

ACCEPTED MANUSCRIPT • OPEN ACCESS

Deep learning based synthetic CT from cone beam CT generation for abdominal paediatric radiotherapy

To cite this article before publication: Adam Szmul *et al* 2023 *Phys. Med. Biol.* in press <https://doi.org/10.1088/1361-6560/acc921>

Manuscript version: Accepted Manuscript

Accepted Manuscript is “the version of the article accepted for publication including all changes made as a result of the peer review process, and which may also include the addition to the article by IOP Publishing of a header, an article ID, a cover sheet and/or an ‘Accepted Manuscript’ watermark, but excluding any other editing, typesetting or other changes made by IOP Publishing and/or its licensors”

This Accepted Manuscript is © 2023 The Author(s). Published on behalf of Institute of Physics and Engineering in Medicine by IOP Publishing Ltd.



As the Version of Record of this article is going to be / has been published on a gold open access basis under a CC BY 4.0 licence, this Accepted Manuscript is available for reuse under a CC BY 4.0 licence immediately.

Everyone is permitted to use all or part of the original content in this article, provided that they adhere to all the terms of the licence <https://creativecommons.org/licenses/by/4.0>

Although reasonable endeavours have been taken to obtain all necessary permissions from third parties to include their copyrighted content within this article, their full citation and copyright line may not be present in this Accepted Manuscript version. Before using any content from this article, please refer to the Version of Record on IOPscience once published for full citation and copyright details, as permissions may be required. All third party content is fully copyright protected and is not published on a gold open access basis under a CC BY licence, unless that is specifically stated in the figure caption in the Version of Record.

View the [article online](#) for updates and enhancements.

TITLE PAGE**TITLE**

Deep learning based synthetic CT from Cone Beam CT generation for abdominal paediatric radiotherapy

AUTHOR LIST

Dr Adam Szmul, a.szmul@ucl.ac.uk, Wellcome / EPSRC Centre for Interventional and Surgical Sciences, University College London, London, UK; Centre for Medical Image Computing, Department of Medical Physics and Biomedical Engineering, University College London, London, UK; ORCID: 0000-0002-3068-1068

Ms Sabrina Taylor, sabrina.taylor.18@ucl.ac.uk, Centre for Medical Image Computing, Department of Medical Physics and Biomedical Engineering, University College London, London, UK; ORCID:

Dr Pei Lim, pei.lim@nhs.net, Department of Oncology, University College London Hospitals NHS Foundation Trust, London, UK; ORCID: 0000-0001-9633-7908

Ms Jessica Cantwell, jessica.cantwell@newcastle.edu.au, Radiotherapy, University College London Hospitals NHS Foundation Trust, London, UK; ORCID:

Ms Isabel Moreira, i.silva@ucl.ac.uk, Centre for Medical Image Computing, Department of Medical Physics and Biomedical Engineering, University College London, London, UK; ORCID:

Ms Ying Zhang, ying.zhang.18@ucl.ac.uk, Department of Medical Physics and Biomedical Engineering, University College London, London, UK; ORCID:

Mr Derek D'Souza, derek.dsouza@nhs.net, Radiotherapy Physics Services, University College London Hospitals NHS Foundation Trust, London, UK; ORCID:0000-0002-4393-7683

Mr Syed Moinuddin, syed.moinuddin@nhs.net, Radiotherapy, University College London Hospitals NHS Foundation Trust, London, UK; ORCID:

1
2
3
4
5 Dr Mark Gaze, mgaze@nhs.net, Department of Oncology, University College London Hospitals NHS
6 Foundation Trust, London, UK; ORCID: 0000-0002-8344-7902
7
8
9

10 Dr Jennifer Gains, jenny.gains@nhs.net, Department of Oncology, University College London Hospitals
11 NHS Foundation Trust, London, UK; ORCID:
12
13
14

15 Dr Catarina Veiga, c.veiga@ucl.ac.uk, Centre for Medical Image Computing, Department of Medical
16 Physics and Biomedical Engineering, University College London, London, UK; ORCID: 0000-0002-
17 4132-2554
18
19
20
21
22

23 **CORRESPONDING AUTHOR**

24 Dr Adam Szmul
25 Wellcome / EPSRC Centre for Interventional and Surgical Sciences
26 University College London
27 90 High Holborn
28 Floor 1
29 London WC1V 6LJ
30 Tel: +44 (0) 203 549 5530
31 a.szmul@ucl.ac.uk
32
33
34
35
36
37
38
39
40
41

42 **KEYWORDS**

43 Cone-beam computed tomography (CBCT), computer tomography (CT), cycle-consistent generative
44 adversarial networks (cycleGANs), synthetic images, childhood cancer, abdominal neuroblastoma,
45 paediatric radiotherapy
46
47
48
49
50
51

52 **ACKNOWLEDGEMENTS**

53 This work was supported by the Radiation Research Unit at the Cancer Research UK City of London Centre
54 Award [C7893/A28990]. AS was supported by the National Institute for Health Research University College
55
56
57
58
59
60

1
2
3 London Hospitals Biomedical Research Centre and by the Wellcome Trust Institutional Strategic Support
4 Fund (WISSF). MG was supported by the National Institute for Health Research Biomedical Research
5 Centre of University College London Hospitals. CV was supported by the Royal Academy of Engineering
6 under the Research Fellowship scheme (RF\201718\17140).
7
8
9

10 11 12 13 **ETHICAL STATEMENT**

14 The data for this study was requested and approved in line with the internal information governance
15 procedures of the University College London Hospital NHS Foundation Trust Radiotherapy Department.
16
17
18
19
20

21 **NOVELTY & SIGNIFICANCE**

22 We present a novel method for synthetic CT generation from CBCTs. The proposed method was tailored
23 to address challenges in paediatric applications. We adapted cycle consistent Generative Adversarial
24 Networks (cycleGAN), introducing global residuals in the generators with structure-consistency loss and
25 weakly paired data approach via smart data selection. All these innovations improved the performance of
26 CBCT-to-CT synthesis in the proposed application outperforming a baseline cycleGAN method. Extensive
27 evaluation showed that the proposed method led to more realistic CT synthesis in terms of intensity
28 distance to ground truth images, allowing for more accurate structure segmentation and in radiotherapy
29 specific water equivalent thickness metric.
30
31
32
33
34
35
36
37
38
39
40
41
42
43
44
45
46
47
48
49
50
51
52
53
54
55
56
57
58
59
60

ABSTRACT

Objective: Adaptive radiotherapy workflows require images with the quality of computed tomography (CT) for re-calculation and re-optimisation of radiation doses. In this work we aim to improve quality of cone beam CT (CBCT) images for dose calculation using deep learning.

Approach: We propose a novel framework for CBCT-to-CT synthesis using cycle-consistent Generative Adversarial Networks (cycleGANs). The framework was tailored for paediatric abdominal patients, a challenging application due to the inter-fractional variability in bowel filling and smaller patient numbers. We introduced the concept of global residuals only learning to the networks and modified the cycleGAN loss function to explicitly promote structural consistency between source and synthetic images. Finally, to compensate for the anatomical variability and address the difficulties in collecting large datasets in the paediatric population, we applied a smart 2D slice selection based on the common field-of-view across the dataset (abdomen). This acted as a weakly paired data approach that allowed us to take advantage of scans from patients treated for a variety of malignancies (thoracic-abdominal-pelvic) for training purposes. We first optimised the proposed framework and benchmarked its performance on a development dataset. Later, a comprehensive quantitative evaluation was performed on an unseen dataset, which included calculating global image similarity metrics, segmentation-based measures and proton therapy-specific metrics.

Main results: We found improved performance, compared to a baseline implementation, on image-similarity metrics such as Mean Absolute Error calculated for a matched virtual CT (55.0 ± 16.6 proposed vs 58.9 ± 16.8 baseline). There was also a higher level of structural agreement for gastrointestinal gas between source and synthetic images measured through dice similarity overlap (0.872 ± 0.053 proposed vs 0.846 ± 0.052 baseline). Differences found in water-equivalent thickness metrics were also smaller for our method ($3.3 \pm 2.4\%$ proposed vs $3.7 \pm 2.8\%$ baseline).

Significance: Our findings indicate that our innovations to the cycleGAN framework improved the quality and structure consistency of the synthetic CTs generated.

MAIN BODY

35 1 Introduction

Abdominal irradiation is commonly used in the treatment of young patients with a variety of tumours, including abdominal neuroblastoma and Wilms' tumour (Bölling, Willich, and Ernst 2010). The success of radiotherapy relies on the accurate delivery of radiation dose to the target volume with minimal toxicity to surrounding normal tissues. Anatomical variations throughout the course of radiotherapy may lead to reduced tumour coverage and increased radiation to healthy tissues, thereby affecting the efficacy of radiotherapy (Berger et al. 2017). Although approaches to adapt radiotherapy plans have been extensively researched in the head and neck, thoracic and pelvis for adult populations (Ghilezan, Yan, and Martinez 2010; Sonke and Belderbos 2010; Thörnqvist et al. 2013; Morgan and Sher 2020; Tocco et al. 2020), fewer studies have exclusively focused in abdominal malignancies (Feng Liu et al. 2012; Schlaich et al. 2013) especially in younger populations (Laskar et al. 2015; Guerreiro et al. 2019). The abdominal and lower abdominal region is particularly susceptible to daily anatomical variations due to the presence of organs with variable filling, including the gastrointestinal (GI) tract, bowel, bladder, and rectum (Berger et al. 2017). GI air volumes were shown to vary by up to $\pm 80\%$ throughout radiotherapy in adult pancreatic cancer patients (Estabrook et al. 2018). In children with abdominal cancers, average GI air volume changes of 99.4 ± 126.9 ml (range: 216.7–454.7 ml) have been reported (Guerreiro et al 2019) as well as evidence that younger children are the most predisposed to variability (Lim et al. 2021; Sabrina Taylor et al. 2021). The presence or absence of GI air is reflected into substantial local tissue density changes. Density changes are particularly detrimental to proton beam therapy (PBT) treatments, a favourable radiation modality in abdominal paediatric cancer due to its tissue-sparing capabilities and potential for reducing long term side effects (Guerreiro et al. 2019; Lim et al. 2021; Sophie Taylor et al. 2021). A study on adult cervical cancer patients receiving PBT found correlations between dose degradation and volume, thickness, and width of bowel gas (Berger et al. 2017). The impact of GI air variation on radiotherapy plan robustness was shown to be more pronounced in PBT than x-ray intensity modulated arc therapy plans in both paediatric and adult cancers (Lim et al. 2021; Ashida et al. 2020; Mondlane et al. 2017).

1
2
3 Image-guided radiotherapy (IGRT) technologies, such as cone-beam-CT (CBCT), show the patients'
4 anatomy immediately before treatment. IGRT enables the monitoring of patient's anatomical variations
5 between when their planning CT was acquired and subsequent treatment fraction delivery sessions, and
6
7
8 65 potentially allows one to adjust the treatment to the observed anatomical variations (Nazmy et al. 2012).
9
10 A limitation to the direct use of CBCT in adaptive pathways is that the imaging quality of CBCT scans is
11 considered significantly inferior to the planning CT scans in terms of contrast-to-noise ratio and
12 prevalence of imaging artefacts such as streaks. In abdominal scans, streak artifacts can be attributed
13 (amongst others) to x-ray scatter and internal motion (Siewerdsen and Jaffray 2001; Peroni et al. 2012).
14
15
16 70 This has led to growing interest in developing methodology to make the quality of CBCT scans comparable
17 to that of CT.
18
19
20
21

22 The most established technique to generate synthetic CTs (synCT) scans with the image quality of CT is
23 based on deformable image registration (DIR) (Giacometti, Hounsell, and McGarry 2020), in which the
24
25 75 planning CT is deformed to match the CBCT's geometry (Peroni et al. 2012; Veiga et al. 2014; Landry et al.
26 2015). The main disadvantage of DIR-based approaches is that they cannot properly account for non-
27 deformable changes between consecutive scans, such as collapsing lungs or variable GI air volume and its
28 location. Post-processing methods may be paired with DIR to minimize gross anatomical mismatch to a
29 certain degree (Veiga et al. 2016). Most DIR algorithms also do not meet speed requirements to be usable
30
31
32
33 80 in real-time applications, particularly without GPU implementations (Shams et al. 2010; Fu et al. 2020).
34 An alternative approach is to directly apply scattering-corrections to CBCT images (Mainegra-Hing and
35 Kawrakow 2010; Park et al. 2015; Hansen et al. 2018). Deep learning (DL) methods became an emerging
36 and active field in recent years for medical image synthesis tasks such as CBCT-to-CT translation (Yu et al.
37 2020; Wang et al. 2021). Compared to classical approaches, data-driven methods have shown
38 encouraging performance metrics and can be applied to unseen datasets quickly. Their main
39 disadvantages are the significant efforts required for appropriate data collection and its curation (Wang
40 et al. 2021).
41
42 85
43
44
45
46
47
48

49 There is a wide breadth of previous work on using DL in image-to-image translation. The most popular
50 90 application of such task is MRI-to-CT conversion (Maspero, Houweling, et al. 2020; Florkow et al. 2020).
51 There are also successful implementations of DL-based solutions for PET attenuation correction
52 (Ladefoged et al. 2018) and CT synthesis from CBCT in the context of adaptive radiotherapy (Kurz et al.
53 2019; Y. Liu et al. 2020). Different frameworks utilising this approach have been proposed, such as paired
54
55
56
57
58
59
60

1
2
3 U-nets (Kida et al. 2018; Y. Li et al. 2019; Landry et al. 2019; L. Chen et al. 2020), paired pix2pix generative
4
5 95 adversarial networks (GANs) (Yang Zhang et al. 2020), paired cycle-consistent GANs (cycleGANs) (Harms
6
7 et al. 2019; Y. Liu et al. 2020; Eckl et al. 2020; Yang Zhang et al. 2020), and unpaired cycleGANs (Kurz et al.
8
9 2019; Liang et al. 2019; Gao et al. 2021). While all methods use a data-driven approach to map image
10
11 intensities between different imaging modalities, unpaired cycleGANs frameworks are of particular
12
13 100 offer good synthesis performance. This is of great interest in the context of CBCT-to-CT synthesis, where
14
15 simultaneous acquisition of both scans is unfeasible in practice, and using scans acquired closely in time
16
17 reduces but does not eliminate anatomical mismatch. Previous studies implementing paired approaches
18
19 used CTs and CBCTs acquired on the same day and applied DIR to compensate for residual anatomical
20
21 mismatch resulting from patient's position in the different scanners and potential internal anatomy
22 105 changes (L. Chen et al. 2020). These datasets still do not represent ideally paired examples and introduce
23
24 uncertainties. Additional challenges must also be considered when applying methodologies to younger
25
26 populations. IGRT protocols that include regular CT or CBCT imaging are rarely used in children's
27
28 treatments (Hua et al. 2019) due to concerns with the long term side effects associated with diagnostic
29
30 110 in adult populations and even harder to justify in younger patients, limiting the amount of data available
31
32 for training. Approaches capable of utilizing large and representative but unpaired datasets are very
33
34 advantageous in DL-based image synthesis frameworks, which usually benefit from larger dataset
35
36 (Shorten and Khoshgoftaar 2019; Brown et al. 2020).

37
38 115 Promising results were reported by several groups in the application of unpaired cycleGANs for CBCT-to-
39
40 CT synthesis (Kurz et al. 2019; Liang et al. 2019; Maspero, Houweling, et al. 2020; Uh et al. 2021). However,
41
42 there are still challenges in achieving CT-like quality in synthetic images and completely removing CBCT
43
44 artefacts. A well-known limitation of unpaired cycleGANs in medical image synthesis is that structural
45
46 consistency between source and synthetic images cannot be guaranteed, leading to incorrect anatomical
47 120 information in the synthetic images. Therefore, the original cycleGAN framework is not well suited for
48
49 CBCT-to-CT synthesis without addressing this limitation.

50
51 Data from younger cohorts require techniques specifically developed to account for the variability
52
53 presented due to disease, presentation, growth, and development from young age to adulthood.
54
55 125 Paediatric patients are a very diverse population, which likely reflects into a more challenging learning
56
57
58
59
60

1
2
3 task (Ladefoged et al. 2018; Maspero, Houweling, et al. 2020; Florkow et al. 2020). Childhood cancer is a
4 rare disease, making it more difficult to gather large, representative cohorts across all age groups for DL
5 applications (Guerreiro et al. 2019). CBCT imaging frequency may vary greatly between primary care
6 centres and types of radiotherapy used, often at the discretion of the treating physician (Nazmy et al.
7
8
9
10 130 2012). Low yearly patient numbers, combined with challenges in collecting imaging datasets in children,
11 make large datasets availability scarce, particularly for single institutions (Florkow et al. 2020). To address
12 limitations in data available for disease specific cohorts, combining multiple datasets from different
13 anatomical sites has been previously proposed; however, this has been achieved simply by including well-
14 balanced number of cases per patient group in training and evaluation (Maspero, Houweling, et al. 2020;
15
16
17
18 135 Uh et al. 2021). Transfer learning from adult cohorts is a viable option as well (Ladefoged et al. 2018).
19 However, there are intrinsic differences in the paediatric cancer population in comparison with adults that
20 will likely affect model generalizability. Differences include treatment strategies, such as the common use
21 of shunts and anaesthesia, and the inherent anatomical differences across developmental stages such as
22 variation in composition and shape of tissues and organs (White et al. 1991; Bolch et al. 2020).
23
24
25
26

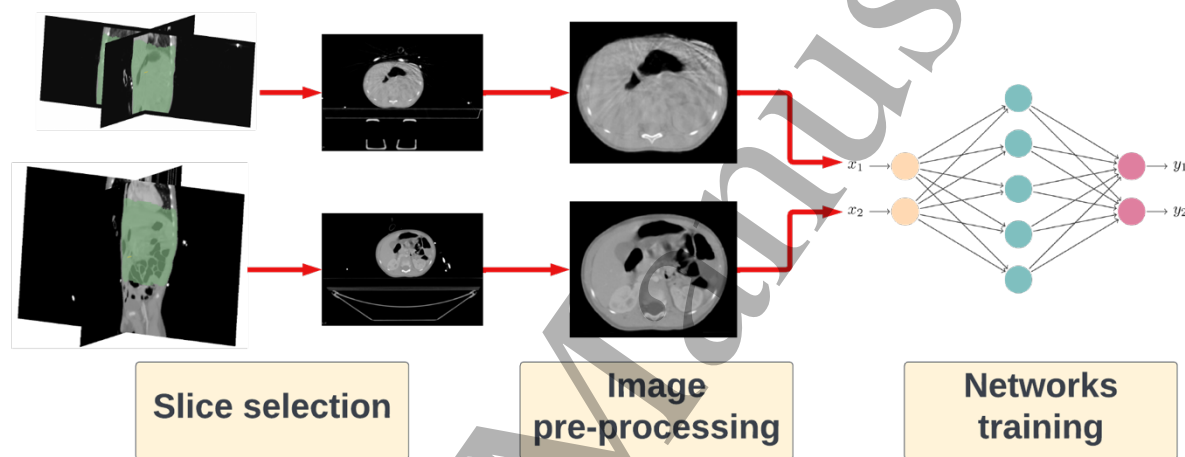
27 140
28 Key challenges remain that impede the usability of cycleGANs for CBCT-to-CT synthesis in clinical settings.
29 These challenges include how to ensure that structural consistency in the synthetic images is preserved
30 while removing unwanted artifacts, how to achieve large and representative sample sizes – particularly in
31 scarce data settings (as paediatrics), and how to define adequate ground-truths for the validation of novel
32 synthesis methods when paired data is not available. In this work we propose and evaluate a novel
33
34
35 145 framework for CBCT-to-CT synthesis tailored for paediatric abdominal patients, a challenging application
36 both due to inter-fractional variability in gastrointestinal filling and small patient numbers. This study
37 focuses on exploring improvements to the original cycleGAN framework and training data selection
38 techniques aiming to addresses the outlined challenges in the proposed application. Preliminary results
39
40
41
42
43 150 of this study were presented in conference publications (Szmul, Taylor, Lim, Cantwell, D'Souza,
44 Moinuddin, et al. 2021; Szmul et al. 2022). The key novel aspects of our framework are: (1) application of
45 the global residuals only learning approach, (2) incorporating structural consistency metrics to promote
46 anatomical plausibility of synthesized images, (3) a novel smart data selection process to efficiently
47 combine data from multiple patient groups (weakly-paired approach), (4) an automated pipeline for
48
49
50
51
52 155 evaluation of synthetic images.
53
54
55
56
57
58
59
60

2 Methods and Materials

2.1 A framework for CBCT-to-CT synthesis using cycleGANs

160

The developed synthetic CT (synCT) generation pipeline included the following key steps: (1) smart slice selection strategy, (2) image pre-processing and (3) network training and inference (Figure 1).



165 *Figure 1- Overview of the CBCT-to-CT synthesis framework developed, highlighting the three main steps: training data selection, image pre-processing and the networks training.*

2.1.1 Smart slice selection via weakly paired data approach

170 CT and CBCT scans do not cover the same sections of the body due to the reduced field-of-view of CBCT
 171 (in comparison to CT) and the intra-patient variability in the location of the imaging isocentre. We propose
 172 a weakly-paired data approach to compensate for the intra and inter-subject variability in imaged
 173 anatomical location. The CT and CBCT scans were spatially normalized to a common reference space and
 174 only slices from the same body regions were sampled (Figure 2). We used as reference space an atlas-
 175 based paediatric average anatomy and a co-registration strategy developed and evaluated in our previous
 176 study (Veiga et al. 2021). A region of interest was defined by fusing the co-registered CBCT body contours
 177 on the average space and applying a thresholded majority voting. The created common field-of-view mask
 178 was then propagated back to each subject's space. Slices in the individual CT/CBCT scans located outside

of this mask were excluded from all experiments. The top and bottom four slices of each CBCT scans were typically truncated and thus were also excluded from all experiments.

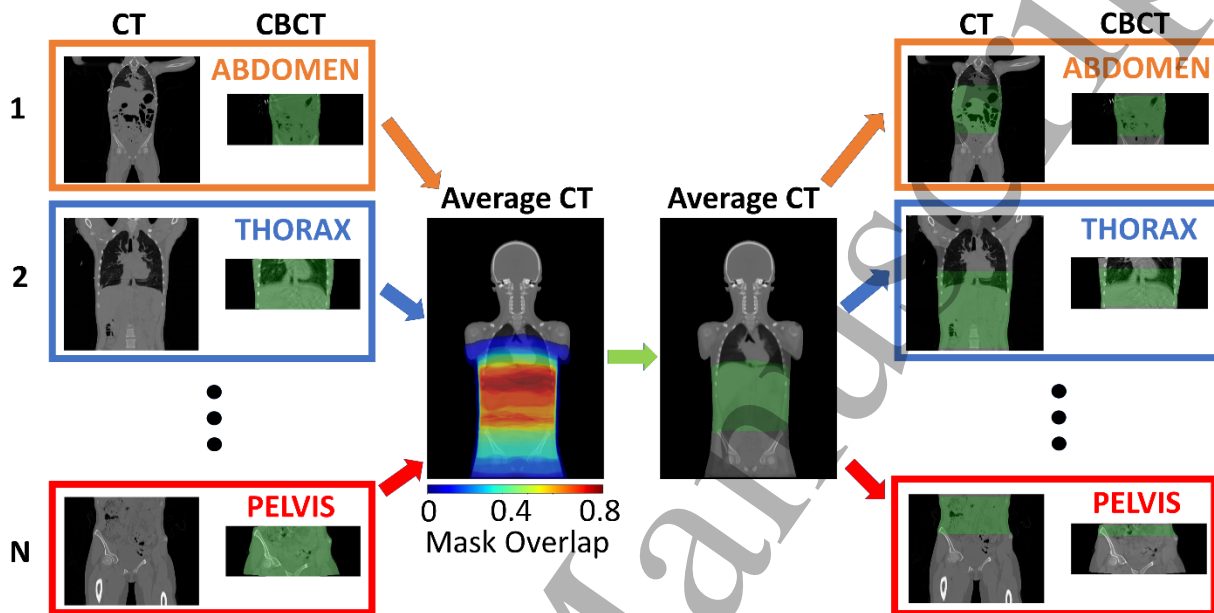


Figure 2- Overview of the proposed weakly-paired data approach. The CT and CBCT scans were spatially normalized to a paediatric anatomical atlas, and a common field-of-view mask generated after fusing the CBCT body contours. The mask was then mapped back to each individual space and used to include CT and CBCT slices only within the region of interest from training. This approach effectively adjusted the field-of-view of all images to the abdominal region only, allowing us to best utilize data from multiple patient groups.

2.1.2 Image pre-processing steps

Pre-processing steps were applied to all CT and CBCT scans before presenting them to the networks for training and inference. The images were corrected to exclude surrounding equipment and elements external to the subject – such as the treatment couch, anaesthesia equipment, shunts, feeding tubes, and/or lines, some of which may introduce high-intensity artifacts. External components were defined as regions outside the body contour and were replaced with air equivalent intensity (-1000HU). Internal regions with high-intensity artefacts were segmented by applying thresholding ($\text{HU} \geq 1700$) followed by morphological operations, and subsequently overwritten with water equivalent intensities (0HU). Finally, the image intensities were clipped to the range of $[-1000, 1000]$ and normalized to $[-1, 1]$. The proposed adjustments of intensities aimed at preventing the networks from generating elements such as tubes and

1
2
3 200 internal lines in the synthetic images that were not present in the source images but commonly present
4 in the training data.
5
6
7

8 To account for significant variation in body size and shape across ages, an axial normalisation pre-
9 processing step was also employed. The training CT/CBCT slices were axially normalized by fitting the body
10 contour to a fixed size of 256×256 and resizing the corresponding image slice. This step was done by
11 205 finding the longest profile between x and y axis on each slice; a margin of 10 pixels padded with -1000HU
12 was included to each slice to allow for additional variation in shape during data augmentation. The
13 determined distance was used to calculate the required scaling factor to best fit the slice to the fixed
14 image size of 256×256. The same scaling factor was applied on both directions to preserve body shape.
15
16
17
18
19
20 210 The images requiring resampling were interpolated using spline interpolation, using intensity clipping to
21 ensure the intensity ranges do not extend the normalisation ranges. The spatial normalisation of the body
22 aimed at artificially reducing the anatomical variability in size across the population while preserving
23 shape variability. In our preliminary investigations we have found consistent improvements by including
24 axial normalisation – without this step the results were often unrealistic and the body contours could be
25 distorted, particularly for smaller patients (Szmul, Taylor, Lim, Cantwell, D'Souza, Syed, et al. 2021). Our
26 observations were also confirmed by (Uh et al. 2021).
27
28
29
30
31
32

33 2.1.3 Design of the cycleGAN network 34 35 36

37 220 A 2D cycleGAN approach for CBCT-to-CT synthesis was implemented in this study, known for its good
38 performance in unpaired data style conversion (Figure 3). We followed closely the implementation
39 presented in (Zhu et al. 2017). The cycleGAN framework consists of two arms with a pair of a generator
40 and a discriminator in each of the arms. One arm converts CBCT to CT, and the counterpart generates
41 CBCT from CT. The task of a generator in this configuration is to, conditioned on an input image from one
42 modality, provide a corresponding image in another imaging modality.
43
44
45
46
47
48
49
50
51
52
53
54
55
56
57
58
59
60

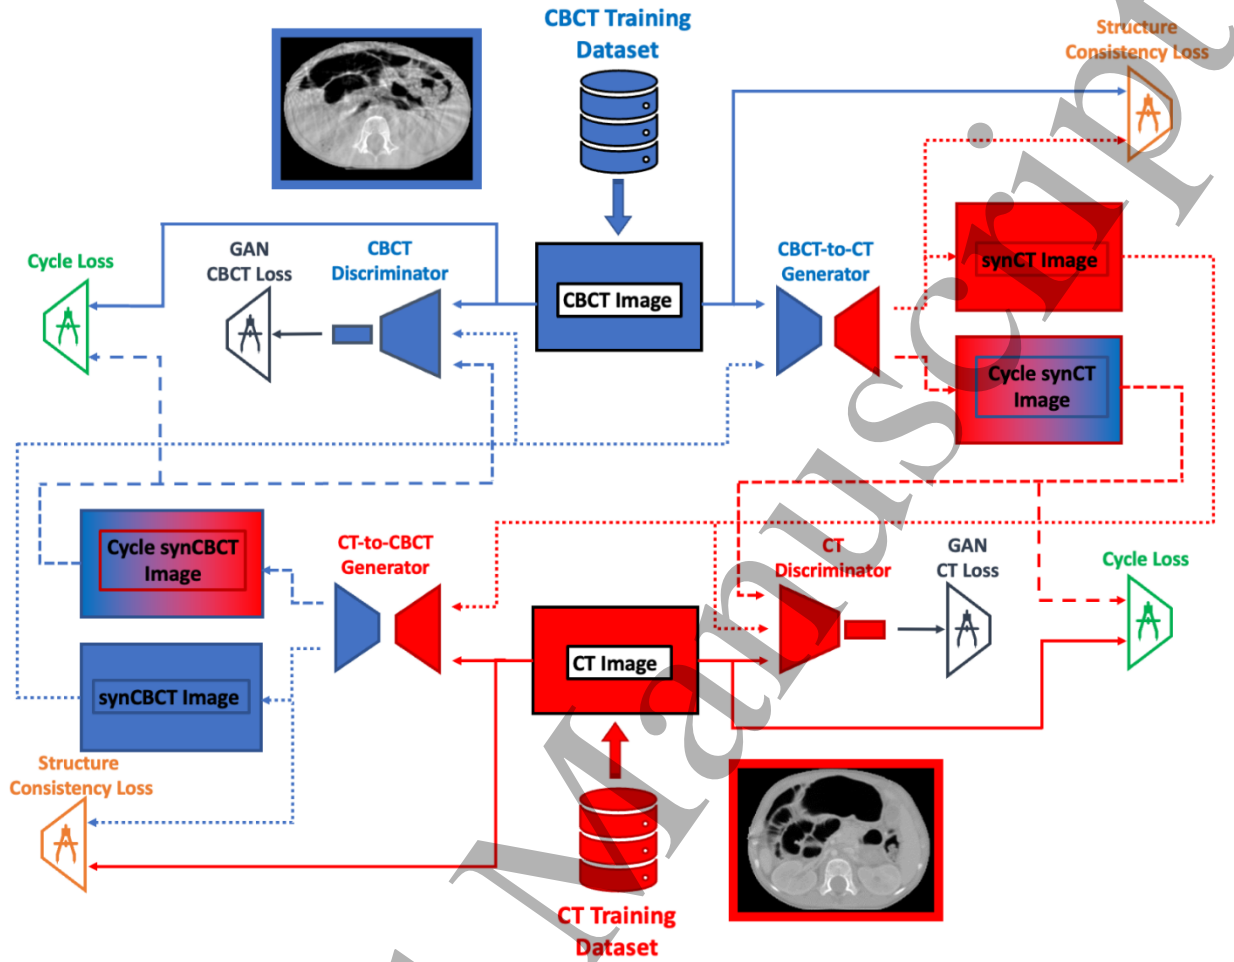
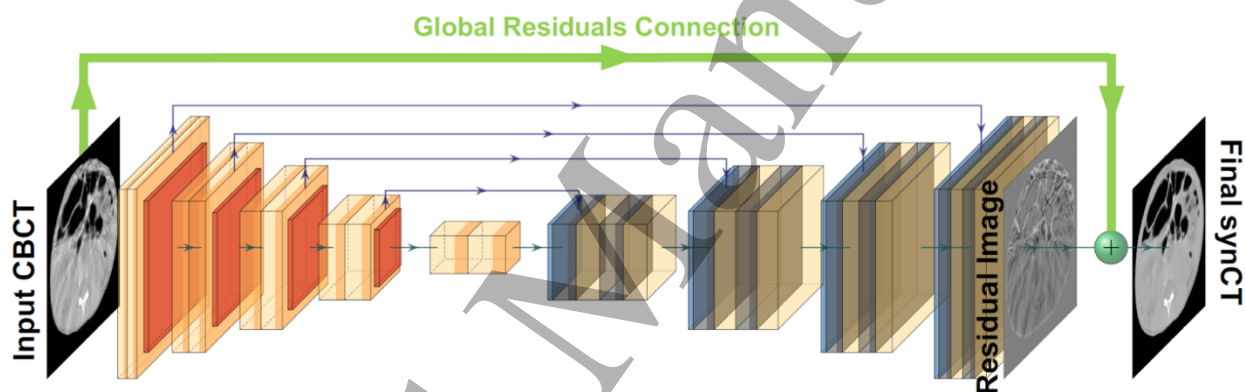


Figure 3. Overview of the proposed cycleGAN architecture. The main difference from the original cycleGAN implementation is the inclusion of the Structure Consistency Loss between raw input images and their synthetic contour parts, shown in orange colour. The paths of the different types of images are shown with dedicated line types: solid for the raw images, dotted for synthetic images and dashed for cycle images. Blue and red colours correspond to the modality arm paths, blue to the CBCT arm path and red to the CT arm path. For clarity the identity loss was not included in the diagram but described in the main body.

We investigated two different generator architectures: UNet-based (Ronneberger, Fischer, and Brox 2015) and ResNet-based (He et al. 2015). The ResNet-based generator is, similarly to UNet, an encoder-decoder architecture. However, it does not apply skip connections and uses residual blocks at the bottleneck stage. The implementation is based on works of (Johnson, Alahi, and Fei-Fei 2016; Isola et al. 2018), and has been successfully used in style adaptation and domain adaptation. For the rest of the paper, we will refer to these as UNet and ResNet generators respectively, to keep it consistent with the cycleGAN implementation of (Zhu et al. 2017). The choice of widely used generator architectures allowed our study to focus on improvements to the synthesis framework only.

1
2
3 A novel aspect of our cycleGAN CBCT-to-CT synthesis framework was incorporating the concept of global
4 residuals only learning in the generators. Global residual only learning is a common technique in computer
5 vision image restoration (Yulun Zhang et al. 2021), denoising (K. Zhang et al. 2017) and enhancement (Kim,
6 245 Lee, and Lee 2016). Unlike other pairs of very distinct imaging modalities (for example, CT and MR), CBCT
7 and CT share many acquisition similarities such that CBCT may be considered as distorted CT image to a
8 certain extent. Therefore, instead of learning the whole image, our proposed network focused only on
9 predicting the unwanted elements in the images, which were then combined with the source image to
10 produce the synthetic counterpart (Figure 4). This approach effectively redefines the network's aim from
11 synthesizing to refining/enhancing images. The proposed concept of global residuals was applied to both
12 generator architectures.
13
14
15 250
16
17
18
19
20
21
22



23
24
25
26
27
28
29
30
31
32
33
34 255 *Figure 4 Visualization of the concept of the only global residuals learning, by applying the global residual connection which*
35 *combines the input image with the output of the generator network. For illustrative purpose we show global residuals only with*
36 *UNet architecture, but the same concept was applied to ResNet.*
37
38
39

40
41 260 The task of the discriminators is to distinguish between real and synthetic images implemented as binary
42 classifiers with a binary cross-entropy loss function. The core loss in the framework is a GAN loss function
43 (\mathcal{L}_{GAN}), which is used to reward the generator for delivering outputs closer to the target domain, while
44 discriminator is rewarded for distinguishing between real and generated data. We applied an extended
45 version of the loss using Least Squares GAN loss functions (\mathcal{L}_{GAN}) which also notes how far from the
46 decision boundary the new generated image was, when evaluated by the discriminator (Mao et al. 2017).
47
48
49 265 CycleGANs utilise two arms, synthesizing CT into CBCTs and CBCT into CT, therefore there are two losses
50 corresponding to each arm. The cycle consistency loss (\mathcal{L}_{cycle}) evaluates how similar the input image is to
51 itself after going through both generators. This loss indirectly promotes structural consistency between
52 source and synthetic images but cannot guarantee it (Yunhao Ge et al. 2019). Consequently, studies have
53
54
55
56
57
58
59
60

1
2
3 reported for example differences in body outline in CT images generated with CBCT-to-CT cycleGAN
4 methodology (Kurz et al. 2019). To promote preservation of the CBCT anatomy, the original formulation
5 270 methodology (Kurz et al. 2019). To promote preservation of the CBCT anatomy, the original formulation
6 was here extended by introducing a structure consistency loss ($\mathcal{L}_{structure}$) in the form of locally
7 normalized cross correlation (LNCC) (Hermosillo, Chefd'Hotel, and Faugeras 2002) between the original
8 images and their synthesized counterparts. A similar approach was applied by Hiasa et al. 2018a) in
9 multimodal image synthesis, where cross correlation was calculated between gradients of CT and MRI
10 images. In our case, before the LNCC was calculated, the images were smoothed with gaussian filter and
11 gaussian noise was added to focus the attention of the measure into higher level structures. Additional
12 regularisation was promoted by the identity loss ($\mathcal{L}_{identity}$), which penalises applying changes to an image
13 275 introduced to the generator if the image is already in the target domain (for instance, the output of CT to
14 CBCT generator when a CBCT image is the input should be the same image). L1 was used as the identity
15 and cycle consistency losses. The overall proposed cycleGAN framework is optimised based on the total
16 loss function stated in Equation 1. Complete details of each loss function can be found in Supplementary
17 Data A.

$$\begin{aligned}
 \mathcal{L}(G_{CT \rightarrow CBCT}, G_{CBCT \rightarrow CT}, D_{CT}, D_{CBCT}) = & \\
 285 \quad & \mathcal{L}_{GAN_{CT \rightarrow CBCT}}(G_{CT \rightarrow CBCT}, D_{CBCT}) + \\
 & \mathcal{L}_{GAN_{CBCT \rightarrow CT}}(G_{CBCT \rightarrow CT}, D_{CT}) + \\
 & \lambda_{cycle} \mathcal{L}_{cycle}(G_{CBCT \rightarrow CT}, G_{CT \rightarrow CBCT}) + \\
 & \lambda_{identity} \mathcal{L}_{identity}(G_{CBCT \rightarrow CT}, G_{CT \rightarrow CBCT}) + \\
 & \lambda_{structure} \mathcal{L}_{structure}(G_{CBCT \rightarrow CT}, G_{CT \rightarrow CBCT})
 \end{aligned} \tag{1}$$

290 2.2 Data and data split

30
31
32
33
34
35
36
37
38
39
40
41
42
43
44
45
46
47
48
49
50
51
52
53
54
55
56
57
58
59
300
Scans from 63 patients aged 2 to 24 years old historically treated with radiation therapy were used in this
study. The data for this study was requested and approved in line with the internal information
governance procedures of the University College London Hospital NHS Foundation Trust Radiotherapy
Department. The smart slice selection process described in section 2.1.1 allowed us to make use of data
from multiple treatment groups. Thus, we included not only subjects irradiated to the abdomen (68%) but
also to the thoracic (8%) and pelvic (24%) region to increase the dataset size for training in this scarce data
domain. One planning CT and one to ten weekly CBCTs were gathered per patient, from a variety of
scanners and on-board imaging systems. The planning CTs were acquired on a SOMATON Confidence

(Siemens), LightSpeed RT16 or Discovery 710 (GE Medical Systems) with 120 kVp and field of view of 28 – 50 cm, resulting in reconstructed images with $0.93 (\pm 0.1) \times 0.93 (\pm 0.1) \times 2.29 (\pm 0.25) \text{ mm}^3$ resolution. Contrast enhancement was used in 70% of the CT scans. The CBCTs were acquired with the on-board imaging of the Varian Medical Systems Truebeam or Clinac with 125 kVp, 15 – 80 mA, 13 – 18 ms, half-fan mode, and shifted panels. The CBCT scans had a field of view of 41 – 46 cm and were reconstructed with a resolution of $0.91 (\pm 0.1) \times 0.91 (\pm 0.1) \times 1.99 (\pm 0.001) \text{ mm}^3$. In total 63 CT and 209 CBCT scans were available for the development and evaluation of the proposed CBCT-to-CT synthesis framework. In addition to the CT and CBCT scans, corresponding segmentations of the body, GI air, bone, and lung volumes were used for evaluation purposes. The volumes were first semi-automatically generated, and then manually edited and revised slice by slice using ITK-Snap (Yushkevich et al. 2006). Post-processing was applied to all contours to reduce common manual segmentation errors, such as discarding small volumes located outside the body.

The whole dataset was divided into 50 development cases and 13 cases for testing. The development dataset contained scans from thoracic-abdominal-pelvic subjects and was randomly split as 40 and 10 for training and validation. The testing dataset consisted of scans from abdominal subjects only, and cases were selected to achieve a well-balanced representation of ages and genders, while prioritizing the use of cases with multiple CBCTs for development purposes. The split between the datasets was made based on non-image characteristics only to minimise selection bias (i.e., the scans were not visually inspected during the splitting process).

We have created three training datasets by applying different slice selection approaches to our imaging data, which resulted in different number of slices being used for training. In the first dataset no data selection threshold was applied (naïve sampling). Two other datasets were created using the smart data selection via weakly paired approach, with two different thresholds to create the datasets: 1% and 40%. The 1% mask excluded rarely represented slices (hereby referred as “no outliers”) in a systematic way (such as for instance slices containing the neck or inferior members), which could accidentally be included in the naïve sampling. The 40% mask effectively adjusted the field-of-view of all images to the abdominal region, as shown previously in Figure 2 (hereby referred as “smart data selection”). A complete breakdown of the number of CT/CBCT slices included per approach can be found in Supplementary Data B.

2.3 Network training and inference

We trained a total of 16 cycleGAN frameworks by varying the generator architectures (Unet and ResNet), the slice selection strategy (naïve, no outliers and smart data selection), with and without global residuals, and including (or not) the structure consistency loss in the frameworks.

For the ResNet networks we followed the implementation used in (Johnson, Alahi, and Fei-Fei 2016; Isola et al. 2018). It consists of a down-sampling (encoder-like) section, 9 residual blocks at the bottleneck stage, and an up-sampling (decoder-like) section afterwards. The Unet-like architecture was a standard implementation by (Ronneberger, Fischer, and Brox 2015), with 5 down/up sampling levels, LeakyReLU activation function used in the downspampling blocks and ReLu in upsampling blocks, following (Isola et al. 2018). Both architectures were used with 64 initial filters, instance normalisation and dropout probability set to 0.5. The implementation of the discriminator followed PatchGAN discriminator with 70x70 overlapping patches, initially introduced in (Isola et al. 2018) and successfully applied by (Uh et al. 2021). The discriminator had 3 layers of depth, with 64 filters in the initial layer. Leaky ReLu was applied as the activation function. The cycleGAN frameworks were trained for 200 epochs with a batch size of 4 and diminishing learning rate on an in-house high-performance computing (HPC) facility with graphical processor units (Titan X 12GB GPU cards). The initial learning rate was set to 0.002 for the initial 100 epochs and was linearly decreased to 0 for the following 100 epochs. Adam optimiser was used with momentum $\beta_1=0.5$. Before calculating the Structure Consistency Loss, the images were smoothed with gaussian filter (kernel size 7, $\sigma=1$) and random gaussian noise was added (with the magnitude of 0.001). The losses were calculated with fixed weights for all experiments: $\lambda_{\text{cycle}}=10$, $\lambda_{\text{identity}}=0.5$, $\lambda_{\text{structure}}=1$ (when no structure consistency was used this weight was set to 0). For data augmentation, the following transformations were applied: random flipping (left and right), elastic deformations with the spacing set to (64, 64) and magnitude range to (3, 3), rotations by up to 15 degrees, and gamma contrast adjustments in a range [0.7, 1.3]. The probability of all random augmentations was set to 0.5.

At the inference stage, unseen datasets were fed to the network subjected to the same pre-processing steps as the training data (Section 2.1.2). The generated synCT were then resized back to their original size, their intensities were rescaled to [-1000, 1000] and stacked back into a 3D volume for evaluation.

2.4 Evaluation experiments

2.4.1 Definition of the ground-truth

Evaluating the quality of synthetic images is a challenging task due to the lack of a real ground-truth, which in this application would consist of pairs of simultaneously acquired CT and CBCT scans. Thus, we opted to evaluate the synCTs against two complementary image ground-truths: the raw CBCT and a virtual CT (vCT) matched to the anatomy of the raw CBCT. The vCT consisted of the planning CT deformably registered to the CBCTs using the open-source NiftyReg (Rueckert et al. 1999; Ourselin et al. 2001; Modat et al. 2010; 2014) with additional post-processing steps to account for the variable position of the GI air between scans (complete details in Supplementary Data C). Figure 5 shows an example of a planning CT and corresponding CBCT and vCT. The use of those two ground-truths enables the quantitative evaluation of both anatomical and intensity consistency of the synthetic images.




Figure 5- Example of paired CT, CBCT and virtual CT slices. The virtual CT was created by registering CT to CBCT and semi-automatically correcting for GI air differences. The air regions inside body in CT were replaced with water equivalent values prior to the registration, while air regions from CBCT were subsequently embedded into the virtual CT (that is indicated by white arrows in CT and in corresponding places in vCT).

2.4.2 Experiments

A total of three sets of experiments were conducted to optimise and evaluate the proposed CBCT-to-CT synthesis method.

1
2
3
4
5
6
7 390 to-CT synthesis framework. A total of 16 framework configurations were quantitatively evaluated
8 calculating global image similarity metrics between the gold-standards and the synthesized images
9 (details in Section 2.3.3).

10
11
12
13
14
15 395 more detail the optimal configuration of the framework on five-folds of the development dataset and
16 compared it against a baseline configuration. These experiments aimed to ascertain on the level of
17 overfitting of the method to the training dataset. The proposed cycleGANs was benchmarked against a
18 baseline implementation of the cycleGANs (also defined in Experiment 1). Both methods were evaluated
19 in terms of global image similarity with CBCT and vCT (details in Section 2.3.3).

20
21
22
23 400
24
25 *Experiment 3: Comprehensive quantitative evaluation on the unseen testing dataset.* The baseline and the
26 identified optimal configuration were retrained on the whole development dataset and used for inference
27 in the testing dataset. This experiment aimed to evaluate how the model generalises to unseen data, and
28 to validate its performance in detail in the proposed application. The proposed cycleGAN was again
29 compared to a baseline configuration of the cycleGAN. The performance of the networks was assessed
30 through global image similarity metrics, segmentation-based measures and radiotherapy-specific metrics
31 (details in Section 2.3.3).

32 405 33 34 35 36 37 38 2.4.3 Evaluation metrics 39

40 410
41
42 The quality of the CBCT-to-CT synthesis achieved by different methods was evaluated using metrics that
43 quantify how close the synthetic CT is to the original planning CT intensities and to the CBCT anatomy.
44 These metrics may be grouped into three separate categories: (1) global image similarity (experiments 1
45 – 3), (2) segmentation-based measures (experiment 3) and (3) radiotherapy-specific metrics (experiment
46 3). Global image similarity metrics as well radiotherapy-specific metrics were calculated in 2D (i.e., slice
47 by slice) in all experiments, while segmentation-based metrics were calculated in 3D. The metrics are
48 415 briefly described in the following paragraphs.

We used three global image similarity metrics when comparing pairs of images: mean absolute error (MAE), normalised cross-correlation (NCC) and Root Mean Square Error (RMSE). Those measures were calculated between the evaluated image and two ground truth images: CBCT and vCT. The metrics were chosen to provide complementary information on the similarity between synCTs in pixel intensities (MAE and RMSE) and structural agreement (NCC) with the CBCT/vCT. It is expected that MAE and RSME between synCT and CBCT will always reflect some disagreement from locally incorrect CBCT intensities. However, they may also reflect differences due to structural inconsistencies. Therefore, due to their complementary nature, all metrics were reported and analyzed for both ground-truths. Voxels outside the body contour were excluded from analysis based on semiautomatic binary masks. Intensities in all images were clipped to [-1000, 1000].

For segmentation-based evaluation, we used complementary measures of volume similarity between different types of tissue within the body contour: GI air, bones, soft tissues (muscles and fat) and lungs. Ground truth segmentations of these volumes on CBCTs were compared with corresponding structures automatically segmented on the synthetic CTs using the Dice Similarity Coefficient (DSC), the Hausdorff Distance (HD) and average pixel intensity (HU) as metrics. To automate the tissue segmentation we implemented and trained a patch-based 3D-UNet (Çiçek et al. 2016) using MONAI library (<https://monai.io/>) (Consortium 2020) (complete implementation details in Supplementary Data D). The training dataset consisted of a total of 183 CBCT and 50 CT scans, and corresponding ground truth labels, from the same subjects included in development dataset of the cycleGAN networks. The training dataset consisted of both CT and CBCT images. The assumption is that while synCTs are expected to have CT-like quality, some features/artifacts typical of CBCT may not be completely removed. Our initial experiments using only CTs led, for example, to gross segmentations errors in the presence of streak artifacts caused by high-intensity elements.

For radiotherapy-specific evaluation, we calculated polar water equivalent thickness (WET) differences between the vCT and synthetic CT methods (ΔWET). The WET is the thickness of water that would cause a proton beam to lose the same energy as if it had crossed a certain medium. The WET for a given beam direction (θ) was calculated as:

$$WET^{\theta} = \sum_{i,j,k \in S} RSP_{i,j,k} \times d_{i,j,k}$$

where S is a set of voxels that contains the beam path, $RSP_{i,j,k}$ is the relative stopping power (estimated from CT numbers using a standard calibration curve), and $d_{i,j,k}$ is the path length of the beam inside voxel (i, j, k) estimated by a ray tracing algorithm (R. Zhang et al. 2010; Lui, C F 2018). The WET between the beam entrance point and centre-of-mass of the body contour was calculated slice-by-slice, considering a complete arc with steps of one degree. The WET differences ($\Delta WET_{A,B}^\ominus$) between two scans (A and B) was reported as the RSME value

$$\Delta WET_{A,B}^\ominus = RMSE \left\{ 100 \times \frac{WET_B^\ominus - WET_A^\ominus}{WET_A^\ominus} \right\}$$

ΔWET was calculated overall and for each gantry angle individually, to quantify the impact that synthesis errors have on calculation of clinical dose distributions and ascertain if some beam angles were more affected than others.

3 Results

3.1 Experiment 1: Ablation study of the proposed modifications

Table 1 presents the numerical results in terms of the global similarity measures for the ablation study of the proposed modifications, where we tested a total of 16 configurations for the network. We systematically added the proposed modifications (global residual learning, smart data selection and structure consistency loss) and observed steadily improved performance of the framework, regardless of the generator architecture that was used. The largest improvements were observed by introducing the global residuals to the generators. Configurations that used the UNet architecture always outperformed the Resnet-based frameworks.

Table 1- Results of the ablation study of the proposed modification in terms of global similarity metrics for two generator networks architectures. The proposed configuration is highlighted in bold font in the grey box.

Architecture	Global residuals	Data Selection	Structure consistency	MAE_vCT [HU]	NCC_vCT [1]	RMSE_vCT [HU]	MAE_CBCT [HU]	NCC_CBCT [1]	RMSE_CBCT [HU]
ResNet	No	None	No	78.9±30.0	0.91±0.06	148.0±62.4	86.8±31.5	0.91±0.06	148.1±60.8
	No	No outliers	No	71.7±21.7	0.92±0.04	136.7±38.3	71.8±18.6	0.92±0.04	130.2±34.2
	Yes	None	No	58.7±15.8	0.96±0.02	102.3±29.8	63.8±17.3	0.96±0.02	107.3±31.3
	Yes	None	Yes	54.5±15.5	0.97±0.01	90.5±22.4	48.5±11.9	0.97±0.02	87.8±25.9
	Yes	No outliers	No	61.9±17.5	0.9±0.05	98.3±25.3	71.0±16.4	0.92±0.04	95.0±20.8
	Yes	No outliers	Yes	54.2±17.0	0.92±0.04	84.3±23.6	40.2±10.3	0.96±0.02	58.7±14.5
	Yes	Smart data selection	No	59.4±13.9	0.96±0.01	101.9±21.3	69.2±13.9	0.95±0.02	107.9±23.3
	Yes	Smart data selection	Yes	53.5±16.0	0.97±0.01	90.7±22.2	51.9±12.6	0.97±0.02	90.8±24.9
UNet	No	None	No	61.2±23.3	0.95±0.05	110.6±48.7	63.9±26.4	0.94±0.05	111.7±50.7
	No	No outliers	No	59.2±15.2	0.95±0.03	107.1±29.1	59.6±18.9	0.95±0.03	104.9±34.2
	Yes	None	No	54.4±12.8	0.97±0.01	90.0±20.0	61.4±16.3	0.96±0.02	97.1±28.0
	Yes	None	Yes	53.1±15.3	0.97±0.01	87.5±21.9	47.3±14.1	0.97±0.02	85.8±28.4
	Yes	No outliers	No	54.1±14.1	0.97±0.02	89.7±22.9	56.9±16.9	0.96±0.02	93.6±31.3
	Yes	No outliers	Yes	53.0±15.1	0.97±0.01	87.6±21.8	47.0±14.8	0.97±0.02	85.3±28.6
	Yes	Smart data selection	No	52.3±12.4	0.97±0.01	86.4±19.6	57.0±16.3	0.97±0.02	92.2±29.4
UNet	Yes	Smart data selection	Yes	51.7±15.0	0.97±0.01	85.0±20.9	45.0±12.6	0.97±0.02	82.8±27.5
MAE vCT/CBCT - Mean Absolute Error calculated with respect to vCT/CBCT									
NCC vCT/CBCT - Normalised Cross Correlation calculated with respect to vCT/CBCT									
RMSE vCT/CBCT - Root Mean Squared Error calculated with respect to vCT/CBCT									

475

This first experiment allowed us to narrow down into an optimal configuration, as well as a baseline configuration for comparison purposes, to use in the following experiments. The optimal configuration used the following settings: UNet as generator, applying the global residual learning, the smart data selection strategy and using structural constrains. We will refer to this configuration as the “proposed” approach for the rest of the manuscript. Likewise, the “baseline” cycleGAN configuration settings chosen were: Unet generator (Zhu et al. 2017), without structure consistency loss and global residuals learning trained with the “no outliers” slice selection method.

485

3.2 Experiment 2: Five-fold cross validation and benchmarking of the optimal configuration

In 5-fold cross validation study we investigated the proposed synCT generation configuration against the baseline cycleGAN configuration. The proposed method performed consistently between different folds. Figure 6 shows an example of inference results on the same slice for different folds, where in 4 folds the slice was used for training and in one it was part of the validation subset (4th fold). Only small levels of inconsistency were observed between folds, mostly originating from differences in where contrast enhancement is added by the networks (e.g: brightness of the kidneys and liver). These visual findings were also confirmed with the global similarity metrics, which were consistent for different folds and outperformed considerably the baseline method (Figure 7). To ascertain on the level of overfitting of the method to training dataset, the evaluation data is shown for both training and validation datasets – the results between the two datasets were comparable on average, with the validation set having larger standard deviations. The proposed method resulted not only in better mean values compared to the baseline method, but also in lower standard deviations. Overall, the numerical data suggests improved consistency with our method in agreement with the visual inspection findings.





Figure 6- Comparison of synthetic results between different folds when 5-fold cross validation was applied. The folds when the case was in the validation subset is shown in bold (fold 4). For the remaining images the case was used for training. No significant visual differences in anatomy were observed. In lower rows difference images between the outputs and CBCT and virtual CT are shown.



510 *Figure 7- Box plot showing the distribution of the numerical results between the proposed and baseline synCT for MAE (A), RMSE (B) and NCC (C) calculated between the synCTs and ground truths (CBCT and virtual CT).*

3.3 Experiment 3: Comprehensive quantitative evaluation on the unseen testing dataset

515 3.3.1 Global image similarity

The final experiments were conducted after re-training the proposed and baseline configuration on the whole development dataset and applying them to a previously unseen testing dataset. For clarity, due to the smart slice selection step being included, the final proposed network was effectively trained only on slices coming from the abdomen region (although some slices are from patients treated primarily for different anatomical regions). Figure 8 shows a visual comparison between the two configurations for two example slices from different subjects (A and B). The baseline configuration had worse performance in terms of anatomical realism and consistency when compared to both ground truths (CBCT and virtual CT). Our proposed method generated synthetic images that more closely matching the anatomy of the source CBCT. Upon visual inspection, the baseline method was found to commonly remove or add vertebrae and/or introducing inexistent bowel pockets, as well as contrast from shunts to the synCTs. Such patterns of failure were not observed for the proposed synCT method, which successfully preserved the anatomy from CBCT while improving the overall image quality. A limitation seen in both synCT methods was that occasionally the generated synCTs were contrast-enhanced, when the original CBCT or matching CT was not (and vice-versa), likely because the training set included both scans with and without contrast agent injection. That has led to inconsistencies in contrast between adjacent slices in some cases. The additional analysis of cross-sectional intensity profiles highlights how the baseline cycleGAN method was unable to preserve the structural information, adding up inexistent air pockets. The proposed method follows more closely the patterns observed in CBCT and virtual CT.

535

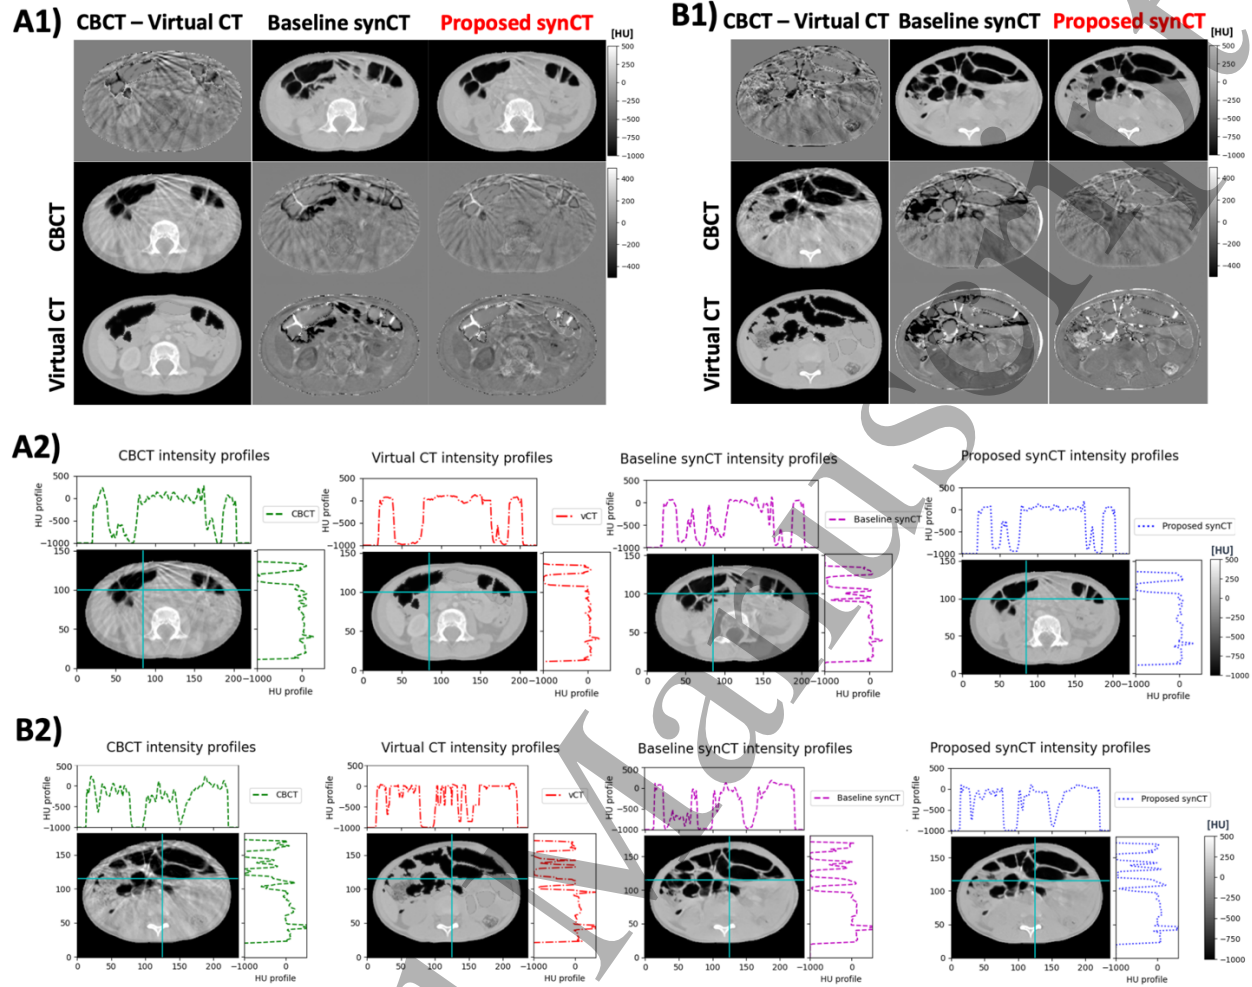


Figure 8- Visual comparison of the results between the proposed and baseline synCTs with respect to the virtual CT and CBCT for two cases A and B. We show difference images to highlight where disagreements appear in A1 and B1. The difference between the virtual CT and CBCT highlights the challenge with establishing reliable ground truth images. For both cases we show intensity profiles for cross-sections for the CBCT, virtual CT, baseline synCT and the proposed synCT in A2 and B2.

Visual inspection was followed by numerical evaluation in terms of global similarity measures. The proposed method achieved better numerical results than the baseline method for all scores (Table 2). These results were only slightly worse than those reported for the during 5-fold cross validation, indicating that the method did not overfit to the training dataset and generalized well to unseen cases.

Table 2- Numerical evaluation of the proposed method and the baseline in terms of global similarity measures.

	virtual CT			CBCT		
	MAE [HU]	NCC [1]	RMSE [HU]	MAE [HU]	NCC [1]	RMSE [HU]
Baseline synCT	58.9±16.8	0.96±0.02	102.6±31.7	63.4±15.9	0.95±0.02	110.0±32.2
Proposed synCT	55.0±16.6	0.97±0.02	89.8±23.8	49.8±10.9	0.97±0.02	88.6±24.9

550 3.3.2 Segmentation-based measures

A 3D UNet segmentation was trained to identify four different types of tissue in CT and CBCT within the body contour: soft tissues, skeleton, GI air and lungs. This network was used to automatically segment each testing CT, CBCT and synCT (Figure 9).




Figure 9 Segmentation results for CT, CBCT, and proposed synCT with comparison to the ground truth segmentation overlaid on CBCT.

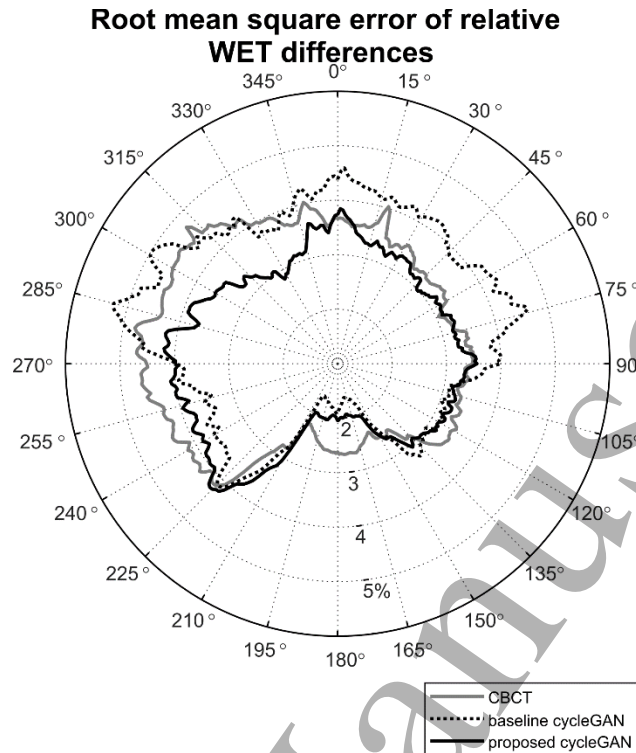
560 Table 3 presents the similarity between automated and ground truth segmentations in terms of DSC, HD and average intensity for each tissue type. The DSC and HD values reported for planning CT and CBCT represent how well the automated segmentation method performed on real images from the testing dataset, with DSC ranging from 0.889 – 0.978 for both scans and tissue types. The DSC and HD values reported for the synCTs allowed to assess the structural similarity (versus the source CBCTs). The values of DSC and HD reported for the synCT, while expectedly inferior to those on real CT and CBCT, were of similar magnitude indicating that the automated segmentation worked well on all datasets. The proposed synCT approach outperformed the baseline synCT demonstrating better structural similarity to the CBCT segmentations. The average intensity data allowed us to assess the intensity similarity between the synCTs and planning CTs for different tissue types. The proposed synCT method resulted in more closely matched mean HU value to planning CT for soft tissue and lungs. However, the proposed synCT method was found to slightly underestimate the skeleton Hus and to overestimate GI air intensities. Both methods achieved similar levels of standard deviations for all investigated tissue types (Supplementary Data D, Figure S1).

575 *Table 3- Numerical evaluation of the baseline and proposed synCT in terms of structure-based metrics. DSC and HD were calculated between ground truth contours (on CT and CBCT) and automatically generated contours (on CT, CBCT and synCT). The mean HU values were calculated within the ground truth contours (on CT and CBCT).*

		Planning CT	CBCT	Baseline synCT	Proposed synCT
Soft tissues	DSC	0.973±0.007	0.978±0.006	0.971±0.007	0.974±0.007
	HD	1.5±0.6	1.4±0.5	2.3±1.0	2.1±1.1
	Mean HU	2±1	0±26	24±20	9±20
Skeleton	DSC	0.896±0.011	0.889±0.022	0.849±0.024	0.862±0.020
	HD	1.5±0.8	2.0±2.5	5.9±6.6	3.4±5.2
	Mean HU	350±30	359±33	328±27	313±27
GI air	DSC	0.910±0.055	0.908±0.037	0.846±0.052	0.872±0.053
	HD	3.0±4.9	2.7±3.2	9.3±7.2	6.1±6.5
	Mean HU	-814±76	-737±93	-775±113	-756±95
Lungs	DSC	0.956±0.017	0.928±0.034	0.898±0.055	0.898±0.059
	HD	2.1±0.9	2.2±0.6	3.9±2.2	3.5±2.0
	Mean HU	-526±99	-573±67	-497±74	-547±71

3.3.3 Radiotherapy-specific metrics

580 The RMS error in relative WET differences between vCT and CBCT, baseline synCT and proposed synCT was 3.6±2.6%, 3.7±2.8% and 3.3±2.4%, respectively, when considering all gantry angles and all individual slices. The ΔWET for different gantry angles is shown in Figure 10. The WET measured for the proposed synCT best matches the vCT WET for anterior and anterior-oblique angles (0° to 90° and 270° to 360°), likely due to its better representation of the location of GI air pockets. The differences are smaller in the posterior direction, with baseline synCT being more similar to vCT in terms of WET for gantry angles between 160° and 270°. This is likely due to HU differences in high-intensity regions for those angles, such as vertebrae and liver and kidneys (in contrast enhanced scans).



590 *Figure 10- Polar map of the RMS value of the relative water equivalent thickness WET differences between virtual CT (ground-truth) and CBCT and deep-learning synthetic images.*

4 Discussion

595 We proposed and thoroughly evaluated a novel method for generating synthetic CTs from CBCTs based on cycleGANs. To the best of our knowledge, this is the first study to incorporate structural consistency loss and global residual learning with weakly paired data approach into CBCT-to-CT synthesis. Our novel smart slice selection framework was shown to facilitate training in diverse populations by allowing one to combine data from multiple patient cohorts in an optimal and efficient strategy. Global residuals learning combined with structural consistency loss helped to improve the structure correspondence between the input and output images, producing synthetic CT images that more closely preserve the structural information of the CBCT.

605 To the best of our knowledge, (Uh et al. 2021) was the only other study that has also investigated using cycleGANs for CBCT-to-CT synthesis in children and young adults undergoing PBT. In their study models were trained for the abdomen and pelvis data using a configuration of cycleGAN very close to what we

LNCC between the input and synthetic images. Although LNCC is commonly used in many medical image analysis applications, to the best of our knowledge its application with the image processing for CBCT-to-CT synthesis has not been proposed before. Alternative approaches to explicitly encourage direct correspondence between source and synthetic images included using similarity measures such as Modality Independent Neighbourhood Descriptor (Shrivastava et al. 2017), mean absolute error, gradient difference (Wang et al. 2019; Dai et al. 2020), correlation coefficients (Y Ge et al. 2019) and gradient correlation (Hiasa et al. 2018b). (L. Chen et al. 2020) combined mean absolute error loss with structure dissimilarity loss to encourage whole structure wise similarity. (Ouyang et al. 2019) employed a feature-matching technique where a new objective function was specified. In this case the generator encouraged the synthetic images to match the expected value of features on the intermediate layers instead of forcing it on the final output of the discriminator. The MaxInfo loss function proposed in (X. Chen et al. 2016) was applied to ensure correct pixel intensities of synthetic images. This function can be considered as an equivalent of histogram matching via informative maximising, working as an additional constrain to make the images follow the same intensity distribution. An alternative approach is to employ some sort of shape-consistency loss, which promotes similarity between annotated features, such as body contours (Yunhao Ge et al. 2019), organ sub-volumes (Z. Zhang, Yang, and Zheng, n.d.; Cai et al. 2019) or different tissue types (Fang Liu et al. 2018) between source and synthetic images. This is often employed in conjunction with segmentation networks to facilitate feature identification during the cycleGANS training. Changes to the generator architecture to capture multi-scale information have also been proposed, together with changes to the loss function to generate less blurry images (Lei et al. 2019). Others have used attention gates (Oktay et al. 2018) incorporated in the generated architecture to learn structural variations, improving prediction of the image intensities and organ boundaries (Y. Liu et al. 2020). Some solutions however are only valid when paired data is used for training, as the synthetic image is compared to some sort of ground-truth image. Our proposed approach with LNCC is task dedicated, making use of a well-suited similarity measure known to work well for two image modalities between which a linear relationship can be established. Additionally, the embedding gaussian noise and smoothing of the images helped to enhance global structure preservation. Our results showed that the introduction of the structure consistency loss improved the results for every configuration tested.

To address limitations in data available for disease specific cohorts, we proposed a weakly paired data method for training data. (Maspero, Houweling, et al. 2020) also proposed combining multiple datasets from different anatomical sites to generate a single, generalised network capable of performing on

multiple regions. In their study, scans from 33 head and neck, 33 lung and 33 breast subjects, with 15/8/10 split per site were used training validation and testing of a cycleGAN. The reported differences between single-site networks and a combined network trained on a dataset of combined sites were of up to 3HU (mean value), and no statistical significance was reported. Similarly, (Uh et al. 2021) used data from two patient groups to correct CBCT scans of children and young adults (28 abdominal and 36 pelvic, 64 cases in total). They found that the model trained on the combined dataset significantly outperformed the abdomen and pelvis models in terms of mean absolute HU error of the corrected CBCT from 14 testing patients (47 ± 7 HU versus 51 ± 8 HU). It is possible that the relatively small number of cases per site used for training led to the conclusion that including all the sites in the training results in a better method performance. Moreover, both studies used relatively well-balanced datasets with roughly equal number of cases for each anatomical site in included (Maspero, Houweling, et al. 2020) and 2:3 (Uh et al. 2021)). Our proposed approach, where images are initially registered to a common reference space and then sampled only from a chosen anatomical region is a less naïve technique of combining cases from different anatomical sites. Multiple anatomical regions can be combined in a systematic manner, such that all regions are similarly represented during training. The weakly paired approach improved the performance in the ablation study regardless of the chosen generator architecture and other settings. This demonstrates that it is advantageous to carefully consider presentation strategies of the data to the networks. There is merit into making the training datasets representative, not only larger, as more data did not necessarily lead to better performance.

Training an unsupervised model is a difficult task. We opted to train all models for a fixed number of epochs, while reducing learning rate, stabilizing the networks. Since all methods compared were based on the same generator architecture, therefore it was a reliable comparison of their performance. Early stopping, based on a metric calculated on a validation subset, could be an alternative, however the choice of the metric or a combination of multiple metrics becomes another challenge and variable within the framework. Our approach with a fixed number of epochs made it also consistent with other works in the field (Maspero, Bentvelzen, et al. 2020; Uh et al. 2021).

The acquisition of paired CBCT and CT images is not easily feasible. With paired data the synthesis could be framed as a supervised regression problem. (Qiu et al. 2021) performed deformable image registration between planning CTs and CBCTs, which resulted in pseudo paired data and allowed for fully supervised

1
2
3 learning. Such an approach is heavily dependent on the quality of the registration between planning CT
4 and CBCT. That might be challenging and in the case of abdomen with significant bowel gas changes,
5 require additional postprocessing steps and still not always successful. These challenges were
6 demonstrated in our work when generating our virtual CTs for evaluation purposes. While the DIR-based
7 vCT grossly corrected for both structural and intensity differences between pairs of CT and CBCT scans, it
8 was not a perfect representation of a CT-like scan paired to a CBCT. Therefore, unsupervised approaches,
9 where paired data are not required are better suited for the proposed task.
10
11
12
13
14
15

16
17 710 The segmentation-based evaluation allowed to explore in more detail how well different types of tissue
18 are represented in the synCTs. It is important to note that our aim was not to develop a segmentation
19 algorithm per se, but to generate quantitative and automated measures of the structural and intensity
20 quality of the synCT. There were challenges associated with manual and automated segmentation of each
21 tissue type. Segmentation of GI air pockets is associated with uncertainties due to scatter and motion
22 artifacts. It is not always clear from a low quality CBCT where to draw a boundary between gas and tissue.
23 For example, we could observe that the automated GI air segmentation was able to split air regions into
24 individual air pockets, while manual segmentations was more likely to connect them into larger pockets.
25 Furthermore, visual inspection indicated mismatch between the methods in identifying small pockets of
26 gas. Errors in the automatic skeleton segmentation mostly originated from contrast agents been classified
27 as skeleton. Our pre-processing step, where most of contrast agent areas were replaced with water
28 equivalent HU values, contributed to minimise issues in skeleton segmentation. The lung volumes were
29 underrepresented in the testing dataset since only a small fraction of the abdominal scans will contain
30 this tissue type. Motion artifacts at the diaphragm also contributed to making it increasingly challenging
31 to segment the lungs both manually and automatically. DSC and HD metrics were calculated on small
32 sections of the total lung size which may be reflected in the scores. While these uncertainties may impact
33 the results, in our opinion the methodology employed is accurate enough to compare structural similarity
34 between the different synthesis methods.
35
36
37
38
39
40
41
42 725
43
44
45
46
47

48
49 The analysis of the HUs for individual body tissue types indicated that some challenges remain in terms of
50 generating synthetic CTs with the same intensity information of the matched planning CT. For instance,
51 the mean value of the skeleton in CT was 359 HU, whereas in the proposed synCT it was 313 HU. Intensity
52 mismatch may propagate to dosimetric errors. While our method led to relatively small WET differences
53 (when compared to the baseline synCT) particularly for anterior proton beams, there were no clear
54
55
56
57
58
59
60

1
2
3 advantages for anterior proton beams. These angles could be associated path lengths that cross the
4
5 735 vertebrae and organs such as liver and kidneys (in contrast enhanced scans), where our HUs errors were
6 more pronounced. The observed differences could potentially be addressed by applying global histogram
7 normalisations (Zimmerman et al. 1988; Sandfort et al. 2019) as a post processing or by introducing
8 additional loss function within the optimisation framework (W. Li et al. 2019; Afifi, Brubaker, and Brown
9 2020).

10
11
12
13 740
14
15 One of the limitations observed in our study was occasional inconsistency between adjacent slices. Our
16 models were trained with real world CT and CBCT images from a variety of scanners and acquisition
17 settings. CT images used for training in our approach were both with and without contrast agent injection,
18 which could lead to inconsistencies where some slices were generated with contrast and others not. This
19
20
21 745 was not entirely unexpected, as our approach was 2D with no explicit adjacent slice consistency enforced.
22 With access to larger datasets, this could be solved by improving the selection of CT scans used for training
23 and ensuring that only the inclusion of scans with or without contrast. As an alternative approach to
24 address this limitation, we are considering exploring 3D approach as that will improve consistency but
25 would come with additional burden with larger memory and computational requirements. The pre-
26 processing method applied to all images corrected for the presence of elements such as anaesthesia
27 equipment, shunts, and lines. Since these elements are very common in our cohort, we realised during
28 our preliminary studies that the networks learned to spontaneously generate them in the synCT even
29 when they were not present in the source CBCT. We expect that this is more likely to happen when the
30 750 networks do not enforce structure similarity, so it is likely that this step is not required to train our
31 proposed network. This should be investigated in the future. Clinically, these elements are avoided by the
32 treatment beams if possible but, if they are inside the treatment volumes, the typical procedure is to
33 override their density for accurate dose calculations. The impact of incorrect representation in synCTs
34 could be corrected simply by propagating contours from the source CBCT. The proposed framework could
35 also be further improved by finetuning hyperparameters of the optimisation stage (learning rate and its
36 scheduler, optimiser and its parameters, etc) at data presentation (data augmentation parameters) and
37 760 loss calculation (individual loss weights). We decided to keep them fixed and for all the experiments after
38 initial parameters search for more interpretable results presentation.

39
40
41
42
43
44
45
46
47
48
49
50
51
52
53
54
55 765 In the future it would be interesting to investigate the performance of the cycleGAN models in low dose
56 CBCTs. In paediatric radiotherapy lower dose protocols are of interest (Bryce-Atkinson et al. 2021) but
57
58
59
60

1
2
3 may potentially result in lower quality and even more challenging task. (Gao et al. 2021) investigated
4 different GAN configurations in CBCT-to-CT synthesis in the thorax, noting that the increased imaging
5 artifacts inherent to lower dose CBCT protocols will disturb image translation tasks.
6
7
8
9

770

11 5 Conclusions

12
13
14
15
16 The proposed framework showed improved quality of synCTs generated from CBCTs when employing
17 strategies to preserve structural consistency and to account for variable field-of-view in the training
18 dataset. The reformulation of the problem from generating synthetic images to refining image quality by
19 775 applying global residuals only learning lead to the biggest improvements. Our study demonstrated the
20 advantages of thought-through data pre-processing and presentation to the AI method to improve its
21 performance on challenging real-world applications, with scarce and diverse data. A multi-step and multi-
22 layer evaluation allowed us to prove that the proposed method results in more realistic synCT generation.
23
24
25
26
27 780 Further evaluation using metrics of anatomical plausibility and realism, as well as impact on dose
28 calculations, is needed to provide insights into clinical utility.
29
30
31
32
33
34
35

36 785 References

- 37
38
39 Afifi, Mahmoud, Marcus A. Brubaker, and Michael S. Brown. 2020. "HistoGAN: Controlling Colors of
40 GAN-Generated and Real Images via Color Histograms."
41 <https://doi.org/10.48550/ARXIV.2011.11731>.
42
43 790 Alaei, Parham, and Emiliano Spezi. 2015. "Imaging Dose from Cone Beam Computed Tomography in
44 Radiation Therapy." *Physica Medica* 31 (7): 647–58.
45 <https://doi.org/10.1016/j.ejmp.2015.06.003>.
46
47 Ashida, Ryo, Mitsuhiro Nakamura, Michio Yoshimura, and Takashi Mizowaki. 2020. "Impact of
48 Interfractional Anatomical Variation and Setup Correction Methods on Interfractional Dose
49 795 Variation in IMPT and VMAT Plans for Pancreatic Cancer Patients: A Planning Study." *Journal of
50 Applied Clinical Medical Physics* 21 (7): 49–59. <https://doi.org/10.1002/acm2.12883>.
51
52 Berger, Thomas, Jørgen Breede Baltzer Petersen, Jacob Christian Lindegaard, Lars Ulrik Fokdal, and Kari
53 Tanderup. 2017. "Impact of Bowel Gas and Body Outline Variations on Total Accumulated Dose
54 800 with Intensity-Modulated Proton Therapy in Locally Advanced Cervical Cancer Patients." *Acta
55 Oncologica* 56 (11): 1472–78. <https://doi.org/10.1080/0284186X.2017.1376753>.
56
57
58
59
60

- 1
2
3 Bolch, W.E., K. Eckerman, A. Endo, J.G.S. Hunt, D.W. Jokisch, C.H. Kim, K-P. Kim, et al. 2020. "ICRP
4 Publication 143: Paediatric Reference Computational Phantoms." *Annals of the ICRP* 49 (1): 5–
5 297. <https://doi.org/10.1177/0146645320915031>.
6
7 Bölling, Tobias, Normann Willich, and Iris Ernst. 2010. "Late Effects of Abdominal Irradiation in Children:
8 805 A Review of the Literature." *Anticancer Research* 30 (1): 227–31.
9
10 Brown, Tom B., Benjamin Mann, Nick Ryder, Melanie Subbiah, Jared Kaplan, Prafulla Dhariwal, Arvind
11 Neelakantan, et al. 2020. "Language Models Are Few-Shot Learners." *ArXiv:2005.14165 [Cs]*,
12 July. <http://arxiv.org/abs/2005.14165>.
13
14 Bryce-Atkinson, Abigail, Rianne De Jong, Tom Marchant, Gillian Whitfield, Marianne C. Aznar, Arjan Bel,
15 810 and Marcel van Herk. 2021. "Low Dose Cone Beam CT for Paediatric Image-Guided
16 Radiotherapy: Image Quality and Practical Recommendations." *Radiotherapy and Oncology* 163
17 (October): 68–75. <https://doi.org/10.1016/j.radonc.2021.07.027>.
18
19 Cai, Jinzheng, Zizhao Zhang, Lei Cui, Yefeng Zheng, and Lin Yang. 2019. "Towards Cross-Modal Organ
20 Translation and Segmentation: A Cycle- and Shape-Consistent Generative Adversarial Network."
21 815 *Medical Image Analysis* 52 (February): 174–84. <https://doi.org/10.1016/j.media.2018.12.002>.
22
23 Chen, Liyuan, Xiao Liang, Chenyang Shen, Steve Jiang, and Jing Wang. 2020. "Synthetic CT Generation
24 from CBCT Images via Deep Learning." *Medical Physics* 47 (3): 1115–25.
25 <https://doi.org/10.1002/mp.13978>.
26
27 Chen, Xi, Yan Duan, Rein Houthoofd, John Schulman, Ilya Sutskever, and Pieter Abbeel. 2016. "InfoGAN:
28 820 Interpretable Representation Learning by Information Maximizing Generative Adversarial Nets."
29 *ArXiv:1606.03657 [Cs, Stat]*, June. <http://arxiv.org/abs/1606.03657>.
30
31 Çiçek, Özgün, Ahmed Abdulkadir, Soeren S. Lienkamp, Thomas Brox, and Olaf Ronneberger. 2016. "3D
32 U-Net: Learning Dense Volumetric Segmentation from Sparse Annotation."
33 <https://doi.org/10.48550/ARXIV.1606.06650>.
34
35 Consortium, The MONAI. 2020. "Project MONAI." Zenodo. <https://doi.org/10.5281/ZENODO.4323059>.
36
37 Dai, Xianjin, Yang Lei, Yabo Fu, Walter J. Curran, Tian Liu, Hui Mao, and Xiaofeng Yang. 2020.
38 "Multimodal MRI Synthesis Using Unified Generative Adversarial Networks." *Medical Physics* 47
39 (12): 6343–54. <https://doi.org/10.1002/mp.14539>.
40
41 Eckl, Miriam, Lea Hoppen, Gustavo R. Sarria, Judit Boda-Heggemann, Anna Simeonova-Chergou, Volker
42 830 Steil, Frank A. Giordano, and Jens Fleckenstein. 2020. "Evaluation of a Cycle-Generative
43 Adversarial Network-Based Cone-Beam CT to Synthetic CT Conversion Algorithm for Adaptive
44 Radiation Therapy." *Physica Medica: PM: An International Journal Devoted to the Applications of
45 Physics to Medicine and Biology: Official Journal of the Italian Association of Biomedical Physics
46 (AIFB)* 80 (December): 308–16. <https://doi.org/10.1016/j.ejmp.2020.11.007>.
47
48 Estabrook, Neil C., Jonathan B. Corn, Marvene M. Ewing, Higinia R. Cardenas, and Indra J. Das. 2018.
49 835 "Dosimetric Impact of Gastrointestinal Air Column in Radiation Treatment of Pancreatic Cancer."
50 *The British Journal of Radiology* 91 (1083): 20170512. <https://doi.org/10.1259/bjr.20170512>.
51
52 Florkow, Mateusz C., Filipa Guerreiro, Frank Zijlstra, Enrica Seravalli, Geert O. Janssens, John H. Maduro,
53 Antje C. Knopf, et al. 2020. "Deep Learning-Enabled MRI-Only Photon and Proton Therapy
54 840 Treatment Planning for Paediatric Abdominal Tumours." *Radiotherapy and Oncology, Physics
55 Special Issue: ESTRO Physics Research Workshops on Science in Development*, 153 (December):
56 220–27. <https://doi.org/10.1016/j.radonc.2020.09.056>.
57
58 Fu, Yabo, Yang Lei, Tonghe Wang, Walter J. Curran, Tian Liu, and Xiaofeng Yang. 2020. "Deep Learning in
59 845 Medical Image Registration: A Review." *Physics in Medicine and Biology* 65 (20): 20TR01.
60 <https://doi.org/10.1088/1361-6560/ab843e>.
61
62 Gao, Liugang, Kai Xie, Xiaojin Wu, Zhengda Lu, Chunying Li, Jiawei Sun, Tao Lin, Jianfeng Sui, and Xinye
63 Ni. 2021. "Generating Synthetic CT from Low-Dose Cone-Beam CT by Using Generative
64
65
66
67
68
69
70

- Adversarial Networks for Adaptive Radiotherapy." *Radiation Oncology* 16 (1): 202.
<https://doi.org/10.1186/s13014-021-01928-w>.
- 850 Ge, Y, Z Xue, T Cao, and S Liao. 2019. "Unpaired Whole-Body MR to CT Synthesis with Correlation Coefficient Constrained Adversarial Learning." *Proceedings SPIE Medical Imaging* 10949.
<https://doi.org/10.1117/12.2512479>.
- Ge, Yunhao, Dongming Wei, Zhong Xue, Qian Wang, Xiang Zhou, Yiqiang Zhan, and Shu Liao. 2019. "Unpaired Mr to CT Synthesis with Explicit Structural Constrained Adversarial Learning." In *2019 IEEE 16th International Symposium on Biomedical Imaging (ISBI 2019)*, 1096–99.
855 <https://doi.org/10.1109/ISBI.2019.8759529>.
- Ghilezan, Michel, Di Yan, and Alvaro Martinez. 2010. "Adaptive Radiation Therapy for Prostate Cancer." *Seminars in Radiation Oncology* 20 (2): 130–37.
<https://doi.org/10.1016/j.semradonc.2009.11.007>.
- 860 Giacometti, Valentina, Alan R. Hounsell, and Conor K. McGarry. 2020. "A Review of Dose Calculation Approaches with Cone Beam CT in Photon and Proton Therapy." *Physica Medica: PM: An International Journal Devoted to the Applications of Physics to Medicine and Biology: Official Journal of the Italian Association of Biomedical Physics (AIFB)* 76 (August): 243–76.
<https://doi.org/10.1016/j.ejmp.2020.06.017>.
- 865 Guerreiro, Filipa, Cornel Zachiu, Enrica Seravalli, Cássia O. Ribeiro, Geert O. Janssens, Mario Ries, Baudouin Denis de Senneville, et al. 2019. "Evaluating the Benefit of PBS vs. VMAT Dose Distributions in Terms of Dosimetric Sparing and Robustness against Inter-Fraction Anatomical Changes for Pediatric Abdominal Tumors." *Radiotherapy and Oncology: Journal of the European Society for Therapeutic Radiology and Oncology* 138 (September): 158–65.
870 <https://doi.org/10.1016/j.radonc.2019.06.025>.
- Hansen, David C., Guillaume Landry, Florian Kamp, Minglun Li, Claus Belka, Katia Parodi, and Christopher Kurz. 2018. "ScatterNet: A Convolutional Neural Network for Cone-beam CT Intensity Correction." *Medical Physics* 45 (11): 4916–26. <https://doi.org/10.1002/mp.13175>.
- 875 Harms, Joseph, Yang Lei, Tonghe Wang, Rongxiao Zhang, Jun Zhou, Xiangyang Tang, Walter J. Curran, Tian Liu, and Xiaofeng Yang. 2019. "Paired Cycle-GAN-Based Image Correction for Quantitative Cone-Beam Computed Tomography." *Medical Physics* 46 (9): 3998–4009.
<https://doi.org/10.1002/mp.13656>.
- He, Kaiming, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. 2015. "Deep Residual Learning for Image Recognition." *ArXiv:1512.03385 [Cs]*, December. <http://arxiv.org/abs/1512.03385>.
- 880 Hermosillo, Gerardo, Christophe Chefd'Hotel, and Olivier Faugeras. 2002. "Variational Methods for Multimodal Image Matching." *International Journal of Computer Vision* 50 (3): 329–43.
<https://doi.org/10.1023/A:1020830525823>.
- Hiasa, Yuta, Yoshito Otake, Masaki Takao, Takumi Matsuoka, Kazuma Takashima, Aaron Carass, Jerry L. Prince, Nobuhiko Sugano, and Yoshinobu Sato. 2018a. "Cross-Modality Image Synthesis from Unpaired Data Using CycleGAN." In *Simulation and Synthesis in Medical Imaging*, edited by Ali Gooya, Orcun Goksel, Ipek Oguz, and Ninon Burgos, 31–41. Lecture Notes in Computer Science. Cham: Springer International Publishing. https://doi.org/10.1007/978-3-030-00536-8_4.
- 885 ———. 2018b. "Cross-Modality Image Synthesis from Unpaired Data Using CycleGAN." In *Simulation and Synthesis in Medical Imaging*, edited by Ali Gooya, Orcun Goksel, Ipek Oguz, and Ninon Burgos, 31–41. Lecture Notes in Computer Science. Cham: Springer International Publishing.
890 https://doi.org/10.1007/978-3-030-00536-8_4.
- Hua, C.H., T.Z. Vern-Gross, C.B. Hess, A.J. Olch, V. Sathiaseelan, P. Alaei, J. Deng, et al. 2019. "Practice Patterns of Pediatric Image-Guided Radiotherapy (IGRT): A Children's Oncology Group Report." *International Journal of Radiation Oncology*Biophysics*Physics* 105 (1): S187.
895 <https://doi.org/10.1016/j.ijrobp.2019.06.234>.

- 1
2
3 Isola, Phillip, Jun-Yan Zhu, Tinghui Zhou, and Alexei A. Efros. 2018. "Image-to-Image Translation with
4 Conditional Adversarial Networks." *ArXiv:1611.07004 [Cs]*, November.
5 <http://arxiv.org/abs/1611.07004>.
6
- 7 Johnson, Justin, Alexandre Alahi, and Li Fei-Fei. 2016. "Perceptual Losses for Real-Time Style Transfer
8 and Super-Resolution." *ArXiv:1603.08155 [Cs]*, March. <http://arxiv.org/abs/1603.08155>.
9 900
- 10 Kida, Satoshi, Takahiro Nakamoto, Masahiro Nakano, Kanabu Nawa, Akihiro Haga, Jun'ichi Kotoku,
11 Hideomi Yamashita, and Keiichi Nakagawa. 2018. "Cone Beam Computed Tomography Image
12 Quality Improvement Using a Deep Convolutional Neural Network." *Cureus* 10 (4).
13 <https://doi.org/10.7759/cureus.2548>.
14 905
- 14 Kim, Jiwon, Jung Kwon Lee, and Kyoung Mu Lee. 2016. "Accurate Image Super-Resolution Using Very
15 Deep Convolutional Networks." In *2016 IEEE Conference on Computer Vision and Pattern
16 Recognition (CVPR)*, 1646–54. Las Vegas, NV, USA: IEEE.
17 <https://doi.org/10.1109/CVPR.2016.182>.
18
- 18 Kurz, Christopher, Matteo Maspero, Mark H. F. Savenije, Guillaume Landry, Florian Kamp, Marco Pinto,
19 Minglun Li, Katia Parodi, Claus Belka, and Cornelis A. T. van den Berg. 2019. "CBCT Correction
20 Using a Cycle-Consistent Generative Adversarial Network and Unpaired Training to Enable
21 Photon and Proton Dose Calculation." *Physics in Medicine & Biology* 64 (22): 225004.
22 <https://doi.org/10.1088/1361-6560/ab4d8c>.
23 910
- 23 Ladefoged, Claes Nøhr, Lisbeth Marner, Amalie Hindsholm, Ian Law, Liselotte Højgaard, and Flemming
24 Littrup Andersen. 2018. "Deep Learning Based Attenuation Correction of PET/MRI in Pediatric
25 Brain Tumor Patients: Evaluation in a Clinical Setting." *Frontiers in Neuroscience* 12: 1005.
26 <https://doi.org/10.3389/fnins.2018.01005>.
27 915
- 28 Landry, Guillaume, George Dedes, Christoph Zöllner, Josefine Handrack, Guillaume Janssens, Jonathan
29 Orban de Xivry, Michael Reiner, et al. 2015. "Phantom Based Evaluation of CT to CBCT Image
30 Registration for Proton Therapy Dose Recalculation." *Physics in Medicine and Biology* 60 (2):
31 595–613. <https://doi.org/10.1088/0031-9155/60/2/595>.
32 920
- 32 Landry, Guillaume, David Hansen, Florian Kamp, Minglun Li, Ben Hoyle, Jochen Weller, Katia Parodi,
33 Claus Belka, and Christopher Kurz. 2019. "Comparing Unet Training with Three Different
34 Datasets to Correct CBCT Images for Prostate Radiotherapy Dose Calculations." *Physics in
35 Medicine and Biology* 64 (3): 035011. <https://doi.org/10.1088/1361-6560/aaf496>.
36 925
- 36 Laskar, Siddhartha, Prakash Pandit, Suman Mallik, Pramod Tike, Suresh Chaudhari, Nehal Khanna, and
37 Tushar Vora. 2015. "Adaptive Radiation Therapy for Pediatric Head and Neck Malignancies:
38 Dosimetric Implications." *Practical Radiation Oncology* 5 (2): e87-94.
39 <https://doi.org/10.1016/j.prro.2014.07.005>.
40 930
- 41 Lei, Yang, Joseph Harms, Tonghe Wang, Yingzi Liu, Hui-Kuo Shu, Ashesh B. Jani, Walter J. Curran, Hui
42 Mao, Tian Liu, and Xiaofeng Yang. 2019. "MRI-Only Based Synthetic CT Generation Using Dense
43 Cycle Consistent Generative Adversarial Networks." *Medical Physics* 46 (8): 3565–81.
44 <https://doi.org/10.1002/mp.13617>.
45 935
- 46 Li, Wei, Wei Ding, Rajani Sadasivam, Xiaohui Cui, and Ping Chen. 2019. "H i s -GAN: A Histogram-Based
47 GAN Model to Improve Data Generation Quality." *Neural Networks* 119 (November): 31–45.
48 <https://doi.org/10.1016/j.neunet.2019.07.001>.
49 940
- 49 Li, Yinghui, Jinhan Zhu, Zhibin Liu, Jianjian Teng, Qiuying Xie, Liwen Zhang, Xiaowei Liu, Jinping Shi, and
50 Lixin Chen. 2019. "A Preliminary Study of Using a Deep Convolution Neural Network to Generate
51 Synthesized CT Images Based on CBCT for Adaptive Radiotherapy of Nasopharyngeal
52 Carcinoma." *Physics in Medicine & Biology* 64 (14): 145010. [https://doi.org/10.1088/1361-
53 6560/ab2770](https://doi.org/10.1088/1361-6560/ab2770).
54
- 55 Liang, Xiao, Liyuan Chen, Dan Nguyen, Zhiguo Zhou, Xuejun Gu, Ming Yang, Jing Wang, and Steve Jiang.
56 2019. "Generating Synthesized Computed Tomography (CT) from Cone-Beam Computed
57
58
59
60

- Tomography (CBCT) Using CycleGAN for Adaptive Radiation Therapy." *Physics in Medicine & Biology* 64 (12): 125002. <https://doi.org/10.1088/1361-6560/ab22f9>.
- 945 Lim, P.S., V. Rompokos, N. Bizzocchi, C. Gillies, A. Gosling, G. Royle, Y.-C. Chang, M.N. Gaze, and J.E. Gains. 2021. "Pencil Beam Scanning Proton Therapy Case Selection for Paediatric Abdominal Neuroblastoma: Effects of Tumour Location and Bowel Gas." *Clinical Oncology* 33 (3): e132–42. <https://doi.org/10.1016/j.clon.2020.08.012>.
- 950 Liu, Fang, Hyungseok Jang, Richard Kijowski, Tyler Bradshaw, and Alan B. McMillan. 2018. "Deep Learning MR Imaging–Based Attenuation Correction for PET/MR Imaging." *Radiology* 286 (2): 676–84. <https://doi.org/10.1148/radiol.2017170700>.
- Liu, Feng, Beth Erickson, Cheng Peng, and X. Allen Li. 2012. "Characterization and Management of Interfractional Anatomic Changes for Pancreatic Cancer Radiotherapy." *International Journal of Radiation Oncology, Biology, Physics* 83 (3): e423-429. <https://doi.org/10.1016/j.ijrobp.2011.12.073>.
- 955 Liu, Yingzi, Yang Lei, Tonghe Wang, Yabo Fu, Xiangyang Tang, Walter J. Curran, Tian Liu, Pretesh Patel, and Xiaofeng Yang. 2020. "CBCT-Based Synthetic CT Generation Using Deep-Attention CycleGAN for Pancreatic Adaptive Radiotherapy." *Medical Physics* 47 (6): 2472–83. <https://doi.org/10.1002/mp.14121>.
- 960 Lui, C F. 2018. "Density Heterogeneity Tool for Optimising Beam Angle Selection in Proton Therapy." *University College London Master's Thesis*.
- Mainegra-Hing, Ernesto, and Iwan Kawrakow. 2010. "Variance Reduction Techniques for Fast Monte Carlo CBCT Scatter Correction Calculations." *Physics in Medicine and Biology* 55 (16): 4495–4507. <https://doi.org/10.1088/0031-9155/55/16/S05>.
- 965 Mao, Xudong, Qing Li, Haoran Xie, Raymond Y. K. Lau, Zhen Wang, and Stephen Paul Smolley. 2017. "Least Squares Generative Adversarial Networks." *ArXiv:1611.04076 [Cs]*, April. <http://arxiv.org/abs/1611.04076>.
- Maspero, Matteo, Laura G. Bentvelzen, Mark H. F. Savenije, Filipa Guerreiro, Enrica Seravalli, Geert O. Janssens, Cornelis A. T. van den Berg, and Marielle E. P. Philippens. 2020. "Deep Learning-Based Synthetic CT Generation for Paediatric Brain MR-Only Photon and Proton Radiotherapy." *Radiotherapy and Oncology* 153 (December): 197–204. <https://doi.org/10.1016/j.radonc.2020.09.029>.
- 970 Maspero, Matteo, Antonetta C. Houweling, Mark H. F. Savenije, Tristan C. F. van Heijst, Joost J. C. Verhoeff, Alexis N. T. J. Kotte, and Cornelis A. T. van den Berg. 2020. "A Single Neural Network for Cone-Beam Computed Tomography-Based Radiotherapy of Head-and-Neck, Lung and Breast Cancer." *Physics and Imaging in Radiation Oncology* 14 (April): 24–31. <https://doi.org/10.1016/j.phro.2020.04.002>.
- 975 Modat, Marc, David M. Cash, Pankaj Daga, Gavin P. Winston, John S. Duncan, and Sébastien Ourselin. 2014. "Global Image Registration Using a Symmetric Block-Matching Approach." *Journal of Medical Imaging* 1 (2): 024003. <https://doi.org/10.1117/1.JMI.1.2.024003>.
- 980 Modat, Marc, Gerard R. Ridgway, Zeike A. Taylor, Manja Lehmann, Josephine Barnes, David J. Hawkes, Nick C. Fox, and Sébastien Ourselin. 2010. "Fast Free-Form Deformation Using Graphics Processing Units." *Computer Methods and Programs in Biomedicine* 98 (3): 278–84. <https://doi.org/10.1016/j.cmpb.2009.09.002>.
- 985 Mondlane, Gracinda, Michael Gubanski, Pehr A. Lind, Ana Ureba, and Albert Siegbahn. 2017. "Comparison of Gastric-Cancer Radiotherapy Performed with Volumetric Modulated Arc Therapy or Single-Field Uniform-Dose Proton Therapy." *Acta Oncologica* 56 (6): 832–38. <https://doi.org/10.1080/0284186X.2017.1297536>.
- 990 Morgan, Howard E., and David J. Sher. 2020. "Adaptive Radiotherapy for Head and Neck Cancer." *Cancers of the Head & Neck* 5 (1): 1. <https://doi.org/10.1186/s41199-019-0046-z>.

- 1
2
3
4
5
6
7
8
9
10
11
12
13
14
15
16
17
18
19
20
21
22
23
24
25
26
27
28
29
30
31
32
33
34
35
36
37
38
39
40
41
42
43
44
45
46
47
48
49
50
51
52
53
54
55
56
57
58
59
60
- 995 Nazmy, Mohamed Soliman, Yasser Khafaga, Amr Mousa, and Ehab Khalil. 2012. "Cone Beam CT for Organs Motion Evaluation in Pediatric Abdominal Neuroblastoma." *Radiotherapy and Oncology: Journal of the European Society for Therapeutic Radiology and Oncology* 102 (3): 388–92. <https://doi.org/10.1016/j.radonc.2011.11.009>.
- 1000 Oktay, Ozan, Jo Schlemper, Loic Le Folgoc, Matthew Lee, Mattias Heinrich, Kazunari Misawa, Kensaku Mori, et al. 2018. "Attention U-Net: Learning Where to Look for the Pancreas." *ArXiv:1804.03999 [Cs]*, May. <http://arxiv.org/abs/1804.03999>.
- 1005 Ourselin, S, A Roche, G Subsol, X Pennec, and N Ayache. 2001. "Reconstructing a 3D Structure from Serial Histological Sections." *Image and Vision Computing* 19 (1–2): 25–31. [https://doi.org/10.1016/S0262-8856\(00\)00052-4](https://doi.org/10.1016/S0262-8856(00)00052-4).
- 1010 Ouyang, Jiahong, Kevin T. Chen, Enhao Gong, John Pauly, and Greg Zaharchuk. 2019. "Ultra-Low-Dose PET Reconstruction Using Generative Adversarial Network with Feature Matching and Task-Specific Perceptual Loss." *Medical Physics* 46 (8): 3555–64. <https://doi.org/10.1002/mp.13626>.
- 1015 Park, Yang-Kyun, Gregory C. Sharp, Justin Phillips, and Brian A. Winey. 2015. "Proton Dose Calculation on Scatter-Corrected CBCT Image: Feasibility Study for Adaptive Proton Therapy." *Medical Physics* 42 (8): 4449–59. <https://doi.org/10.1118/1.4923179>.
- 1020 Peroni, Marta, Delia Ciardo, Maria Francesca Spadea, Marco Riboldi, Stefania Comi, Daniela Alterio, Guido Baroni, and Roberto Orecchia. 2012. "Automatic Segmentation and Online VirtualCT in Head-and-Neck Adaptive Radiation Therapy." *International Journal of Radiation Oncology, Biology, Physics* 84 (3): e427-433. <https://doi.org/10.1016/j.ijrobp.2012.04.003>.
- 1025 Qiu, Richard L. J., Yang Lei, Aparna H. Kesarwala, Kristin Higgins, Jeffrey D. Bradley M.d, Walter J. Curran, Tian Liu, and Xiaofeng Yang. 2021. "Chest CBCT-Based Synthetic CT Using Cycle-Consistent Adversarial Network with Histogram Matching." In *Medical Imaging 2021: Image Processing*, 11596:115961Z. International Society for Optics and Photonics. <https://doi.org/10.1117/12.2581094>.
- 1030 Ronneberger, Olaf, Philipp Fischer, and Thomas Brox. 2015. "U-Net: Convolutional Networks for Biomedical Image Segmentation." *ArXiv:1505.04597 [Cs]*, May. <http://arxiv.org/abs/1505.04597>.
- 1035 Rueckert, D., L.I. Sonoda, C. Hayes, D.L.G. Hill, M.O. Leach, and D.J. Hawkes. 1999. "Nonrigid Registration Using Free-Form Deformations: Application to Breast MR Images." *IEEE Transactions on Medical Imaging* 18 (8): 712–21. <https://doi.org/10.1109/42.796284>.
- 1040 Sandfort, Veit, Ke Yan, Perry J. Pickhardt, and Ronald M. Summers. 2019. "Data Augmentation Using Generative Adversarial Networks (CycleGAN) to Improve Generalizability in CT Segmentation Tasks." *Scientific Reports* 9 (1): 16884. <https://doi.org/10.1038/s41598-019-52737-x>.
- 1045 Schlaich, Fabian, Stephan Brons, Thomas Haberer, Jürgen Debus, Stephanie E. Combs, and Klaus-Josef Weber. 2013. "Comparison of the Effects of Photon versus Carbon Ion Irradiation When Combined with Chemotherapy in Vitro." *Radiation Oncology (London, England)* 8 (November): 260. <https://doi.org/10.1186/1748-717X-8-260>.
- 1050 Shams, Ramtin, Parastoo Sadeghi, Rodney Kennedy, and Richard Hartley. 2010. "A Survey of Medical Image Registration on Multicore and the GPU." *IEEE Signal Processing Magazine* 27 (2): 50–60. <https://doi.org/10.1109/MSP.2009.935387>.
- 1055 Shorten, Connor, and Taghi M. Khoshgoftaar. 2019. "A Survey on Image Data Augmentation for Deep Learning." *Journal of Big Data* 6 (1): 60. <https://doi.org/10.1186/s40537-019-0197-0>.
- 1060 Siewerdsen, J. H., and D. A. Jaffray. 2001. "Cone-Beam Computed Tomography with a Flat-Panel Imager: Magnitude and Effects of x-Ray Scatter." *Medical Physics* 28 (2): 220–31. <https://doi.org/10.1118/1.1339879>.
- 1065 Sonke, Jan-Jakob, and José Belderbos. 2010. "Adaptive Radiotherapy for Lung Cancer." *Seminars in Radiation Oncology* 20 (2): 94–106. <https://doi.org/10.1016/j.semradonc.2009.11.003>.

- 1
2
3 Szmul, Adam, S Taylor, P Lim, J Cantwell, D D'Souza, S Moinuddin, M Gaze, J Gains, and C Veiga. 2022.
4 1040 "OC-0773 CBCT-to-CT Synthesis Using Weakly-Paired Cycle-Consistent Generative Adversarial
5 Networks." *Radiotherapy and Oncology* 170: S690–92.
6
7 Szmul, Adam, Sabrina Taylor, Pei Lim, Jessica Cantwell, Derek D'Souza, Syed Moinuddin, Mark Gaze,
8 Jennifer Gains, and Catarina Veiga. 2021. "Developing a Framework for CBCT-to-CT Synthesis in
9 Paediatric Abdominal Radiotherapy." *Medical Image Understanding and Analysis*.
10 1045 Szmul, Adam, Sabrina Taylor, Pei Lim, Jessica Cantwell, Derek D'Souza, Moinuddin Syed, Mark Gaze,
11 Jennifer Gains, and Catarina Veiga. 2021. "Developing a Framework for CBCT-to-CT Synthesis in
12 Paediatric Abdominal Radiotherapy." *Medical Image Analysis and Understanding*.
13 Taylor, Sabrina, Pei Lim, Jessica Cantwell, Syed Moinuddin, D. D'Souza, Mark Gaze, Jennifer Gains, and
14 Catarina Veiga. 2021. "Identifying Predictors of Inter-Fractional Anatomical Variations in
15 1050 Paediatric Abdominal Radiotherapy" 68: S30–31.
16 Taylor, Sophie, Pei Lim, Reem Ahmad, Ammar Alhadi, William Harris, Vasilis Rompokos, Derek D'Souza,
17 Mark Gaze, Jennifer Gains, and Catarina Veiga. 2021. "Risk of Radiation-Induced Second
18 Malignant Neoplasms from Photon and Proton Radiotherapy in Paediatric Abdominal
19 Neuroblastoma." *Physics and Imaging in Radiation Oncology* 19 (July): 45–52.
20 1055 <https://doi.org/10.1016/j.phro.2021.06.003>.
21 Thörnqvist, Sara, Ludvig P. Muren, Lise Bentzen, Liv B. Hysing, Morten Høyer, Cai Grau, and Jørgen B. B.
22 Petersen. 2013. "Degradation of Target Coverage Due to Inter-Fraction Motion during Intensity-
23 Modulated Proton Therapy of Prostate and Elective Targets." *Acta Oncologica (Stockholm,*
24 *Sweden)* 52 (3): 521–27. <https://doi.org/10.3109/0284186X.2012.752860>.
25 1060 Tocco, Boris R., Amar U. Kishan, Ting Martin Ma, Linda G. W. Kerkmeijer, and Alison C. Tree. 2020. "MR-
26 Guided Radiotherapy for Prostate Cancer." *Frontiers in Oncology* 10 (December): 616291.
27 <https://doi.org/10.3389/fonc.2020.616291>.
28 Uh, Jinsoo, Chuang Wang, Sahaja Acharya, Matthew J. Krasin, and Chia-ho Hua. 2021. "Training a Deep
29 Neural Network Coping with Diversities in Abdominal and Pelvic Images of Children and Young
30 Adults for CBCT-Based Adaptive Proton Therapy." *Radiotherapy and Oncology* 160 (July): 250–
31 1065 58. <https://doi.org/10.1016/j.radonc.2021.05.006>.
32 Veiga, Catarina, Guillaume Janssens, Ching-Ling Teng, Thomas Baudier, Lucian Hotoiu, Jamie R.
33 McClelland, Gary Royle, et al. 2016. "First Clinical Investigation of Cone Beam Computed
34 Tomography and Deformable Registration for Adaptive Proton Therapy for Lung Cancer."
35 1070 *International Journal of Radiation Oncology, Biology, Physics* 95 (1): 549–59.
36 <https://doi.org/10.1016/j.ijrobp.2016.01.055>.
37 Veiga, Catarina, Pei Lim, Virginia Marin Anaya, Edward Chandy, Reem Ahmad, Derek D'Souza, Mark
38 Gaze, Syed Moinuddin, and Jennifer Gains. 2021. "Atlas Construction and Spatial Normalisation
39 to Facilitate Radiation-Induced Late Effects Research in Childhood Cancer." *Physics in Medicine*
40 *& Biology* 66 (10): 105005. <https://doi.org/10.1088/1361-6560/abf010>.
41 1075 Veiga, Catarina, Jamie McClelland, Syed Moinuddin, Ana Lourenço, Kate Ricketts, James Annkah, Marc
42 Modat, Sébastien Ourselin, Derek D'Souza, and Gary Royle. 2014. "Toward Adaptive
43 Radiotherapy for Head and Neck Patients: Feasibility Study on Using CT-to-CBCT Deformable
44 Registration for 'Dose of the Day' Calculations." *Medical Physics* 41 (3): 031703.
45 1080 <https://doi.org/10.1118/1.4864240>.
46 Wang, Tonghe, Yang Lei, Yabo Fu, Jacob F. Wynne, Walter J. Curran, Tian Liu, and Xiaofeng Yang. 2021.
47 "A Review on Medical Imaging Synthesis Using Deep Learning and Its Clinical Applications."
48 *Journal of Applied Clinical Medical Physics* 22 (1): 11–36. <https://doi.org/10.1002/acm2.13121>.
49 1085 Wang, Tonghe, Yang Lei, Zhen Tian, Xue Dong, Yingzi Liu, Xiaojun Jiang, Walter J. Curran, Tian Liu, Hui-
50 Kuo Shu, and Xiaofeng Yang. 2019. "Deep Learning-Based Image Quality Improvement for Low-
51
52
53
54
55
56
57
58
59
60

- 1
2
3 Dose Computed Tomography Simulation in Radiation Therapy." *Journal of Medical Imaging*
4 (*Bellingham, Wash.*) 6 (4): 043504. <https://doi.org/10.1117/1.JMI.6.4.043504>.
- 5
6 White, D. R., E. M. Widdowson, H. Q. Woodard, and J. W. Dickerson. 1991. "The Composition of Body
7 Tissues (II). Fetus to Young Adult." *The British Journal of Radiology* 64 (758): 149–59.
8 1090 <https://doi.org/10.1259/0007-1285-64-758-149>.
- 9 Yu, Biting, Yan Wang, Lei Wang, Dinggang Shen, and Luping Zhou. 2020. "Medical Image Synthesis via
10 Deep Learning." *Advances in Experimental Medicine and Biology* 1213: 23–44.
11 https://doi.org/10.1007/978-3-030-33128-3_2.
- 12
13 1095 Yushkevich, Paul A., Joseph Piven, Heather Cody Hazlett, Rachel Gimpel Smith, Sean Ho, James C. Gee,
14 and Guido Gerig. 2006. "User-Guided 3D Active Contour Segmentation of Anatomical Structures:
15 Significantly Improved Efficiency and Reliability." *NeuroImage* 31 (3): 1116–28.
16 <https://doi.org/10.1016/j.neuroimage.2006.01.015>.
- 17
18 Zhang, Kai, Wangmeng Zuo, Yunjin Chen, Deyu Meng, and Lei Zhang. 2017. "Beyond a Gaussian
19 1100 Denoiser: Residual Learning of Deep CNN for Image Denoising." *IEEE Transactions on Image*
20 *Processing* 26 (7): 3142–55. <https://doi.org/10.1109/TIP.2017.2662206>.
- 21
22 Zhang, Rui, Phillip J Taddei, Markus M Fitzek, and Wayne D Newhauser. 2010. "Water Equivalent
23 Thickness Values of Materials Used in Beams of Protons, Helium, Carbon and Iron Ions." *Physics*
24 *in Medicine and Biology* 55 (9): 2481–93. <https://doi.org/10.1088/0031-9155/55/9/004>.
- 25
26 1105 Zhang, Yang, Ning Yue, Min-Ying Su, Bo Liu, Yi Ding, Yongkang Zhou, Hao Wang, Yu Kuang, and Ke Nie.
27 2020. "Improving CBCT Quality to CT Level Using Deep-Learning with Generative Adversarial
28 Network." *Medical Physics*, December. <https://doi.org/10.1002/mp.14624>.
- 29
30 Zhang, Yulun, Yapeng Tian, Yu Kong, Bineng Zhong, and Yun Fu. 2021. "Residual Dense Network for
31 Image Restoration." *IEEE Transactions on Pattern Analysis and Machine Intelligence* 43 (7):
32 2480–95. <https://doi.org/10.1109/TPAMI.2020.2968521>.
- 33
34 1110 Zhang, Zizhao, Lin Yang, and Yefeng Zheng. n.d. "Translating and Segmenting Multimodal Medical
35 Volumes With Cycle- and Shape-Consistency Generative Adversarial Network," 10.
- 36
37 Zhu, Jun-Yan, Taesung Park, Phillip Isola, and Alexei A. Efros. 2017. "Unpaired Image-to-Image
38 Translation Using Cycle-Consistent Adversarial Networks." In *2017 IEEE International Conference*
39 *on Computer Vision (ICCV)*, 2242–51. <https://doi.org/10.1109/ICCV.2017.244>.
- 40
41
42
43
44
45
46
47
48
49
50
51
52
53
54
55
56
57
58
59
60 1115 Zimmerman, J.B., S.M. Pizer, E.V. Staab, J.R. Perry, W. McCartney, and B.C. Brenton. 1988. "An
Evaluation of the Effectiveness of Adaptive Histogram Equalization for Contrast Enhancement."
IEEE Transactions on Medical Imaging 7 (4): 304–12. <https://doi.org/10.1109/42.14513>.

1
2
3
4
5
6
7
8
9
10
11
12
13
14
15
16
17
18
19
20
21
22
23
24
25
26
27
28
29
30
31
32
33
34
35
36
37
38
39
40
41
42
43
44
45
46
47
48
49
50
51
52
53
54
55
56
57
58
59
60

Accepted Manuscript