthineers AG.

References

1. Bouwman, J.G., Bakker, C.J.G.: Alias subtraction more efficient than conventional zero-padding in the fourier-based calculation of the susceptibility induced perturbation of the magnetic field in MR. Magnetic Resonance in Medicine **68**(2) (2012)
2. Cao, Z., Gao, X., Chang, Y., Liu, G., Pei, Y.: A novel approach for eliminating metal artifacts based on MVCBCT and CycleGAN. Frontiers in Oncology **12** (2022)
3. Dayarathna, S., Islam, K.T., Uribe, S., Yang, G., Hayat, M., Chen, Z.: Deep learning based synthesis of MRI, CT and PET: Review and analysis. Medical Image Analysis **92** (2024)
4. De Guio, F., Benoit-Cattin, H., Davenel, A.: Signal decay due to susceptibility-induced intravoxel dephasing on multiple air-filled cylinders: MRI simulations and experiments. MAGMA **21**(4) (2008)
5. Fusella, M., Andres, E.A., Villegas, F., Milan, L., Janssen, T., Dal Bello, R., Garibaldi, C., Placidi, L., Cusumano, D.: Results of 2023 survey on the use of synthetic computed tomography for magnetic resonance imaging-only radiotherapy: Current status and future steps. Physics and Imaging in Radiation Oncology **32** (2024)
6. Hargreaves, B.A., Worters, P.W., Pauly, K.B., Pauly, J.M., Koch, K.M., Gold, G.E.: Metal-induced artifacts in MRI. American Journal of Roentgenology **197**(3) (2011)
7. Huijben, E.M., Terpstra, M.L., Pai, S., Thummerer, A., Koopmans, P., Afonso, M., Van Eijnatten, M., Gurney-Champion, O., Chen, Z., Zhang, Y., et al.: Generating synthetic computed tomography for radiotherapy: SynthRAD2023 challenge report. Medical Image Analysis **97** (2024)
8. International Commission on Radiation Units and Measurements: Prescribing, recording, and reporting intensity-modulated photon-beam therapy (IMRT). Tech. Rep. 83, ICRU (2010)

9. Isola, P., Zhu, J.Y., Zhou, T., Efros, A.A.: Image-to-image translation with conditional adversarial networks. In: Conference on Computer Vision and Pattern Recognition (2017)
10. Jungmann, P.M., Agten, C.A., Pfirrmann, C.W., Sutter, R.: Advances in MRI around metal. Journal of Magnetic Resonance Imaging **46**(4) (2017)
11. Keskin, K., Sanson, A.R., Hargreaves, B.A., Nayak, K.S.: Open-source simulator of imaging near metal at arbitrary magnetic field strengths. Magnetic Resonance in Medicine **95**(4), 2370–2383 (2026)
12. Koivula, L., Seppälä, T., Collan, J., Visapää, H., Tenhunen, M., Korhonen, A.: Synthetic computed tomography based dose calculation in prostate cancer patients with hip prostheses for magnetic resonance imaging-only radiotherapy. Physics and Imaging in Radiation Oncology **27** (2023)
13. Korsholm, M.E., Waring, L.W., Edmund, J.M.: A criterion for the reliable use of MRI-only radiotherapy. Radiation Oncology **9**(1) (2014)
14. Krug, R., Larson, P.E.Z., Wang, C., Burghardt, A.J., Kelley, D.A., Link, T.M., Zhang, X., Vigneron, D.B., Majumdar, S.: Ultrashort echo time MRI of cortical bone at 7 tesla field strength: a feasibility study. Journal of Magnetic Resonance Imaging **34**(3) (2011)
15. Lam, F., Sutton, B.P.: Intravoxel B0 inhomogeneity corrected reconstruction using a low-rank encoding operator. Magnetic Resonance in Medicine **84**(2) (2020)
16. Lapaeva, M., La Greca Saint-Esteven, A., Wallimann, P., Andratschke, N., Guckenberger, M., Günther, M., Tanadini-Lang, S., Dal Bello, R.: A comprehensive comparative study of generative adversarial network architectures for synthetic computed tomography generation in the abdomen. Medical Physics **52**(8) (2025)
17. Law, M.W.K., Tse, M.Y., Ho, L.C.C., Lau, K.K., Wong, O.L., Yuan, J., Cheung, K.Y., Yu, S.K.: A study of Bayesian deep network uncertainty and its application to synthetic CT generation for MR-only radiotherapy treatment planning. Medical Physics **51**(2) (2024)
18. Liang, Y., Xu, H., Tang, W., Du, X.: The impact of metal implants on the dose and clinical outcome of radiotherapy. Molecular and Clinical Oncology **21**(4) (2024)
19. Papanastasiou, G., Dikaios, N., Huang, J., Wang, C., Yang, G.: Is attention all you need in medical image analysis? A review. Journal of Biomedical and Health Informatics **28**(3) (2023)
20. Park, T., Liu, M.Y., Wang, T.C., Zhu, J.Y.: Semantic image synthesis with spatially-adaptive normalization. In: Conference on Computer Vision and Pattern Recognition (2019)
21. Rodrigues, O.: Des lois géométriques qui régissent les déplacements d'un système solide dans l'espace, et de la variation des coordonnées provenant de ces déplacements considérés indépendamment des causes qui peuvent les produire. Journal de mathématiques pures et appliquées **5**, 380–440 (1840)
22. Shi, X., Yoon, D., Koch, K.M., Hargreaves, B.A.: Metallic implant geometry and susceptibility estimation using multispectral B0 field maps. Magnetic Resonance in Medicine **77**(6) (2017)
23. Song, C., Huang, Y., Ouyang, W., Wang, L.: Mask-guided contrastive attention model for person re-identification. In: Conference on Computer Vision and Pattern Recognition (2018)
24. Spadea, M.F., Maspero, M., Zaffino, P., Seco, J.: Deep learning based synthetic-CT generation in radiotherapy and PET: A review. Medical Physics **48**(11) (2021)
25. Thummerer, A., Van der Bijl, E., Galapon Jr, A., Verhoeff, J.J., Langendijk, J.A., Both, S., van den Berg, C.N.A., Maspero, M.: SynthRAD2023 grand challenge dataset: Generating synthetic CT for radiotherapy. Medical Physics **50**(7) (2023)

26. Viceconti, M., Lattanzi, R., Antonietti, B., Paderni, S., Olmi, R., Sudanese, A., Toni, A.: CT-based surgical planning software improves the accuracy of total hip replacement preoperative planning. Medical Engineering & Physics **25**(5) (2003)
27. Villegas, F., Dal Bello, R., Alvarez-Andres, E., Dhont, J., Janssen, T., Milan, L., Robert, C., Salagean, G.A.M., Tejedor, N., Trnková, P., et al.: Challenges and opportunities in the development and clinical implementation of artificial intelligence based synthetic computed tomography for magnetic resonance only radiotherapy. Radiotherapy and Oncology **198** (2024)
28. Wang, Q.F., Tang, Y.C., Liao, H.R., Lei, M., Dong, W., Liu, Z.Y., Hao, J., Hu, Z.M.: Prevalence of metal implants among US adults aged 40 years and older. Scientific Reports **15**(1) (2025)
29. Wendling, M., Zijp, L.J., McDermott, L.N., Smit, E.J., Sonke, J.J., Mijnheer, B.J., van Herk, M.: A fast algorithm for gamma evaluation in 3D. Medical Physics **34**(5) (2007)
30. Wieser, H.P., Cisternas, E., Wahl, N., Ulrich, S., Stadler, A., Mescher, H., Müller, L.R., Klinge, T., Gabrys, H., Burigo, L., et al.: Development of the open-source dose calculation and optimization toolkit matRad. Medical Physics **44**(6) (2017)
31. Wyatt, J., McCallum, H.: Applying a commercial atlas-based synthetic computed tomography algorithm to patients with hip prostheses for prostate magnetic resonance-only radiotherapy. Radiotherapy and Oncology **133** (2019)
32. Yablonskiy, D.A.: Quantitation of intrinsic magnetic susceptibility-related effects in a tissue matrix. Phantom study. Magnetic Resonance in Medicine **39**(3) (1998)