Feasibility of Proton Dosimetry Overriding Planning CT with Daily CBCT Elaborated through Generative Artificial Intelligence Tools

Matteo Rossi [1,2,*] [ID], **Gabriele Belotti** [1] [ID], **Luca Mainardi** [1] [ID], **Guido Baroni** [1,3] [ID]
and Pietro Cerveri [1,2,*] [ID]

1   Department of Electronics, Information and Bioengineering, Politecnico di Milano, Milan, Italy
2   Istituto Auxologico Italiano, Milan, Italy
3   Bioengineering Unit, Clinical Department, National Center for Oncological Hadrontherapy (CNAO), Pavia, Italy
*   Correspondence: matteo2.rossi@polimi.it (M.R.); pietro.cerveri@polimi.it (P.C.)

Abstract: Radiotherapy commonly utilizes CBCT for patient positioning and treatment monitoring. CBCT is deemed to be secure for patients, making it suitable for the delivery of fractional doses. However, limitations such as a narrow field of view, beam hardening, scattered radiation artifacts, and variability in pixel intensity hinder the direct use of raw CBCT for dose recalculation during treatment. To address this issue, reliable correction techniques are necessary to remove artifacts and remap pixel intensity into HU values. This study proposes a deep-learning framework for calibrating CBCT images acquired with narrow FOV systems and demonstrates its potential use in proton treatment planning updates. Cycle-consistent GAN processes raw CBCT to reduce scatter and remap HU. Monte Carlo simulation is used to generate CBCT scans, enabling the possibility to focus solely on the algorithm's ability to reduce artifacts and cupping effects without considering intra-patient longitudinal variability and producing a fair comparison between planning CT and calibrated CBCT dosimetry. To showcase the viability of the approach using real-world data, experiments were also conducted using real CBCT. Tests were performed on a publicly available dataset of 40 patients who received ablative radiation therapy for pancreatic cancer. The simulated CBCT calibration led to a difference in proton dosimetry of less than 2%, compared to the planning CT. The potential toxicity effect on the organs at risk decreased from about 50% (uncalibrated) up the 2% (calibrated). The gamma pass rate at 3%/2mm produced an improvement of about 37% in replicating the prescribed dose before and after calibration (53.78% vs 90.26%). Real data also confirmed this with slightly inferior performances for the same criteria (65.36% vs 87.20%). These results may confirm that generative artificial intelligence brings the use of narrow FOV CBCT scans incrementally closer to clinical translation in proton therapy planning updates.

Keywords: deep Learning; image-to-image translation; dosimetry; cycleGAN; CBCT; CT; limited FOV; artifact correction; Hounsfield unit recovery

---

1. Introduction

Medical imaging plays a crucial role in oncology, particularly in radiotherapy. Computed tomography (CT) is the primary modality used in radiation therapy for high-resolution patient geometry and accurate dose calculations [1]. However, CT is associated with high patient exposure to ionizing radiation. Cone-beam computed tomography (CBCT) has the potential to provide faster imaging and reduce patient exposure to non-therapeutic radiation, making it a valuable imaging modality used for patient positioning and monitoring in radiotherapy. CBCT is currently used to monitor and detect changes in patient anatomy throughout the treatment; it is also compatible with fractional dose delivery, making it a patient-safe imaging modality with less additional non-therapeutic dose than traditional CT. However, this modality can introduce scattered radiation image artifacts like shading, cupping, and beam-hardening [2,3]. The artifacts resulting from scattered radiation in CBCT

images can cause fluctuations in pixel values, making it difficult to use these images directly for dose calculation. Consequently, CBCT images cannot be used directly for dose calculations unless correction methods are applied. Reliable correction techniques for calibrating CBCT images to Hounsfield Unit (HU) values used by CT scanners would expand the clinical usage of CBCT in treatment planning and evaluation of tumor shrinkage and organ shift [4–6]. In recent years, traditional approaches, such as anti-scatter grid, partial beam blockers, and scattering estimators [7–9], have been joined to deep learning-based methods, which have been showing interesting potential to improve CBCT quality [10]. Such methods, leveraging mainly convolutional neural networks (CNN) and generative adversarial networks (GAN), were investigated to map the physical model of the x-ray interaction with matter disregarding the underlying complex analytics and avoiding the use of explicit statistical approaches such as Monte Carlo. Aiming at removing scatter and correcting HU units in CBCT scans, many authors explored various types of CNN, ranging from UNet trained with a supervised training approach [11–15] to the more complex cycle-consistent Generative Adversarial Network (cGAN), based on an unsupervised training approach [16–21]. cGAN model consists of two subnetworks, the generator and the discriminator, with opposite roles. While the generator tries to learn how to convert one dataset to another, the discriminator distinguishes between real and synthetic images. This process creates a cycle-consistent loop that improves the generator's ability to produce synthetic images that look just like real ones. Focusing on CBCT-to-CT mapping, Xie et al. proposed a scatter artifact removal CNN based on a contextual loss function trained on the pelvis region of 11 subjects to correct the CBCT artifacts in the pelvic area [22]. Another research focused on a cGAN model to calibrate CBCT HU values in the pelvis region. The model was trained on 49 patients with unpaired data and tested on nine independent subjects, and the authors claimed the method kept the anatomical structure of CBCT images unchanged [18]. Exploring the use of deep residual neural networks in this field, a study demonstrated the capability of such architectures by proposing an iterative tuning-based training, where images with increasing resolutions are used at each step [23]. Likewise, our group recently reported that cGAN has better capability than CNN trained with pure supervised techniques to preserve anatomical coherence [24]. All these contributions, however, did not address the consistency of the treatment planning performed with the corrected CBCT. Conversely, Zhang et al. [25] reported the test of pelvis treatment planning in proton therapy performed on CBCT corrected with CNN. However, they summarized that the dose distribution calculated for traditional photon-based treatment outperformed the one computed for proton therapy. CBCT corrected with a cGAN was applied to evaluate the quality of the proton therapy planning in cancer treatment across different datasets with satisfactory results [13,20]. All the mentioned research works focused on the problem of CBCT-to-CT HU conversion exploiting CBCT with a wide field of view (FOV). However, some systems present in clinical practice have a limited FOV, not sufficient to contain the entire volume of the patient, e.g., in the presence of large regions such as the pelvis or with obese patients [26]. Considering the current use of CBCT for patient positioning purposes, small FOV CBCT systems could be preferred due to their reduced imaging dose, shorter computation time, and increased resolution over the treatment region of interest [27]. However, the limited FOV also causes a truncation problem during reconstruction [28,29]. Consequently, the non-uniqueness of the solution for the iterative reconstruction causes additional bright-band effects that add artifacts to the CBCT [30]. Even in the case of optimal HU calibration and scatter reduction, a CBCT, acquired in a narrow FOV cannot be used for adaptive dose planning. Especially, narrow FOV CBCT lacks important anatomical information (e.g., the air/skin interface) necessary for properly calculating the beam path. The present work aimed to propose a deep-learning framework that elaborates the CBCT to calibrate the HU, remove artifacts due to the conic geometry acquisition, and handle narrow FOV issues to demonstrate the potential use of the corrected CBCT in the context of proton treatment planning updates. The work is part of a larger study carried out in collaboration with the Italian National Center of Hadrontherapy (CNAO, Pavia, Italy) that aims to explore the possibility of using the in-house narrow FOV CBCT system not only for patient positioning but also for dosimetric evaluation without hardware modifications [31]. The deep-learning framework

took its root from the CBCT-to-CT mapping model based on cGAN proposed in [24] that was here extended to address the case of narrow FOV. Tests were carried out on a public dataset of planning CT scans of 40 oncological patients affected by pancreatic cancer. In a first attempt, synthetic raw CBCT volumes were properly generated from CT scans throughout the Monte Carlo simulation. This enabled us to dump anatomical variations usually present in real CBCT with respect to the corresponding planning CT. Moreover, in order to demonstrate the feasibility of the methodology also with real data, we replicated each experiment with the clinical CBCT included within the dataset. As the dataset provided annotation data about the segmented lesion and organs at risk, particle beam dosimetry was computed in the original planning CT and the corrected CBCT volume, verifying the coherency between the two dose distributions. The main contributions of this paper may therefore be summarized as:

- capability of the cGAN to correct CBCT (scatter reduction and HU remapping) when applied to small FOV;
- consistency of the proton dosimetry computed on corrected CBCT with respect to the original planning CT.

2. Materials and Methods

2.1. Dataset description

A publicly available dataset obtained from the Cancer Imaging Archive, called Pancreatic-CT-CBCT-SEG [32], was exploited in this work. The dataset contained CT acquisition from 40 patients who received ablative radiation therapy for locally advanced pancreatic cancer at Memorial Sloan Kettering Cancer Center. Each CT was acquired during a deep inspiration breath-hold verified with an external respiratory monitor. The dataset also included manual segmentations of a region of interest (ROI), defined by expanding the dose planning target volume by 1 cm. Along with the ROI, each scan provided contours of some organ at risk (OAR), namely: i) the stomach with the first two segments of the duodenum, ii) the remainder of the small bowel, and iii) both lungs. The authors reported that the segmentations were performed independently by six radiation oncologists and reviewed by two trained medical physicists. The dataset also provided two CBCT scans for each subject. In the first phase, these CBCTs were not considered because they were obtained at different times with respect to the corresponding CT scan, which could lead to potential changes in patient anatomy between acquisitions. Simulated CBCT scans were considered instead by generating them directly from the corresponding planning CT (implementation is detailed below in section 2.1.1). This way, perfect alignment and anatomical correspondence between the two volumes were both ensured, avoiding the need for additional registration steps (rigid or deformable). To summarize, using simulated CBCT scans allowed the study to focus solely on the algorithm's ability to reduce artifacts and cupping effects without considering intra-patient longitudinal variability. Each experiment was then also replicated with the real small FOV (250 mm diameter) CBCT provided with the dataset by adding an intermediate rigid registration step in the pipeline, in order to demonstrate the feasibility of the methodology also with data from the real world.

2.1.1. CBCT simulation

Synthetic CBCTs were generated from the original available CTs following the approach documented in [33] and replicating the setup and the geometry of CNAO's CBCT acquisition system [31]. Specifically, Monte Carlo (MC) simulations were conducted to generate primary (P_{MC}) and scatter (S_{MC}) X-ray images for each CT scan. All simulations were performed using the GATE open-source software v9.2 (based on Geant4 v11) [34] with fixed forced detection, a variance reduction technique aimed at minimizing computation time. The energy-dependent detector efficiency was based on the design specifications provided by the manufacturer for the Paxscan 4030D (Varian Medical System, Palo Alto, CA). The X-ray fluence spectrum was computed using the open-source software

SpekPy [35], employing 3.2 mm Al filtration at 100 kVp. The A-277 X-ray tube (Varian Medical System, Palo Alto, CA), chosen for this work, features a 7° rhenium-tungsten molybdenum target. Images were produced according to a CBCT scan of 220° and with projection matching the size of the Paxscan 4030D detector (isometric pixel spacing 0.388 mm, detector size 768×1024 pixels). The source-to-detector and source-to-isocenter distances were set to 1600 mm and 1100 mm, respectively. A further acceleration of scatter calculation was achieved by downsampling resolution 8-fold and simulating S_{MC} at 5° steps with a statistical uncertainty $< 5\%$. S_{MC} images were then upsampled and interpolated at the required points to match the corresponding P_{MC} images. Lastly, the final projections $S_{MC} + P_{MC}$ were normalized by the simulated flat field image. CBCT scans were then reconstructed using open-source software RTK [36] at a $1 \times 1 \times 1$ mm resolution, with a $220 \times 220 \times 220$ size in pixel and masked to the axial field-of-view of diameter equal to 204 mm. Some examples of planning CT and simulated CBCT axial slices are visible in Figure 1, along with the intensity profile of the central pixel row. The cupping effect is evident as a shaded portion in the middle of the CBCT and confirmed by the concavity in the intensity profile.
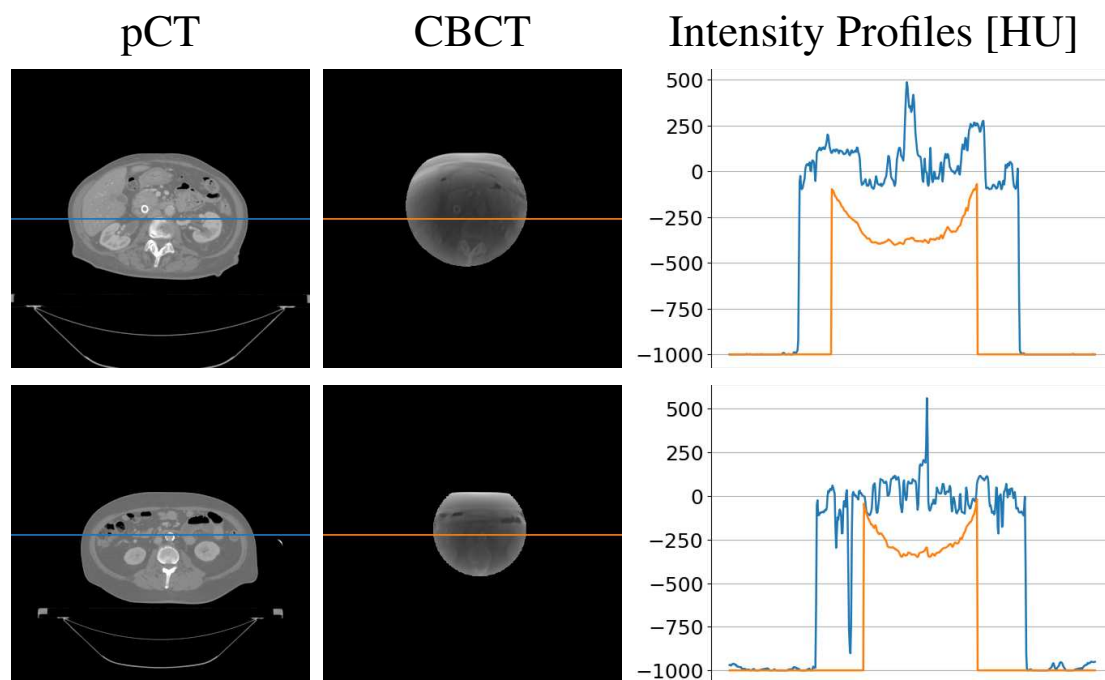
Figure 1. Examples of two CT axial slices with their corresponding simulated CBCT. The intensity profiles of the central row (marked as a line in both images) are plotted in the right panel. Each image is displayed with Window = 1300, Level = 0.

2.2. CBCT-to-CT correction

2.2.1. Neural network architecture and main processing layers

The network implemented for CBCT correction was based on the cycle Generative Adversarial Network (cGAN) [37]. This architecture is based on four concurrent subnetworks, two generators and two discriminators, which work in opposition. While the generators try to learn the mapping to convert CBCT to CT (or CT to CBCT), the discriminator's objective is to distinguish between authentic and network-generated images. This generator-discriminator cycle-consistent loop is designed to improve the generators' ability to produce synthetic images that reproduce with high fidelity the characteristic of the destination modality (e.g., generate a calibrated synthetic CT starting from a scattered CBCT). The network's fundamental processing unit was referred to as ConvBlocks (Figure 2), which were built using a 2D convolution with a 3×3 kernel, followed by an instance normalization layer and a swish activation function. The other basic processing unit for cGAN structure was the InceptionBlock,

consisting of four parallel ConvBlocks, each with an increasing kernel size of dimensions 1×1, 3×3, 7×7, and 11×11, which processed the same input simultaneously with multiple receptive fields. The output of each branch of InceptionBlock was then combined, and the complete set of feature maps was produced as output. The primary objective of this processing block was to execute multi-scale feature extraction from the initial image. The extracted multi-scale features, varying from small to large receptive fields, can produce improved outcomes for image synthesis. The general design of the generator was then carried out as a modified version of the commonly used U-Net architecture. The U-Net model is usually utilized for solving pixel-by-pixel classification challenges in image segmentation [38]. Still, it can also be used to solve image-to-image conversion problems with minor changes. The overall generator structure, depicted in Figure 2, was composed of a contracting and an expanding path. The upper two processing layers of the generator were based on InceptionBlocks, while the deeper two exploited ConvBlocks. Consequently, the network can be broadly top-bottom divided into two segments, each serving distinct functions: i) the inception part (upper layers) extracted global contextual information, whereas ii) the traditional 2D convolution part (bottom layers) was responsible for capturing more intricate context and precise localization. On the other hand, the CNN utilized as the discriminator was responsible for image classification and relied on the PatchGAN architecture [39]. Its architecture was based on four sequential ConvBlocks, each with a kernel size of 4×4. In the initial three ConvBlocks, the convolution was set with stride 2, leading to an output tensor with half the size and twice the features map. In contrast, the last ConvBlock had stride one and maintained the size and the number of feature maps unchanged. A sigmoid activation function was applied to the last layer, generating a 32×32 map used to classify the input image as real or fake.
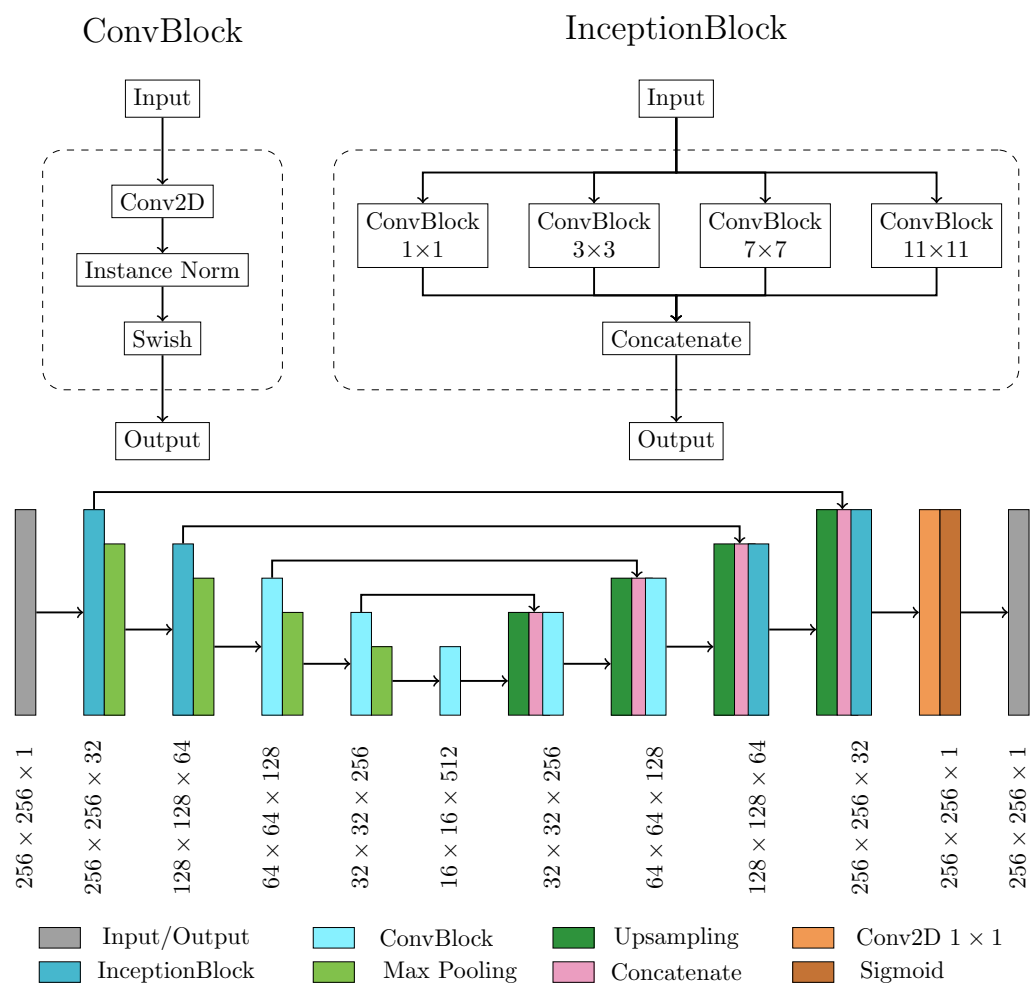
Figure 2. Schematic of the Generator model architecture.

2.2.2. Model training

As formerly stated, the cGAN overall training routine employed two generators and two discriminators, competing against one another to solve the CBCT-to-CT conversion problem. The subnetworks were referred to as generator CT (G_{CT}), generator CBCT (G_{CBCT}), discriminator CT (D_{CT}), and discriminator CBCT (D_{CBCT}). G_{CT} and g_{CBCT} were used to produce generated CT from CBCT, and generated CBCT from CT, respectively, while D_{CT} and D_{CBCT} were used to distinguish the original CT and CBCT from their generated counterparts. The training routine was subdivided into two main steps occurring simultaneously. In the first step, called the generative phase, G_{CT} (G_{CBCT}) took a 2D axial slice of a CBCT (CT) as input and produced a generated CT (generated CBCT) as output. Then, G_{CT} (G_{CBCT}) took the generated CT (generated CBCT) as input and produced a cyclic CBCT (cyclic CT), which was supposed to be equal to the original CBCT (CT). At the same time, during the second step, called the classification phase, D_{CT} (D_{CBCT}) tried to distinguish between real CT (CBCT) and generated CT (generated CBCT). The entire cGAN was implemented in Python, using Keras [40] and TensorFlow [41] frameworks. All the technical details about the network implementation can be found in a previous work of our group [24]. The whole dataset was divided into training, validation, and test sets in proportions of 70%, 15%, and 15%, containing 5698, 1221, and 1221 2D axial slices, respectively.

2.2.3. Performance metrics for model evaluation

The network performances were quantitatively evaluated using the original CT as the ground truth reference. In particular, the metrics evaluated were: i) peak signal-to-noise ratio (PSNR), ii) structural similarity index measure (SSIM), and iii) mean absolute error (MAE) [42]. The PSNR quantifies the quality of images by comparing the mean square error of the images being compared to the maximum possible signal power [43]. It is measured in decibels, and its value increases towards infinity as the difference between the calibrated CBCT and ground-truth CT decreases. Therefore, a larger PSNR value indicates better image quality, while a lower value indicates the opposite. SSIM evaluates the resemblance between two images by analyzing their luminance, contrast, and structure [44]. Compared to PSNR, SSIM is considered a more human-like measure of similarity. The SSIM score ranges from 0 to 1, with a value of 1 indicating the highest level of similarity between the images. MAE was used to quantitatively assess the accuracy of Hounsfield Units (HU) between the generated CBCT and the original CT. The lower value corresponds to the higher level of HU accuracy between the two images. The significance ($p < 0.05$) of the statistical difference between CBCT slices prior to and following calibration was verified using Kruskall-Wallis non-parametric test.

2.2.4. Synthetic CT generation pipeline

Despite the better quality of calibrated CBCT in terms of HU density values, these volumes cannot yet be used for adaptive dose planning due to their limited FOV. In fact, these CBCT acquisitions lack important anatomical information (e.g., the air/skin interface) necessary for the correct calculation of the beam path. In order to overcome this intrinsic limit, the original planning CT was used to provide the missing information. Therefore, synthetic CT (sCT) is defined in this work as an updated version of the original planning CT overridden with the calibrated voxels from daily CBCT acquired during the treatment. Starting from a scattered CBCT, the following procedure was followed in an axial slice-by-slice approach. At first, each pixel in the slice was clipped between values $[-1000; 3000]$ and then normalized in the $[0; 1]$ range. This step is fundamental because the neural network needs value in this range to operate properly. Then, the generator G_{CT} processed the normalized CBCT, producing a corrected version of the same axial slice. It is important to remember that G_{CT} is the only cGAN subnetwork used after completing the training. After neural network processing, the previous pixel clipping guarantees that the normalization can be reversed back to Hounsfield Unit. A rigid registration step between the corrected CBCT and the pCT followed, in order to increase the anatomical correspondence. This step was applied just for real CBCT, since the generated ones already matched the corresponding pCT anatomy. The last

step involved overriding the planning CT pixels with the region acquired with the cone beam modality. Every pixel outside the CBCT FOV belonged to the original CT. The entire pipeline is summarized in Figure 3. In order to evaluate the effective improvement obtained by the corrected CBCT in terms of treatment planning, two versions of sCT were generated for each subject. The first, called sCT corrected (sCTc) was obtained following the mentioned procedure, while the second, called sCT uncorrected (sCTu), was obtained simply by overriding the original CBCT volume without any kind of processing.
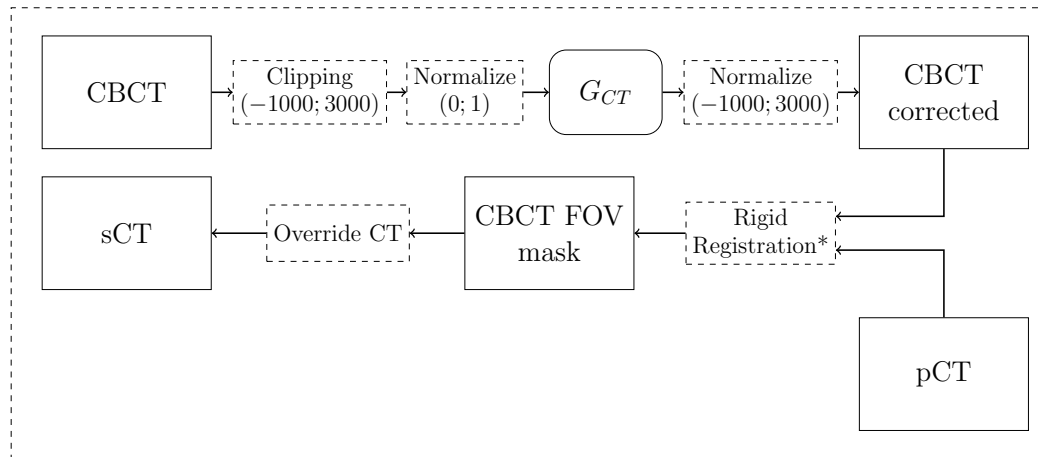
Figure 3. Schematic of sCT generation pipeline. *Rigid registration step is applied only to real CBCT scans since generated ones are intrinsically perfectly aligned.

2.3. Dosimetric analysis

2.3.1. Proton-based treatment planning

The treatment plan for each subject was computed with the matRad package [45], an open-source radiation treatment planning toolkit written in Matlab. In particular, the planning was developed using protons as the radiation mode and optimized using the constant relative biological effectiveness times dose method, which accounts for the varying biological effectiveness of different radiation types and energies. A total of 30 fractions were scheduled for the treatment, with two beams used at gantry angles of 0 degrees (anterior direction) and 270 degrees (right lateral direction). Several constraints were chosen in the planning definition to ensure the safety and efficacy of the treatment. For the bowel and stomach regions, squared overdosing and maximum dose volume histogram constraints were used to limit the radiation dose received by these sensitive areas. The lung regions were also subject to squared overdosing constraints to limit the dose delivered to that areas. Finally, the ROI was subject to squared deviation constraints, which aim to keep the dose distribution as close as possible to the prescribed dose (30 fractions of 2 Gy equivalent) [46]. The reference dose planning was first computed directly on the pCT and used as the ground truth in further comparison. Then, this reference plan was updated, giving either corrected or uncorrected sCT as the new volume.

2.3.2. Dose evaluation

To evaluate the suitability of corrected CBCT scans, various metrics were used, including dose difference pass rate (DPR), dose–volume histogram (DVH) metrics, and Gamma pass rates (GPR). The treatment dose computed for the original planning CT was considered as the prescribed ground truth dose [42]. DPR measures the percentage of pixels that meet a certain dose difference threshold, DVH compares cumulative dose to different structures in relation to volume, while GPR assesses the similarity of two dose distributions based on dose difference and distance-to-agreement criteria. The significance ($p < 0.05$) of the statistical difference in GPR distributions between sCTu and sCTc was verified using Kruskall-Wallis non-parametric test. Mean doses, D5, and D95, measured on ROI, bowel, and stomach were considered to assess the treatment quality and the toxicity control on the organ at risk.

3. Results

At a qualitative inspection, the corrected overwritten CBCT sCT scans were more similar to pCTs with respect to their uncorrected counterpart. The darker region within the CBCT FOV characterized by the cupping artifacts was no longer noticeable (Figure 4). Intensity profiles also confirmed this, showing that the concave shape observable in the unprocessed lines disappeared in the processed ones, now matching the HU values range of the reference pCT and confirming that non-linearity present in the CBCT tissue density was corrected.
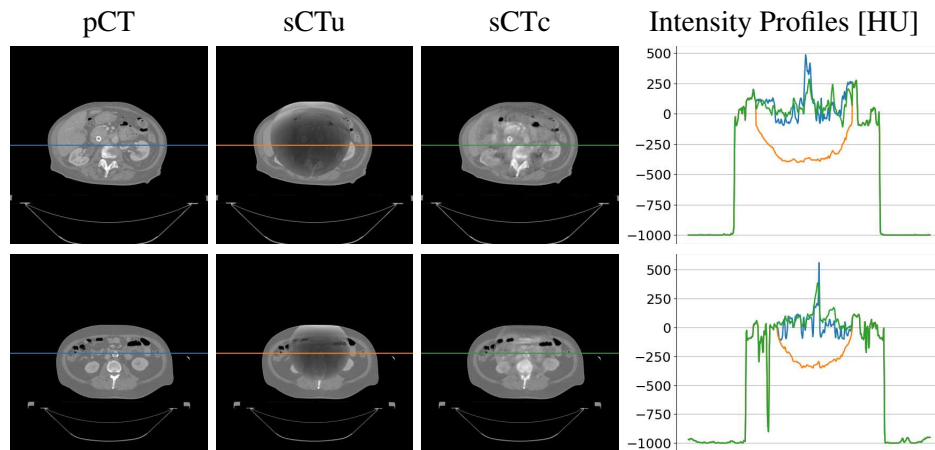
Figure 4. Examples of two CT axial slices with their corresponding sCT generated overriding the uncorrected simulated CBCT (sCTu) and the corrected ones (sCTc). The intensity profiles of the central row (marked as a line in both images) are plotted in the right panel. Each image is displayed with Window = 1300, Level = 0.

3.1. cGAN model evaluation

The results from evaluating the model's performance metrics (Figure 5) showed promising improvements in the quality of the CBCT slices. The original images had a median PSNR of 24.60 dB (IQR 1.40 dB), while the processed images had a median PSNR of 33.41 dB (IQR 3.36 dB), resulting in a relative gain of approximately 37%. In terms of the SSIM score, the original images had a median of 0.90 (IQR 0.03), and the processed CBCTs showed a median of 0.95 (IQR 0.02), which corresponded to a relative enhancement of around 5%. Furthermore, the median MAE for the original images was 148.96 HU (IQR 31.24 HU), whereas the median MAE for the processed images was 43.47 HU (IQR 14.82 HU). These results demonstrate the effectiveness of the cGAN approach in improving CBCT image quality. For all three metrics, a statistical difference was found ($p < 0.01$).

Figure 5. Performance metrics for cGAN model evaluation, computed on the axial slices of the test set before and after model processing.

3.2. Treatment planning evaluation - simulated data

The qualitative comparison between the treatment plans computed for each modality confirmed an ameliorated similarity between sCTc and pCT with respect to their uncorrected counterpart. An example of this can be seen in Figure 6 upper row. It is visible how the beam path computed on sCTu exceeded the ROI releasing a more significant amount of dose in the following tissues. Moreover, it could also be seen that high dose values (red pixels) break over ROI boundaries, indicating that a portion of surrounding healthy areas received overexposure to radiation, as also confirmed by the dose difference computed with respect to the pCT reference plan (cfr. Figure 6 bottom row). Conversely, the sCTc treatment plan corrected that pattern, showing a more similar beam path and a reduced difference with respect to the pCT plan. Likewise, the dose-volume histogram computed for the same test subject confirmed and enforced such a consideration (Figure 7). Observing the ROI lines, the sCTc (orange dotted line) followed the profile of pCT (orange solid line) more closely than the sCTu (orange dashed line). The treatment plan calculated on sCTc and pCT showed a steep slope, indicating that 2 Gy was the dose delivered to almost all the ROI, while sCTu showed a smoother slope, a sign of overdosing in a portion of this area. Concerning the organs at risk (green and light blue lines), this consideration was even more evident, with an overdosing in the order of about two times with respect to the reference plan that is incompatible with clinical practice.
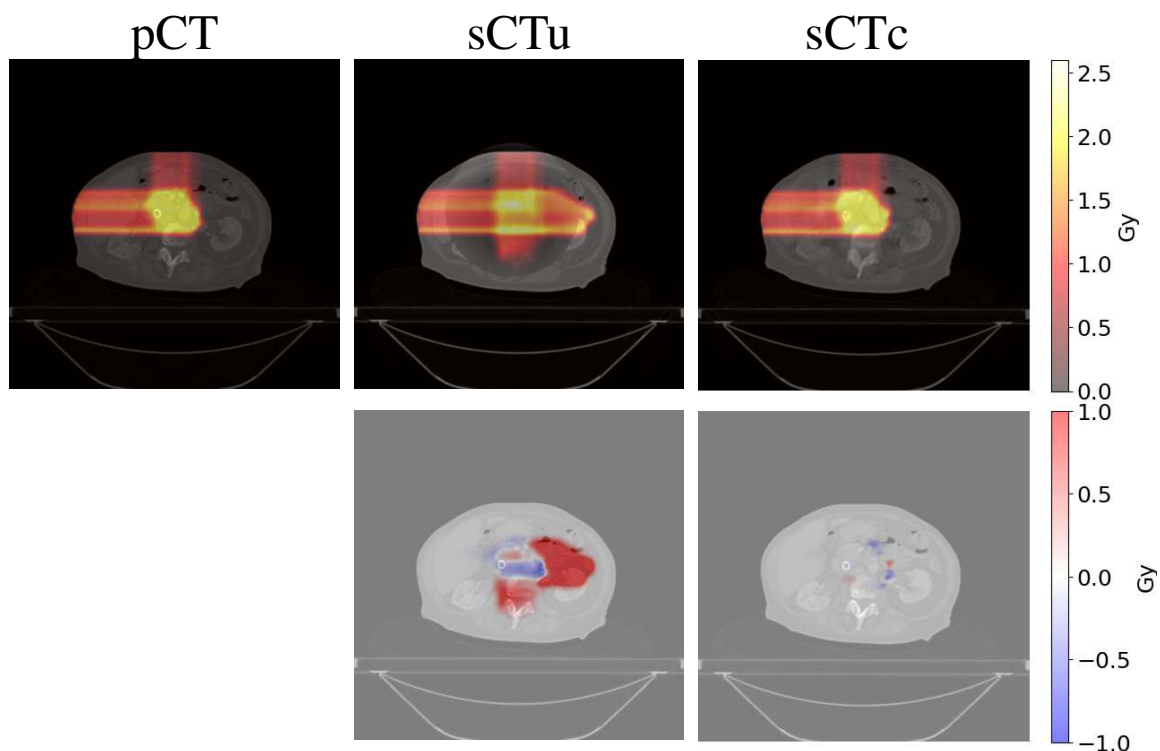
Figure 6. Example of dose planning for an axial slice of a subject for the original pCT, sCTu, and sCTc (upper row). The difference between both sCT treatment plans with respect to the original pCT plan is shown in the bottom row. Synthetic CT scans in this figure were produced using a simulated CBCT as the input.

Figure 7. Example of DVH computed for ROI and two organs at risk (bowel and stomach). Solid line: pCT, dashed line: sCTu, dotted line: sCTc. Synthetic CT scans in this figure were produced using a simulated CBCT as the input.

The GPR results for the entire dataset were computed for different gamma criteria and summarized in Table 1 as median (interquartile range, IQR). The median gamma pass rates for the 1%/1 mm, 2%/2 mm, 3%/2 mm, and 3%/3 mm criteria were consistently higher for the sCTc the sCTu, with the most significant improvement observed for the 3%/3 mm criterion (92.82% vs. 57.57%).

Table 1. Gamma pass rate results for different criteria computed using only the simulated CBCT dataset as input to the framework. Statistical difference in gamma score distributions, between uncorrected and corrected sCT, was found ($p < 0.001$).

Gamma Criterion	sCTu	sCTc
1%/1 mm	44.68 (6.91)	74.37 (4.63)
2%/2 mm	51.72 (8.15)	87.30 (6.32)
3%/2 mm	53.78 (8.53)	90.26 (5.70)
3%/3 mm	57.57 (7.49)	92.82 (5.94)

DPR at 1% was also found to be significantly higher for the processed sCT than the unprocessed ones (93.97% vs. 79.76%), indicating a 14.21% improvement in dose accuracy with the use of sCTc. In regard to mean dose distribution across the overall dataset, the advantage of sCTc was evident with respect to the overdosing of the sCTu in the ROI (Table 2). The relative percentage error decreased from 6% for sCTu up to 2% for sCTc. A greater advantage was achieved in terms of unwanted doses distributed at the bowel as 97% against 2%, with respect to the nominal toxicity in the pCT. Likewise, the relative toxicity in the stomach decreased from 49% up the 2%. For D5, the correction was effective in reducing the overexposure found in the unprocessed sCT. For D95, the correction underestimated the dose of about 5%. As far as OAR is concerned, the correction was again effective in ensuring a low

dose, very similar to that one obtained in the planning CT, at both bowel and stomach. Interestingly, the IQR range of 0.12 Gy (D95) for the stomach potentially delivered when using the sCTu was completely zeroed by the correction.

Table 2. Mean doses, D5, and D95, measured on ROI, bowel, and stomach, computed using only the simulated CBCT dataset as input to the framework. Values are expressed as median(IQR) Gy.

		Mean dose	D5	D95
ROI	**pCT**	1.98 (0.01)	2.04 (0.01)	1.86 (0.07)
	sCTu	2.10 (0.09)	2.61 (0.28)	1.64 (0.10)
	sCTc	1.93 (0.08)	2.09 (0.05)	1.54 (0.28)
Bowel	**pCT**	0.47 (0.44)	2.01 (0.04)	0.00 (0.00)
	sCTu	0.93 (0.52)	2.06 (0.31)	0.00 (0.00)
	sCTc	0.48 (0.40)	2.00 (0.13)	0.00 (0.00)
Stomach	**pCT**	0.65 (0.27)	2.01 (0.03)	0.00 (0.00)
	sCTu	0.97 (0.35)	2.14 (0.24)	0.00 (0.12)
	sCTc	0.64 (0.29)	2.02 (0.11)	0.00 (0.00)

3.3. Treatment planning evaluation - real data

Each result presented in previous sections referred to simulated CBCT. The following results refer to sCT generated using real CBCT as the input data in order to show the quality of the treatment planning with real-world data. As explained in Section 2.2.4, a rigid registration step was added to the pipeline just before the pCT override. Once again, the comparison of the treatment plans generated for each method confirmed an improved similarity between sCTc and pCT with respect to their uncorrected equivalent. Figure 8) depicts an example case in which it is possible to observe again how the beam path exceeded the ROI with a consequent overdosing in the adjacent healthy tissues when computed on sCTu. However, the sCTc treatment plan effectively reduced the overdosing and corrected this pattern, leading to a beam path that closely resembled the pCT plan and reducing the differences (cfr. Figure 8 bottom row). The corresponding dose-volume histogram for the test subject (Figure 9) further supported this observation. Looking at the ROI lines, the sCTc (orange dotted line) followed the profile of pCT (orange solid line) more closely than the sCTu (orange dashed line). Even if the corrected plan slightly overdosed about 60% of the volume, the effect is reduced compared to the uncorrected plan, which presented a smoother slope for the entirety of the ROI, giving 1.25 Gy to the 100% of the volume instead of the 2 Gy of the prescribed dose and overdosing the rest. This observation became even more apparent when considering the organs at risk (represented by the green and light blue lines). Specifically, the corrected treatment plan for the bowel closely replicated the prescribed one, demonstrating a good level of accuracy.

For a quantitative comparison, the GPR results computed on the entire dataset are summarized in Table 3 as median (IQR). Again, even in the presence of real CBCT overridden to the pCT, the median gamma pass rates for all the computed criteria were consistently higher for the sCTc, with the most significant improvement observed for the 3%/2 mm criterion (23% difference).

Regarding the average dose distribution, the superiority of sCTc over sCTu in terms of overdosing within the ROI was evident, as indicated in Table 4. The relative percentage error decreased from 3% for sCTu to 1% for sCTc. Furthermore, a significant advantage was also achieved in terms of undesired doses in the bowel, with error percentages of 31% (sCTu) and 13% (sCTc) in comparison to the nominal toxicity in the pCT. Similarly, the relative toxicity in the stomach decreased from 12% to 3%.

Figure 8. Example of dose planning for an axial slice of a subject for the original pCT, sCTu, and sCTc (upper row). The difference between both sCT treatment plans with respect to the original pCT plan is shown in the bottom row. Synthetic CT scans in this figure were produced using a real CBCT as the input.

Figure 9. Example of DVH computed for ROI and two organs at risk (bowel and stomach). Solid line: pCT, dashed line: sCTu, dotted line: sCTc. Synthetic CT scans in this figure were produced using a real CBCT as the input.

Table 3. Gamma pass rate results for different criteria computed using real CBCT dataset as input to the framework. Statistical difference in gamma score distributions, between uncorrected and corrected sCT, was found ($p < 0.001$).

Gamma Criterion	sCTu	sCTc
1%/1 mm	49.73 (14.84)	71.97 (7.01)
2%/2 mm	61.41 (14.17)	84.37 (5.89)
3%/2 mm	65.36 (14.17)	87.20 (5.79)
3%/3 mm	70.11 (13.66)	89.87 (5.26)

Table 4. Mean doses, D5, and D95, measured on ROI, bowel, and stomach, computed using only the real CBCT dataset as input to the framework. Values are expressed as median(IQR) Gy.

		Mean dose	D5	D95
ROI	pCT	1.98 (0.01)	2.04 (0.01)	1.86 (0.07)
	sCTu	2.03 (0.08)	2.37 (0.21)	1.73 (0.13)
	sCTc	1.96 (0.06)	2.11 (0.07)	1.69 (0.18)
Bowel	pCT	0.47 (0.44)	2.01 (0.04)	0.00 (0.00)
	sCTu	0.62 (0.43)	2.05 (0.22)	0.00 (0.00)
	sCTc	0.41 (0.29)	2.00 (0.12)	0.00 (0.00)
Stomach	pCT	0.65 (0.27)	2.01 (0.03)	0.00 (0.00)
	sCTu	0.73 (0.31)	2.12 (0.15)	0.00 (0.00)
	sCTc	0.63 (0.26)	2.02 (0.06)	0.00 (0.00)

4. Discussion

This work proposed a novel image-processing framework for generating synthetic CT scans, which combines the original planning CT with routine CBCT scans, usable to update the dosimetry plan in proton therapy. The core of the framework was represented by a deep learning model, namely a cycle-consistency GAN, to correct the scatter artifacts in the CBCT images and calibrate intensity values, in the proper HU range. Especially, the framework was shown to properly address CBCT equipment scanning narrow FOV [42]. To the aim, the public dataset of patient-paired CT-CBCT scans, named Pancreatic-CT-CBCT-SEG [32], was exploited because of clinically consistent segmentation of ROI and OAR across all the considered patients. From the CT scans, physically consistent simulated CBCT were generated by means of the Monte Carlo algorithm, also accounting for narrow FOV. The cGAN-based correction mapped synthetic uncorrected CT into synthetic corrected CT. The available segmentations granted replicating the particle beam planning, computed on the original CT, to the corrected sCT for straightforward comparison. The obtained results confirmed, both qualitatively and quantitatively, the capability of the cGAN-based to correct CBCT into CT-compatible images. As shown, the intensity profiles were rectified adequately thanks to the reduction of cupping and truncation artifacts (cfr. Figure 4). A significant increase in PSNR, SSIM, and MAE metrics testified to the effectiveness of the methodology (cfr. Figure 5). Dosimetry computed on uncorrected sCT featured overdosing, especially at OAR (cfr. Tables 1 and 2). An inaccurate assessment of tissue densities was made due to the lack of HU calibration. The difference in grayscale values between CBCT and CT caused a discontinuity in the volume (cfr. Figure 4), leading to errors in particle beam path computation. Conversely, the dosimetric plan computed on the corrected sCT confirmed the consistency of the plan computed on the corresponding CT scan (cfr. Table 1). The American Association of Physicists in Medicine (AAPM) Task Group 218 defined acceptance criteria for tolerance and action limits as exceeding 95% and falling below 90%, respectively, for a 3%/2 mm GPR standard [47]. While not completely in agreement with the upper threshold, the 90.26% found in this work (cfr. Table 1) is to be deemed reasonable. Nonetheless, such value, overcoming the lower 90% action limit threshold, makes the methodology promising for clinical application. In order to avoid confounding factors induced by organ deformation, this work did use simulated CBCT by means of the Monte Carlo approach,

so that differences between the images were exclusively due to artifacts rather than anatomy. This allowed us to use the same lesion and OAR segmentation used to compute the reference treatment plan. However, results obtained using real-world CBCT as input to the framework confirmed the feasibility of the approach (cfr. Table 3). Real data provided in general lower performances when compared to the simulated ones. This was mainly due to the anatomical changes when in the presence of volume acquired on different days during the treatment. It is important to remind that the pelvic site, the subject of this study, contains soft tissue that implies relative movement between organs and the generation of air bubbles [15,48]. Moreover, during the treatment, the patient often loses weight, and the tumor changes shape due to the treatment itself. This makes this framework not always applicable, but a possibly useful aid when changes in internal anatomy are not so evident. For every other case, the traditional acquisition of a revaluation CT remains a better solution. Comparison with works in the literature using deep learning to correct CBCT and test proton dosimetry showcased the consistency of the methodology in terms of GPR 2%/2 mm results, even though slightly smaller values with respect to that reported in some papers (Table 5). Significantly, our results on simulated and real data differed less than expected. Nonetheless, we note that our study applied the correction to small FOV (204 mm diameter for simulated data and 250 mm for real data) while the mentioned works dealt with mainly wider FOV (about 480 mm diameter on average). All the works in this comparison involved large anatomical sites (e.g., pelvis, thorax, and abdomen) and the patient cohort had a similar size. Remarkably, the comparison highlighted the superiority of the generative adversarial networks [20,49] with respect to traditional UNet [13,50].

Table 5. Comparison of dosimetry results with literature outcomes in terms of GPR2%/2 mm in the domain of proton therapy.

Work	Model	Anatomic Site	axial FOV [mm]	Patient cohort	GPR 2%/2 mm
Hansen et al. [50]	Unet	Pelvis	410	30	53%
Landry et al. [13]	Unet	Pelvis	410	42	85%
Thummerer et al. [51]	UNet	Thorax	500	33	90.7%
Kurz et al. [49]	cGAN	Pelvis	550	33	96%
Uh et al. [20]	cGAN	Abdomen/Pelvis	530	50	98.5%
This work - simulated	cGAN	Pelvis	204	40	87.3%
This work - real	cGAN	Pelvis	250	40	84.4%

Concerning the small FOV, it is fundamental to recall that the volume of interest is entirely contained in the CBCT and that the tissues coming from the planning CT are only needed to calculate the beam path when the air/tissue interface is not already embedded in the CBCT [28,29]. In addition, this makes it mandatory to update the segmentation mask. In this work, the FOV was particularly reduced to demonstrate the feasibility of the methodology. In general, the proposed framework can be adopted in all cases in which the district of interest is too large to fit into a single acquisition. Furthermore, the proposed method extends the use of CBCT systems currently used for patient positioning without additional hardware. Considering that daily CBCTs are acquired for patient positioning purposes, the present framework can be used in parallel with the current clinical routine. Moreover, the computation time required to produce an sCT starting from a CBCT is in the order of a few minutes on an average computer, making it compatible with clinical routines. This method will be retrospectively evaluated in a real offline therapy context at the CNAO facility in future work [31]. Correction of CBCTs obtained day-by-day during treatment will be used to assess the evolution of the dose plan without the need to acquire additional CTs and administer further toxicity to the patient [6,20]. Moreover, no additional hardware will be needed to add to the patient positioning setup in order to increase the FOV for dosimetric evaluation.

5. Conclusions

The present study proposed a generative artificial intelligence tool to correct CBCT scans, acquired with narrow FOV systems, enabling the reduction of scatter and the remap pixel intensity in HU. The methodology made feasible treatment planning updates, which brings the use of CBCT images incrementally closer to clinical translation in proton therapy.

Author Contributions: Conceptualization: P.C.; data curation: M.R. and G.Be.; formal analysis: M.R.; funding acquisition: L.M. and P.C.; methodology: M.R. and P.C.; project administration: G. Ba. and P.C.; resources: G. Ba. and P.C.; software: M.R. and G.Be.; supervision: P.C.; validation: M.R.; visualization: M.R.; writing—original draft: M.R.; writing—review and editing: M.R. and P.C. All authors have read and agreed to the published version of the manuscript.

Funding: This work was supported by CNAO Foundation (Pavia, Italy) and by FAIR (Future of Artificial Intelligence Research) funded in the PNRR-PE program by Italian Ministry of University.

Institutional Review Board Statement: Not applicable.

Informed Consent Statement: Not applicable.

Data Availability Statement: The data presented in this study are openly available in The Cancer Imaging Archive (TCIA) at 10.7937/TCIA.ESHQ-4D90. These collections are freely available to browse, download, and use for commercial, scientific and educational purposes as outlined in the Creative Commons Attribution 3.0 Unported License.

Acknowledgments: The authors would like to thank Fabio Casaccio for his support in data retrieval.

Conflicts of Interest: The authors declare no conflict of interest.

Abbreviations

The following abbreviations are used in this manuscript:

CBCT	Cone-Beam Computed Tomography
cGAN	cycle-consistent Generative Adversarial Network
CNN	Convolutional Neural Network
CT	Computed Tomography
D_{CBCT}	Discriminator CBCT
D_{CT}	Discriminator CT
DPR	Dose Difference Pass Rate
DVH	Dose–Volume Histogram
FOV	Field of View
G_{CBCT}	Generator CBCT
G_{CT}	Generator CT
GPR	Gamma Pass Rate
IQR	Interquartile Range
MAE	Mean Absolute Error
MC	Monte Carlo
OAR	Organ at Risk
pCT	planning CT
PSNR	Peak Signal-to-Noise Ratio
ROI	Region of Interest
sCT	synthetic CT
sCTc	corrected sCT
sCTu	uncorrected sCT
SSIM	Structural Similarity Index Measure

References

1. Bortfeld, T. IMRT: a review and preview. *Physics in Medicine and Biology* **2006**, *51*, R363–R379. https://doi.org/10.1088/0031-9155/51/13/r21.

2.   Joseph, P.M.; Spital, R.D. The effects of scatter in x-ray computed tomography. *Medical Physics* **1982**, *9*, 464–472. https://doi.org/10.1118/1.595111.

3.   Schulze, R.; Heil, U.; Groβ, D.; Bruellmann, D.D.; Dranischnikow, E.; Schwanecke, U.; Schoemer, E. Artefacts in CBCT: a review. *Dentomaxillofacial Radiology* **2011**, *40*, 265–273. https://doi.org/10.1259/dmfr/30642039.

4.   Kurz, C.; Kamp, F.; Park, Y.K.; Zöllner, C.; Rit, S.; Hansen, D.; Podesta, M.; Sharp, G.C.; Li, M.; Reiner, M.; et al. Investigating deformable image registration and scatter correction for CBCT-based dose calculation in adaptive IMPT. *Medical Physics* **2016**, *43*, 5635–5646. https://doi.org/10.1118/1.4962933.

5.   Thing, R.S.; Bernchou, U.; Mainegra-Hing, E.; Hansen, O.; Brink, C. Hounsfield unit recovery in clinical cone beam CT images of the thorax acquired for image guided radiation therapy. *Physics in Medicine and Biology* **2016**, *61*, 5781–5802. https://doi.org/10.1088/0031-9155/61/15/5781.

6.   Giacometti, V.; Hounsell, A.R.; McGarry, C.K. A review of dose calculation approaches with cone beam CT in photon and proton therapy. *Physica Medica* **2020**, *76*, 243–276. https://doi.org/10.1016/j.ejmp.2020.06.017.

7.   Sun, M.; Star-Lack, J.M. Improved scatter correction using adaptive scatter kernel superposition. *Physics in Medicine and Biology* **2010**, *55*, 6695–6720. https://doi.org/10.1088/0031-9155/55/22/007.

8.   Sisniega, A.; Zbijewski, W.; Badal, A.; Kyprianou, I.S.; Stayman, J.W.; Vaquero, J.J.; Siewerdsen, J.H. Monte Carlo study of the effects of system geometry and antiscatter grids on cone-beam CT scatter distributions. *Medical Physics* **2013**, *40*, 051915. https://doi.org/10.1118/1.4801895.

9.   Stankovic, U.; Ploeger, L.S.; van Herk, M.; Sonke, J.J. Optimal combination of anti-scatter grids and software correction for CBCT imaging. *Medical Physics* **2017**, *44*, 4437–4451. https://doi.org/10.1002/mp.12385.

10.   Rusanov, B.; Hassan, G.M.; Reynolds, M.; Sabet, M.; Kendrick, J.; Rowshanfarzad, P.; Ebert, M. Deep learning methods for enhancing cone-beam CT image quality toward adaptive radiation therapy: A systematic review. *Medical Physics* **2022**, *49*, 6019–6054. https://doi.org/10.1002/mp.15840.

11.   Kida, S.; Nakamoto, T.; Nakano, M.; Nawa, K.; Haga, A.; Kotoku, J.; Yamashita, H.; Nakagawa, K. Cone Beam Computed Tomography image quality improvement using a deep convolutional neural network. *Cureus* **2018**. https://doi.org/10.7759/cureus.2548.

12.   Jiang, Y.; Yang, C.; Yang, P.; Hu, X.; Luo, C.; Xue, Y.; Xu, L.; Hu, X.; Zhang, L.; Wang, J.; et al. Scatter correction of cone-beam CT using a deep residual convolution neural network (DRCNN). *Physics in Medicine & Biology* **2019**, *64*, 145003. https://doi.org/10.1088/1361-6560/ab23a6.

13.   Landry, G.; Hansen, D.; Kamp, F.; Li, M.; Hoyle, B.; Weller, J.; Parodi, K.; Belka, C.; Kurz, C. Comparing Unet training with three different datasets to correct CBCT images for prostate radiotherapy dose calculations. *Physics in Medicine & Biology* **2019**, *64*, 035011. https://doi.org/10.1088/1361-6560/aaf496.

14.   Chen, L.; Liang, X.; Shen, C.; Jiang, S.; Wang, J. Synthetic CT generation from CBCT images via deep learning. *Medical Physics* **2020**, *47*, 1115–1125. https://doi.org/10.1002/mp.13978.

15.   Rossi, M.; Belotti, G.; Paganelli, C.; Pella, A.; Barcellini, A.; Cerveri, P.; Baroni, G. Image-based shading correction for narrow-FOV truncated pelvic CBCT with deep convolutional neural networks and transfer learning. *Medical physics* **2021**, *48*, 7112–7126. https://doi.org/10.1002/mp.15282.

16.   Kida, S.; Kaji, S.; Nawa, K.; Imae, T.; Nakamoto, T.; Ozaki, S.; Ohta, T.; Nozawa, Y.; Nakagawa, K. Visual enhancement of Cone-beam CT by use of CycleGAN. *Medical Physics* **2020**, *47*, 998–1010. https://doi.org/10.1002/mp.13963.

17.   Eckl, M.; Hoppen, L.; Sarria, G.R.; Boda-Heggemann, J.; Simeonova-Chergou, A.; Steil, V.; Giordano, F.A.; Fleckenstein, J. Evaluation of a cycle-generative adversarial network-based cone-beam CT to synthetic CT conversion algorithm for adaptive radiation therapy. *Physica Medica* **2020**, *80*, 308–316. https://doi.org/10.1016/j.ejmp.2020.11.007.

18.   Dong, G.; Zhang, C.; Liang, X.; Deng, L.; Zhu, Y.; Zhu, X.; Zhou, X.; Song, L.; Zhao, X.; Xie, Y. A Deep Unsupervised Learning Model for Artifact Correction of Pelvis Cone-Beam CT. *Frontiers in Oncology* **2021**, *11*. https://doi.org/10.3389/fonc.2021.686875.

19.   Sun, H.; Fan, R.; Li, C.; Lu, Z.; Xie, K.; Ni, X.; Yang, J. Imaging Study of Pseudo-CT Synthesized From Cone-Beam CT Based on 3D CycleGAN in Radiotherapy. *Frontiers in Oncology* **2021**, *11*. https://doi.org/10.3389/fonc.2021.603844.

20.   Uh, J.; Wang, C.; Acharya, S.; Krasin, M.J.; ho Hua, C. Training a deep neural network coping with diversities in abdominal and pelvic images of children and young adults for CBCT-based adaptive proton therapy. *Radiotherapy and Oncology* **2021**, *160*, 250–258. https://doi.org/10.1016/j.radonc.2021.05.006.

21. Zhao, J.; Chen, Z.; Wang, J.; Xia, F.; Peng, J.; Hu, Y.; Hu, W.; Zhang, Z. MV CBCT-Based Synthetic CT Generation Using a Deep Learning Method for Rectal Cancer Adaptive Radiotherapy. *Frontiers in Oncology* **2021**, *11*. https://doi.org/10.3389/fonc.2021.655325.

22. Xie, S.; Liang, Y.; Yang, T.; Song, Z. Contextual loss based artifact removal method on CBCT image. *Journal of Applied Clinical Medical Physics* **2020**, *21*, 166–177. https://doi.org/10.1002/acm2.13084.

23. Wu, W.; Qu, J.; Cai, J.; Yang, R. Multiresolution residual deep neural network for improving pelvic CBCT image quality. *Medical Physics* **2022**, *49*, 1522–1534. https://doi.org/10.1002/mp.15460.

24. Rossi, M.; Cerveri, P. Comparison of Supervised and Unsupervised Approaches for the Generation of Synthetic CT from Cone-Beam CT. *Diagnostics (Basel, Switzerland)* **2021**, *11*. https://doi.org/10.3390/diagnostics11081435.

25. Zhang, Y.; Yue, N.; Su, M.Y.; Liu, B.; Ding, Y.; Zhou, Y.; Wang, H.; Kuang, Y.; Nie, K. Improving CBCT quality to CT level using deep learning with generative adversarial network. *Medical Physics* **2021**, *48*, 2816–2826. https://doi.org/10.1002/mp.14624.

26. Landry, G.; Hua, C.h. Current state and future applications of radiological image guidance for particle therapy. *Medical Physics* **2018**, *45*, e1086–e1095. https://doi.org/10.1002/mp.12744.

27. Lu, W.; Yan, H.; Zhou, L.; Cervino, L.; Jiang, S.; Jia, X. TU-G-141-05: Limited Field-Of-View Cone-Beam CT Reconstruction for Adaptive Radiotherapy. *Medical Physics* **2013**, *40*, 457–457. https://doi.org/10.1118/1.4815465.

28. Yu, V.Y.; Keyrilainen, J.; Suilamo, S.; Beslimane, I.; Dresner, A.; Halkola, A.; Van der Heide, U.A.; Tyagi, N. A multi-institutional analysis of a general pelvis continuous Hounsfield unit synthetic CT software for radiotherapy. *Journal of applied clinical medical physics* **2021**, *22*, 207–215.

29. Velten, C.; Goddard, L.; Jeong, K.; Garg, M.K.; Tomé, W.A. Clinical Assessment of a Novel Ring Gantry Linear Accelerator-Mounted Helical Fan-Beam kVCT System. *Advances in radiation oncology* **2022**, *7*, 100862.

30. Clackdoyle, R.; Defrise, M. Tomographic Reconstruction in the 21st Century. *IEEE Signal Processing Magazine* **2010**, *27*, 60–80. https://doi.org/10.1109/msp.2010.936743.

31. Fattori, G.; Riboldi, M.; Pella, A.; Peroni, M.; Cerveri, P.; Desplanques, M.; Fontana, G.; Tagaste, B.; Valvo, F.; Orecchia, R.; et al. Image guided particle therapy in CNAO room 2: Implementation and clinical validation. *Physica Medica* **2015**, *31*, 9–15. https://doi.org/10.1016/j.ejmp.2014.10.075.

32. Hong, J.; Reyngold, M.; Crane, C.; Cuaron, J.; Hajj, C.; Mann, J.; Zinovoy, M.; Yorke, E.; LoCastro, E.; Apte, A.P.; et al. Breath-hold CT and cone-beam CT images with expert manual organ-at-risk segmentations from radiation treatments of locally advanced pancreatic cancer (Pancreatic-CT-CBCT-SEG), 2021. https://doi.org/10.7937/TCIA.ESHQ-4D90.

33. Poludniowski, G.; Evans, P.M.; Hansen, V.N.; Webb, S. An efficient Monte Carlo-based algorithm for scatter correction in keV cone-beam CT. *Physics in Medicine and Biology* **2009**, *54*, 3847–3864. https://doi.org/10.1088/0031-9155/54/12/016.

34. Jan, S.; Santin, G.; Strul, D.; Staelens, S.; Assié, K.; Autret, D.; Avner, S.; Barbier, R.; Bardiès, M.; Bloomfield, P.M.; et al. GATE: a simulation toolkit for PET and SPECT. *Physics in Medicine and Biology* **2004**, *49*, 4543–4561. https://doi.org/10.1088/0031-9155/49/19/007.

35. Poludniowski, G.; Omar, A.; Bujila, R.; Andreo, P. Technical Note: SpekPy v2.0—a software toolkit for modeling x-ray tube spectra. *Medical Physics* **2021**, *48*, 3630–3637. https://doi.org/10.1002/mp.14945.

36. Rit, S.; Oliva, M.V.; Brousmiche, S.; Labarbe, R.; Sarrut, D.; Sharp, G.C. The Reconstruction Toolkit (RTK), an open-source cone-beam CT reconstruction toolkit based on the Insight Toolkit (ITK). *Journal of Physics: Conference Series* **2014**, *489*, 012079. https://doi.org/10.1088/1742-6596/489/1/012079.

37. Zhu, J.Y.; Park, T.; Isola, P.; Efros, A.A. Unpaired Image-to-Image Translation Using Cycle-Consistent Adversarial Networks. In Proceedings of the 2017 IEEE International Conference on Computer Vision (ICCV). IEEE, oct 2017. https://doi.org/10.1109/iccv.2017.244.

38. Ronneberger, O.; Fischer, P.; Brox, T. U-Net: Convolutional Networks for Biomedical Image Segmentation, 2015. https://doi.org/10.48550/ARXIV.1505.04597.

39. Isola, P.; Zhu, J.Y.; Zhou, T.; Efros, A.A. Image-to-Image Translation with Conditional Adversarial Networks. In Proceedings of the 2017 IEEE Conference on Computer Vision and Pattern Recognition (CVPR). IEEE, jul 2017. https://doi.org/10.1109/cvpr.2017.632.

40. Chollet, F.; et al. Keras. https://keras.io, 2015.

41. Abadi, M.; Agarwal, A.; Barham, P.; Brevdo, E.; Chen, Z.; Citro, C.; Corrado, G.S.; Davis, A.; Dean, J.; Devin, M.; et al. TensorFlow: Large-Scale Machine Learning on Heterogeneous Systems, 2015. Software available from tensorflow.org.

42. Spadea, M.F.; Maspero, M.; Zaffino, P.; Seco, J. Deep learning based synthetic-CT generation in radiotherapy and PET: A review. *Medical Physics* **2021**, *48*, 6537–6566. https://doi.org/10.1002/mp.15150.

43. Horé, A.; Ziou, D. Image Quality Metrics: PSNR vs. SSIM. In Proceedings of the 2010 20th International Conference on Pattern Recognition, 2010, pp. 2366–2369. https://doi.org/10.1109/ICPR.2010.579.

44. Wang, Z.; Bovik, A.C.; Sheikh, H.R.; Simoncelli, E.P. Image quality assessment: from error visibility to structural similarity. *IEEE Transactions on Image Processing* **2004**, *13*, 600–612. https://doi.org/10.1109/tip.2003.819861.

45. 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* **2017**, *44*, 2556–2568. https://doi.org/10.1002/mp.12251.

46. Dreher, C.; Habermehl, D.; Ecker, S.; Brons, S.; El-Shafie, R.; Jäkel, O.; Debus, J.; Combs, S.E. Optimization of carbon ion and proton treatment plans using the raster-scanning technique for patients with unresectable pancreatic cancer. *Radiation Oncology* **2015**, *10*. https://doi.org/10.1186/s13014-015-0538-x.

47. Miften, M.; Olch, A.; Mihailidis, D.; Moran, J.; Pawlicki, T.; Molineu, A.; Li, H.; Wijesooriya, K.; Shi, J.; Xia, P.; et al. Tolerance limits and methodologies for IMRT measurement-based verification QA: Recommendations of AAPM Task Group No. 218. *Medical Physics* **2018**, *45*. https://doi.org/10.1002/mp.12810.

48. Niu, T.; Al-Basheer, A.; Zhu, L. Quantitative cone-beam CT imaging in radiation therapy using planning CT as a prior: First patient studies. *Medical Physics* **2012**, *39*, 1991–2000. https://doi.org/10.1118/1.3693050.

49. Kurz, C.; Maspero, M.; Savenije, M.H.F.; Landry, G.; Kamp, F.; Pinto, M.; Li, M.; Parodi, K.; Belka, C.; van den Berg, C.A.T. CBCT correction using a cycle-consistent generative adversarial network and unpaired training to enable photon and proton dose calculation. *Physics in Medicine & Biology* **2019**, *64*, 225004. https://doi.org/10.1088/1361-6560/ab4d8c.

50. Hansen, D.C.; Landry, G.; Kamp, F.; Li, M.; Belka, C.; Parodi, K.; Kurz, C. ScatterNet: A convolutional neural network for cone-beam CT intensity correction. *Medical Physics* **2018**, *45*, 4916–4926. https://doi.org/10.1002/mp.13175.

51. Thummerer, A.; Oria, C.S.; Zaffino, P.; Meijers, A.; Marmitt, G.G.; Wijsman, R.; Seco, J.; Langendijk, J.A.; Knopf, A.C.; Spadea, M.F.; et al. Clinical suitability of deep learning based synthetic CTs for adaptive proton therapy of lung cancer. *Medical Physics* **2021**, *48*, 7673–7684. https://doi.org/10.1002/mp.15333.