

# Deep conditional generative model for longitudinal single-slice abdominal computed tomography harmonization

Xin Yu<sup>①,a,\*</sup> Qi Yang<sup>①,a</sup> Yucheng Tang<sup>①,b</sup> Riqiang Gao<sup>①,a</sup> Shunxing Bao<sup>①,b</sup>  
Leon Y. Cai,<sup>c</sup> Ho Hin Lee<sup>①,a</sup> Yuankai Huo<sup>①,a,b</sup> Ann Zenobia Moore,<sup>d</sup>  
Luigi Ferrucci,<sup>d</sup> and Bennett A. Landman<sup>①,a,b,c</sup>

<sup>a</sup>Vanderbilt University, Department of Computer Science, Nashville, Tennessee, United States

<sup>b</sup>Vanderbilt University, Department of Electrical and Computer Engineering, Nashville, Tennessee, United States

<sup>c</sup>Vanderbilt University, Department of Biomedical Engineering, Nashville, Tennessee, United States

<sup>d</sup>National Institute on Aging, Baltimore, Maryland, United States

---

**ABSTRACT.** **Purpose:** Two-dimensional single-slice abdominal computed tomography (CT) provides a detailed tissue map with high resolution allowing quantitative characterization of relationships between health conditions and aging. However, longitudinal analysis of body composition changes using these scans is difficult due to positional variation between slices acquired in different years, which leads to different organs/tissues being captured.

**Approach:** To address this issue, we propose C-SliceGen, which takes an arbitrary axial slice in the abdominal region as a condition and generates a pre-defined vertebral level slice by estimating structural changes in the latent space.

**Results:** Our experiments on 2608 volumetric CT data from two in-house datasets and 50 subjects from the 2015 Multi-Atlas Abdomen Labeling Challenge Beyond the Cranial Vault (BTCV) dataset demonstrate that our model can generate high-quality images that are realistic and similar. We further evaluate our method's capability to harmonize longitudinal positional variation on 1033 subjects from the Baltimore longitudinal study of aging dataset, which contains longitudinal single abdominal slices, and confirmed that our method can harmonize the slice positional variance in terms of visceral fat area.

**Conclusion:** This approach provides a promising direction for mapping slices from different vertebral levels to a target slice and reducing positional variance for single-slice longitudinal analysis. The source code is available at: <https://github.com/MASILab/C-SliceGen>.

© The Authors. Published by SPIE under a Creative Commons Attribution 4.0 International License. Distribution or reproduction of this work in whole or in part requires full attribution of the original publication, including its DOI. [DOI: [10.1117/1.JMI.11.2.024008](https://doi.org/10.1117/1.JMI.11.2.024008)]

**Keywords:** abdominal slice generation; longitudinal data harmonization; body composition

Paper 23272GR received Sep. 13, 2023; revised Jan. 18, 2024; accepted Mar. 14, 2024; published Apr. 2, 2024.

---

## 1 Introduction

Body compositional analysis is an important term to determine an individual's health condition which refers to the percentage of fat, muscle, and bone percentages in the human body.<sup>1</sup> Studying the change of body composition on aging enables better prognosis and early disease detection for

---

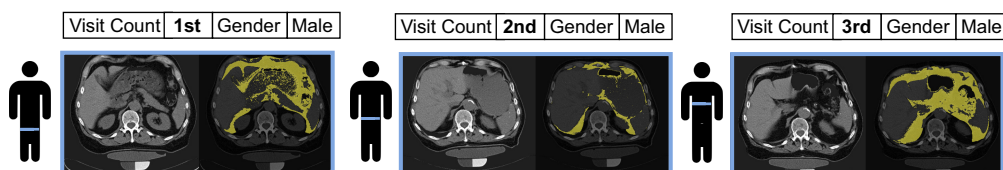
\*Address all correspondence to Xin Yu, [xin.yu@vanderbilt.edu](mailto:xin.yu@vanderbilt.edu)

various diseases, such as heart disease,<sup>2</sup> sarcopenia,<sup>3</sup> and diabetes.<sup>4</sup> Computed tomography body composition is a widely employed technique for assessing body composition.<sup>5</sup>

Existing longitudinal CT scans of the abdomen from the Baltimore longitudinal study of aging (BLSA) dataset<sup>6</sup> provide a valuable opportunity to characterize the relationship between body composition changes and age-related disease, cognitive disease, and metabolic health.<sup>7-12</sup> To minimize radiation exposure for longitudinal imaging and potential risk associated with contrast administration, two-dimensional (2D) non-contrast axial single-slice CT is taken as opposed to three-dimensional (3D) volumetric CT commonly acquired in clinical practice. However, it is difficult to locate the same cross-sectional location in longitudinal imaging, and thus there is substantial variation in the organs and tissues captured in different years, as shown in Fig. 1. The organs and tissues scanned in 2D abdominal slices strongly correlate with body composition measures. Therefore, increased positional variance can make accurately analyzing body composition challenging. Despite this issue, no method has been proposed to address the problem of positional variance in 2D slices.

Our goal is to decrease the effects of positional variance in body composition analysis, to facilitate more precise longitudinal interpretation. A major challenge is that the distance between the scans taken in different years is unknown, as the slice can be taken at any abdominal region. Image registration is a commonly used technique in other contexts for correcting pose or positioning errors. However, this approach is not suitable for addressing out-of-plane motion in 2D acquisitions where the tissues/organs that appear in one scan may not appear in the other scan. Based on Ref. 13, image harmonization methods are categorized into two main groups: deep learning and statistical methods. Notable statistical methods include Combat<sup>14</sup> and its variants,<sup>15-17</sup> ConvBat,<sup>18</sup> and Bayesian factor regression.<sup>19</sup> However, unlike generative models, statistical methods often lack the generative capability crucial for our scenario.

Modern generative models based on deep learning have recently shown significant success in generating and reconstructing high-quality and realistic images.<sup>20-26</sup> The fundamental concept of generative modeling is to train a generative model to learn a distribution so that the generated samples  $\hat{x} \sim p_d(\hat{x})$  are from the same distribution as the training data distribution  $x \sim p_d(x)$ .<sup>27</sup> By learning the joint distribution between the input and target slices, these models can effectively address the limitations of registration. Variational autoencoders (VAEs),<sup>28</sup> which are a type of generative model, consisting of an encoder and a decoder. The encoder encodes inputs to an interpretable latent distribution, and the decoder decodes the samples of the latent distribution to new data. Generative adversarial networks (GANs)<sup>20</sup> are another type of generative model, which contains two sub-models, a generator model that generates new data and a discriminator that distinguishes between real and generated images. By playing this two-player min-max game, GANs can generate realistic images. VAEGAN<sup>29</sup> incorporates GAN into the VAE framework to create better-synthesized images. By using the discriminator to distinguish between real and generated images, VAEGAN can generate more realistic and high-quality images than traditional VAE models. However, original VAEs and GANs suffer from the limitation of lack of control over the generated images. This issue is addressed by conditional GAN (cGAN)<sup>30</sup> and conditional VAE (cVAE)<sup>31</sup> which allow for generating specific images with a condition, providing more control over the generated outputs. However, the majority of these conditional methods necessitate specific target information, such as a target class, semantic map, or heatmap,<sup>32</sup> as a condition during the testing phase, which is not feasible in our scenario since we do not have any direct target information available.



**Fig. 1** An example of a subject with slices acquired at different vertebral levels in different visits. The blue line represents the approximate axial position where the CT scan is taken. The yellow masks represent visceral fat. The shape and size of the captured organs and tissues vary largely among different visits leading to large variations in the visceral fat area.

To provide a condition during testing, we aim to have the network generate a slice of that specific target at a pre-determined vertebral level, which will serve as the generation target. By defining the target slice, the generative model will implicitly learn the organ/tissue composition in the target slice and have this condition learned during training time. We hypothesize that by giving an arbitrary abdominal slice, the model will generate the slice at the target vertebral level while preserving subject-specific information derived from the conditional image such as body habitus. Inspired by Refs. 32–34, we introduce the conditional SliceGen (C-SliceGen) model based on VAEGAN, which enables the generation of subject-specific target vertebral level slices from an arbitrary abdominal slice input. We use 3D volumetric data to train and validate our model since in 3D data the target slice [ground truth (GT)] is available for direct comparison with the generated images. The training datasets include an in-house portal venous phase CT and an in-house non-contrast phase CT volume with 1120 and 1488 subjects, respectively. We further evaluate on the 2015 Multi-Atlas Abdomen Labeling Challenge Beyond the Cranial Vault (BTCV) dataset<sup>35</sup> for external validation. Structural similarity index (SSIM),<sup>36</sup> peak signal-to-noise ratio (PSNR),<sup>37</sup> learned perceptual image patch similarity (LPIPS),<sup>38</sup> and normalized mutual information (NMI)<sup>39</sup> are used for image quality assessment. We further apply our trained model to the BLSA dataset, comprising 1033 subjects, to illustrate our model’s capability in reducing longitudinal variance caused by positional variation. We achieve this by comparing changes in body composition metrics before and after harmonization.

This paper is an extension of our conference version.<sup>40</sup> We focus on improving the generalizability of our model and validate our model’s harmonization capability on the longitudinal single-slice data. The difference can be summarized as follows:

- We revisit the target slice selection method and propose a semi-BPR method that improves the structural similarity of the selected target slice.
- We collect an in-house non-contrast phase CT dataset and validate model performance on different contrast phases to minimize the domain shift problem mentioned in the limitations section in the conference version.
- More metrics are introduced to evaluate our model and generated images in both 3D datasets and 2D single-slice dataset.
- We conduct a comprehensive longitudinal evaluation on the BLSA single-slice dataset with 1033 subjects, which is a significant increase compared to the 20 subjects evaluated in the conference version.
- We conduct an ablation study on validating the most effective distance range between the given and target slice for our proposed method.

Our contributions in this work can be summarized as follows. We present C-SliceGen, a VAEGAN-based generative model for generating subject-specific abdominal slices at predefined vertebral levels. Using an arbitrary axial slice as input, C-SliceGen can implicitly incorporate unknown target slices during testing, producing realistic and structurally similar images. Our experiments demonstrate that the proposed method can harmonize variance in body composition metrics caused by positional variation in the longitudinal setting, facilitating accurate longitudinal analysis.

## 2 Method

### 2.1 Technical Background

Here, we provide a brief review of generative models (VAE and GAN) and their variants, upon which our C-SliceGen is based. While generative models have proven effective in sample generation, it is important to note that no existing method can be directly applied to our specific task.

#### 2.1.1 VAE

VAEs can be expressed probabilistically as  $P(x) = P(z)P(x|z)$ , where  $x$  represents input images and  $z$  represents latent variables. The goal of these models is to maximize the likelihood of  $p(x) = \int p(z)p_{\theta}(x|z)dz$ , where  $z \sim N(0, 1)$  is the prior distribution,  $p_{\theta}(x|z)dz$  is the

posterior distribution, and  $\theta$  represents the decoder parameters. However, it is not feasible to find decoder parameters  $\theta$  that maximize the log-likelihood. Instead, VAEs optimize encoder parameters  $\phi$  by estimating  $p_\theta(x|z)$  using  $q_\phi(z|x)$ , which is assumed to be a Gaussian distribution with  $\mu$  and  $\sigma$  as the outputs of the encoder. VAEs are trained by optimizing the evidence lower bound (ELBO).

$$L_{\text{VAE}}(\theta, \phi, x) = E[\log p_\theta(x|z)] - D_{\text{KL}}[q_\phi(z|x) \| p_\theta(z)], \quad (1)$$

where  $E[\log p_\theta(x|z)]$  represents the reconstruction loss and  $D_{\text{KL}}[q_\phi(z|x) \| p_\theta(z)]$  represents the KL-divergence, which facilitate the posterior distribution to be close to the prior distribution  $p_\theta(z)$ . During testing, new data can be generated by sampling from the normal distribution  $z \sim N(0, 1)$  and inputting it into the decoder. Conditional VAE can be optimized with the following ELBO equation as well with little modification:

$$L_{\text{VAE}}(\theta, \phi, x, c) = E[\log p_\theta(x|z, c)] - D_{\text{KL}}[q_\phi(z|x, c) \| p_\theta(z|c)]. \quad (2)$$

### 2.1.2 GAN

GANs consist of two parts: discriminator and generator. Suppose we have input noise variables  $p_z(z)$ , the generator will map the input noise to data space  $G(z)$  and mix it with the real data  $x$ . The discriminator  $D$ , on the other hand, transforms image data into a probability indicating whether the image belongs to the real data distribution or the generator distribution.<sup>41</sup> To be more specific, the discriminator and the generator play the two-player minmax game with value function  $V(D, G)$  in the following manner:<sup>42</sup>

$$\min_G \max_D V(D, G) = \mathbb{E}_{x \sim p_{\text{data}}(x)}[\log D(x)] + \mathbb{E}_{z \sim p_z(z)}[1 - \log D(G(z))]. \quad (3)$$

The Wasserstein GAN with gradient penalty (WGAN-GP)<sup>43</sup> is an alternative to traditional GAN that enhances the stability of the model during training and addresses problems such as model collapse. The loss function of WGAN-GP can be written as

$$L_{\text{WGAN-GP}} = \mathbb{E}_{\tilde{x} \sim \mathbb{P}_g}[D(\tilde{x})] + \mathbb{E}_{x \sim \mathbb{P}_r}[D(x)] + \lambda \mathbb{E}_{\hat{x} \sim \mathbb{P}_{\hat{x}}}[(\|\nabla_{\hat{x}} D(\hat{x})\|_2 - 1)^2], \quad (4)$$

where  $\mathbb{P}_g$ ,  $\mathbb{P}_r$ , and  $\mathbb{P}_{\hat{x}}$  represent the generator distribution, data distribution, and random sample distribution, respectively.

## 2.2 C-SliceGen

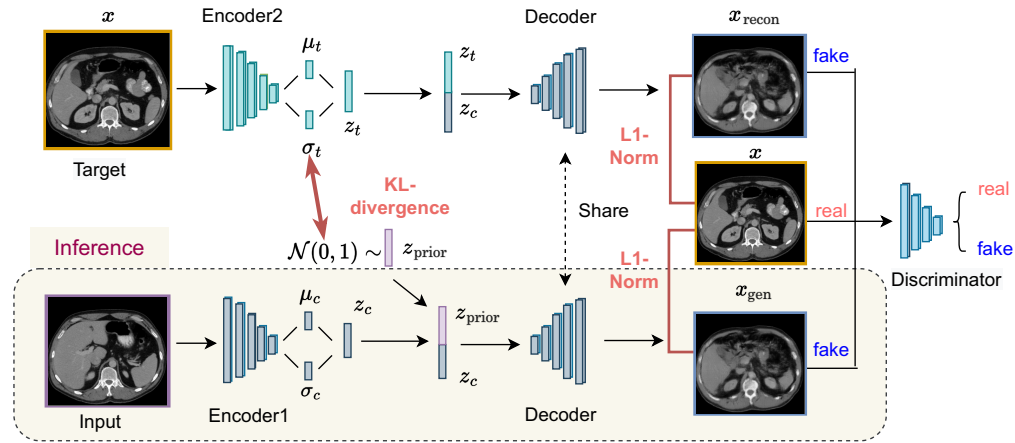
Our task involves generating a new slice at a pre-defined vertebral level from an arbitrary slice that is obtained at any vertebral level within the abdominal region. Our proposed method is shown in Fig. 2, which comprises two encoders, one decoder, and one discriminator.

### 2.3 Target Slice Selection

The first step is to select a target slice for each individual, with the criterion of selecting a slice that is most similar in terms of organ/tissue structure and appearance across all subjects. Choosing a comparable target slice for each individual is a challenging task as it involves taking into account subject-specific variations in organ structure and body composition. We use two methods to select similar target slices for each subject:

**BPR-based method:** We select slices with similar body part regression (BPR) score<sup>44</sup> as the target slices across subjects. BPR gives different scores to different slices in the abdominal region and is efficient in locating slices. We first select a target slice in a reference subject and document its BPR score, and then we select the slices that have the most similar BPR score as the target slices across subjects.

**Registration-based method:** Initially, we select a reference subject's slice as the reference target slice. Subsequently, we register the axial slices of every subject's volume to the reference target slice and identify the slice that has the largest NML score as the subject target slice.



**Fig. 2** The input image is an arbitrarily acquired slice in the abdominal region. During the training phase, target images ( $x$ ) are used as the GT for the generation and reconstruction process. Latent variables, such as  $z_c$ ,  $z_t$ , and  $z_{\text{prior}}$ , are derived from conditional images, target images, and the normal Gaussian distribution, respectively.  $x_{\text{gen}}$  and  $x_{\text{recon}}$  are considered fake images and target images are considered real images for the discriminator.

### 2.3.1 Training

The input image for the model is the arbitrary slice, which provides subject-specific information, including organ shape and tissue localization. Note the assumption that the input is intended to represent the target which is not a random input. We believe that this information remains interpretable after encoding to latent variables  $z_c$  by encoder1. All selected target slices ( $x$ ) should have similar organ/tissue structures and appearances. This information is encoded in the latent variables  $z_t$ . The distribution of  $z_t$  can be expressed as  $q_{\phi}(z_t|x)$ , where  $\phi$  denotes the encoder2 parameters. We combine the organ/tissue structure and appearance of the target slice with the subject-specific information by concatenating the latent variables  $z_c$  and  $z_t$ . This combination facilitates the decoder to reconstruct the target slice for the given individual. To regularize the reconstruction process, we compute the L1-norm between the target slice ( $x$ ) and the reconstructed slice ( $x_{\text{recon}}$ ) using the following equation:

$$L_{\text{recon}} = \|x - x_{\text{recon}}\|. \quad (5)$$

Since no target slice is available during testing, we follow the similar approach as in VAEs. We assume that  $q_{\phi}(z_t|x)$  is a Gaussian distribution with parameters  $\mu_t$  and  $\sigma_t$ , which are the outputs of encoder2. We optimize the KL-divergence to encourage  $q_{\phi}(z_t|x)$  to be close to the prior distribution  $z_{\text{prior}} \sim N(0, 1)$ , which can be written as

$$L_{\text{KL}} = \frac{1}{2} \sum_{k=1}^K (1 + \log(\sigma_k^2) - \mu_k^2 - \sigma_k^2), \quad (6)$$

where  $K$  represents the dimension of the latent space. To mimic the process of image generation in the testing phase, we added another input to the decoder by concatenating  $z_c$  with  $z_{\text{prior}}$  for target slices generation. We denote these generated images as  $x_{\text{gen}}$ .  $x_{\text{gen}}$  is also regularized by L1-Norm with the equation:

$$L_{\text{gen}} = \|x - x_{\text{gen}}\|. \quad (7)$$

The combined loss function of the above-mentioned steps can be expressed as

$$L_{\text{cVAE}} = L_{\text{recon}} + L_{\text{gen}} + L_{\text{KL}}. \quad (8)$$

However, the major drawback of VAEs is that they tend to generate blurry images. On the other hand, GANs can produce images with sharp edges. Following Ref. 29, we add GAN regularization into our model. The discriminator in our proposed C-SliceGen model classifies both the generated and reconstructed images as fake images, while the target images are considered

real images. The decoder acts as the generator for the GAN part. GAN loss adds another constraint to force the generated and target images to be similar. The total loss function for our C-SliceGen model can be written as follows:

$$L = L_{cVAE} + \beta L_{GAN}. \quad (9)$$

The adversarial regularization is adjustable by the weighting factor  $\beta$ .

The weights of the conditional VAE and discriminator are alternately updated. The conditional VAE is updated using the loss in Eq. (8). Subsequently, the conditional VAE transitions to inference mode and its outputs  $x_{\text{gen}}$  and  $x_{\text{recon}}$  are utilized to update the discriminator parameters. The discriminator is then updated based on Eq. (4).

### 2.3.2 Testing

During testing, encoded conditional image  $z_c$  is combined with  $z_{\text{prior}}$ , which is sampled from a normal Gaussian distribution, and then fed into the decoder to generate the target slice.

## 3 Implementation Details

### 3.1 Dataset

We train and evaluate our methods on 3D volumetric CT datasets in both the portal venous phase and the non-contrast phase as well as 2D single-slice CT dataset in the non-contrast phase.

#### 3.1.1 In-house portal venous dataset

This dataset contains 1120 3D portal venous CT volumes from 1120 de-identified subjects from Vanderbilt University Medical Center (VUMC). The data have been approved by the Institutional Review Board (IRB) with IRB #160764. A quality check is performed on every CT scan to ensure normal abdominal anatomy. The dataset is divided into training, validation, and testing with 1029, 8, and 83 subjects, respectively.

#### 3.1.2 In-house non-contrast dataset

To minimize the domain shift problem when applying models trained with the portal venous dataset to non-contrast single-slice data, we further train and validate our method using a 3D non-contrast CT dataset with IRB #172167. This dataset contains 1488 subjects. We split the dataset into training, validation, and testing with 1059, 117, and 312 subjects, respectively.

#### 3.1.3 BTCV dataset

The MICCAI 2015 Multi-Atlas Abdomen Labeling Challenge (BTCV) dataset consisting of 30 portal venous CT volumes for training and 20 for testing is used for the evaluations. We finetuned the model trained with the in-house portal venous dataset with 22 data for training and 8 for validation.

#### 3.1.4 BLSA dataset

We assess the effectiveness of our method in harmonizing the positional variation on single-slice data with the BLSA dataset. To minimize the radiation exposure during longitudinal imaging, the BLSA CT protocol captures single-slice data at specific anatomical landmarks instead of acquiring 3D CT data as is typically done in clinical settings. A total of 1033 subjects have more than one visit with some subjects having up to 12 visits and the median number of scans being three for the past 15 years. The total number of CT axial scans is 4223, and all the scans are in the non-contrast phase.

### 3.2 Metrics

We perform a quantitative evaluation of our C-SliceGen generative models using different target slice selection approaches and varying values of  $\beta$  [as defined in Eq. (9)].

#### 3.2.1 Metrics for 3D datasets

For the 3D volumetric dataset where the target slices are available for comparison, we use four metrics to assess image quality: SSIM, PSNR, LPIPS, and NML. SSIM assesses the image based on three factors: luminance, contrast, and structure, and the final score is derived from the multiplication of those three independent factors. PSNR is most determined by mean squared error. LPIPS is utilized to assess the perceptual similarity of two images by calculating the similarity between the activations of their respective patches through a pre-defined network. This measure is highly correlated with human perception. NML measures the degree of information present in one image that is contained in the other image.<sup>45</sup>

#### 3.2.2 Metrics for 2D single-slice dataset

In the BLSA dataset, each subject only has one axial abdominal CT scan taken per visit, resulting in the absence of GT for direct comparison with the generated target slices. We use NML and coefficient of variation (CV) to evaluate the model performance on harmonizing the positional variation. CV indicates the amount of differences between scans,<sup>46,7</sup> which is defined as

$$CV = \frac{\sigma}{\mu}. \quad (10)$$

### 3.3 Training and Testing

BPR<sup>44</sup> is used to ensure a consistent field of view in the abdominal region for all the 3D volumetric data. We preprocess the data with a soft-tissue CT window range of  $[-125, 275]$  Hounsfield units (HU) and further rescale the data to the range of 0 to 1. The 2D axial CT scans are resized from size  $512 \times 512$  to  $256 \times 256$  before being fed into the models. Pytorch is used to implement the proposed methods, with the Adam optimizer and a learning rate of  $1e-4$ , and a weight decay of  $1e-4$  to optimize the network's total loss when training the model from scratch. When finetuning the BTCV dataset, the learning rate is reduced to  $1e-5$ . The encoder, decoder, and discriminator structures are modified based on Refs. 47 and 48. We adopt common data augmentation methods, such as shift, rotation, and flip, with a probability of 0.5 to facilitate training.

## 4 Results

### 4.1 3D Datasets Evaluation

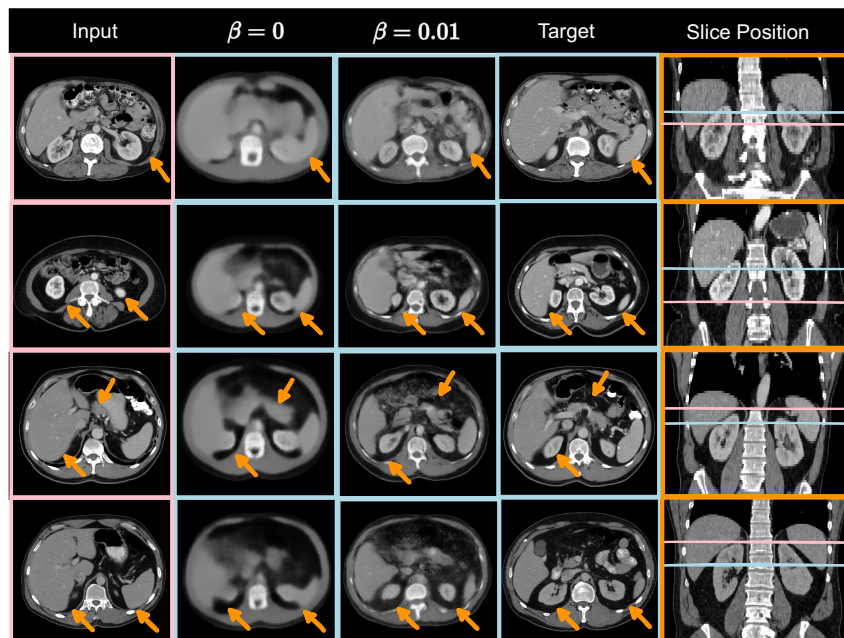
We present the quantitative performance of our model with various metrics on different datasets in Table 1. Comparing with target slices selected by the registration-based method and BPR-based method, the registration-based method achieves better performance on the in-house portal venous dataset while the BPR-based method performs slightly better on the BTCV dataset which might indicate that the BPR-based method and registration-based method have comparable performance on selecting target slices on the portal venous phase CT scans. We show qualitative results on the BTCV test set in Fig. 3, which demonstrates that our model is capable of generating target slices irrespective of whether the conditional slice is at a higher, lower, or similar vertebral level.

On the non-contrast dataset, however, our empirical results indicate that the generated images are not on a similar vertebral level as opposed to what they are supposed to be. We trace back the reason and find that the selected target slices in the non-contrast dataset are not on a similar vertebral level initially, making the network hard to learn the target location and resulting in significant noise in the generated images. To address this issue, we use a semi-BPR method wherein we compare the target slice with the eight axial slices preceding and succeeding it to select the new target slice. The results with the manually corrected target slices are presented in

**Table 1** Quantitative results on two in-house test sets and BTCV test set using different target slice selection methods with different  $\beta$  in Eq. (9) for training.

Method	SSIM $\uparrow$	PSNR $\uparrow$	LPIPS $\downarrow$	NML $\uparrow$
The in-house portal venous dataset				
$\beta = 0$ , registration	<b>0.636</b>	<b>17.634</b>	0.361	0.344
$\beta = 0$ , BPR	0.618	16.470	0.381	0.316
$\beta = 0.01$ , registration	0.615	17.256	<b>0.209</b>	<b>0.377</b>
$\beta = 0.01$ , BPR	0.600	16.117	0.226	0.364
The BTCV dataset				
$\beta = 0$ , registration	0.603	17.367	0.362	0.340
$\beta = 0$ , BPR	<b>0.605</b>	<b>17.546</b>	0.376	0.330
$\beta = 0.01$ , registration	0.583	16.778	<b>0.208</b>	0.422
$\beta = 0.01$ , BPR	0.588	16.932	0.211	<b>0.423</b>
The in-house non-contrast dataset				
$\beta = 0.01$ , semi-BPR	0.570	18.312	0.210	0.405

Note: bold values indicate the best in each column/dataset.



**Fig. 3** The image enclosed within a pink bounding box depicts the input slice, while the image enclosed within a light blue bounding box represents the model outputs and target slice from four different subjects in the BTCV test set. The pink and light blue lines in the rightmost column indicate the axial position of the input and target slices, respectively. The structural differences in organs between the input and target slices are highlighted by orange arrows. The results indicate our model can implicitly learn the subject-specific target slices and generate realistic and structurally similar slices given input slices from arbitrary vertebral levels.



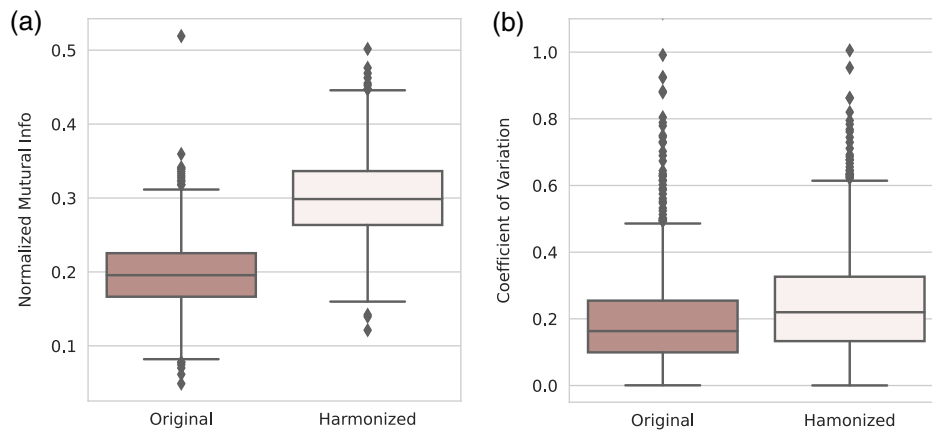
Table 1. Comparing the results from the non-contrast phase dataset and the portal venous phase dataset, two out of four metrics show slightly worse performance while the other two metrics show comparable or even better performance. This may indicate our model has good generalizability on different CT phases.

#### 4.2 2D Single-Slice Evaluation

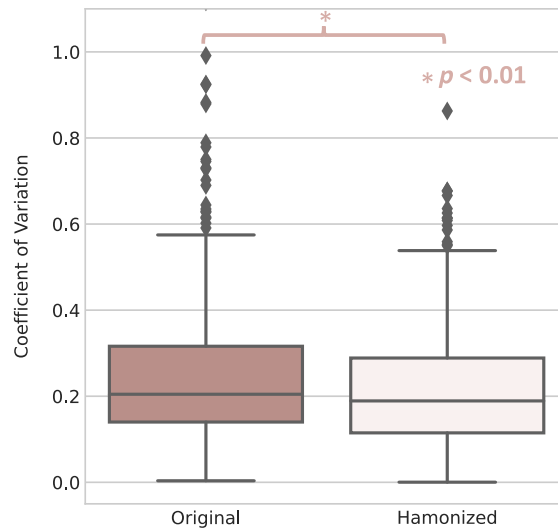
We use NML and CV to evaluate our model's longitudinal variation harmonization capability. We evaluate the model's performance using the visceral fat area as the primary metric. This is because visceral fat is highly susceptible to positional variation as mentioned in Ref. 7, and it is a crucial component of body composition, which indicates an individual's health condition.<sup>7,49</sup>

The BLSA single-slices CT scans are fed into the model trained with non-contrast 3D CT volumes. The generated images are resized to the original image size of  $512 \times 512$ . We use the method in Ref. 7 to extract the segmentation mask of the visceral fat which includes feeding the data in a pre-trained model for inner/outer abdominal wall segmentation and using fuzzy c-means<sup>50,51</sup> to extract the adipose tissues. In the inner abdominal wall segmentation, we observe that the model performs unsatisfactory to exclude the retroperitoneum from both real and generated images. The retroperitoneum is an anatomical region situated behind the abdominal cavity, which comprises the aorta, and left and right kidneys, and often lacks well-defined boundaries, making it difficult to segment accurately. We follow the practice in Ref. 7 and manually assess the results from both the real and generated images to ensure that the retroperitoneum is segmented correctly. We mask the inner abdominal wall with the adipose tissues to get the final visceral fat mask.

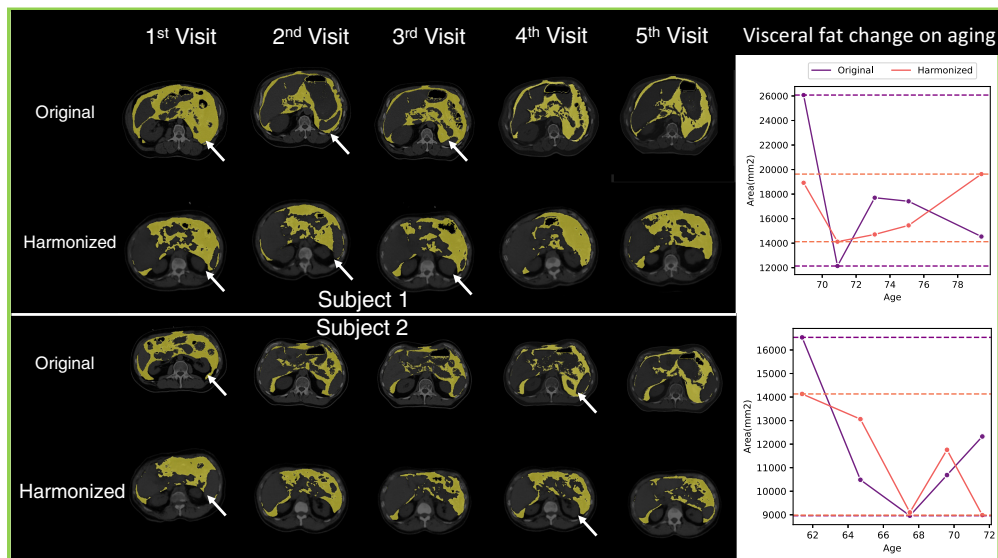
We calculate the NML on every two scans of the same subjects on both original and harmonized images; and the results are shown in Fig. 4(a). According to the result, the harmonized images have higher NML compared with the original images which indicates the generated images of the same subjects share a higher degree of similar information compared to those of the original images. We further evaluate with the CV, while we observe higher CV in the harmonized images compared with the original images, as shown in Fig. 4(b), which implies that differences between slices are increased after generation. As the metrics show contradictory results, we conduct a human assessment of the harmonization results. We find that 431 out of 1033 subjects have at least one scan that is taken in obviously different vertebral levels compared with the other scans. For those 431 subjects, our model can help harmonize the positional variation resulting in a significantly lower CV than the original images with  $p < 0.01$  under the Wilcoxon signed-rank test, as shown in Fig. 5. Our models can effectively reduce the positional



**Fig. 4** Quantitative results of applying the trained model on 1033 subjects from the BLSA dataset. (a) The results with NML as metrics, and (b) the results with CV as metrics. We observed higher NML and CV among the harmonized images. NML and CV show contradictory results where higher NML suggests that the generated images of the same subjects share a higher degree of similar information compared to those of the original images while higher CV implies the differences between slices are increased after harmonization.

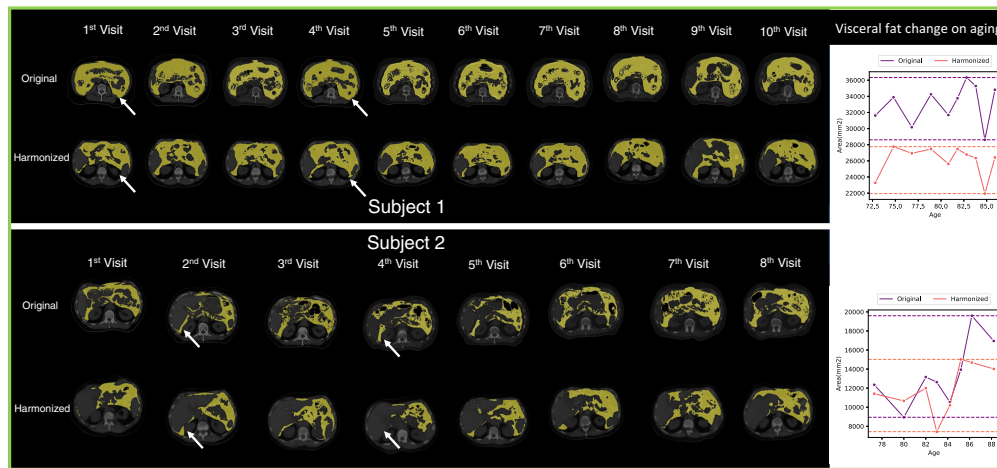


**Fig. 5** CV result of applying the trained model on subjects that have at least one scan that is taken in obvious different vertebral levels. Note: \* represents statistically significant ( $p < 0.01$ ) by Wilcoxon signed-rank test. The result demonstrates that our model is effective in harmonizing the variance caused by positional variation in longitudinal imaging.

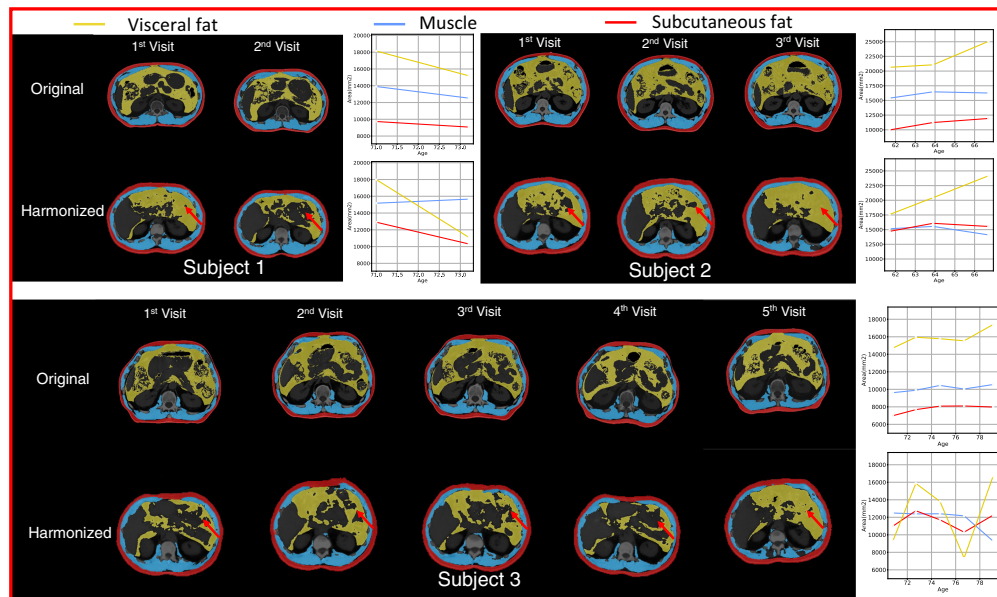


**Fig. 6** The yellow mask represents the visceral fat for two subjects in the BLSA dataset that have five repeated visits. The rightmost column shows the visceral fat area change on aging. The horizontal lines represent the maximum and minimum values of different image types. We highlight the images with different organs or tissues captured compared to most other images by the white arrows. The generated images manage to harmonize organ/tissue differences. We observed the range between the maximum and minimum values of the among different visits decreased after harmonization which implies the harmonized images have less variation.

variance in subjects with both a lesser and greater number of longitudinal visits, as shown in Figs. 6 and 7, respectively. In Fig. 6, our model reduces the variance by 36.3% and 42.5%, respectively. And in Fig. 7, the variance is reduced by 37.8% and 76.9%, respectively. However, for those subjects whose original slice was already taken at a similar vertebral level, our model can introduce additional noise and result in larger variance among scans, as it shown in Fig. 8.



**Fig. 7** Two subjects in the BLSA dataset that have 8 or 10 repeated visits. The yellow mask represents visceral fat. The rightmost column shows the visceral fat area change on aging. Similar to Fig. 6, we highlight the image that is captured in different vertebral levels with white arrows. Subjects with a larger amount of repeated visits are more likely to have images captured at various vertebral levels.



**Fig. 8** Visualization of the original and harmonized images on the subjects that have scans taken at a similar vertebral level. Yellow, blue, and red masks represent visceral fat, muscle, and subcutaneous fat, respectively, which are generated by automatic segmentation methods. When the input images are at a similar vertebral level already, our model can introduce additional noise by predicting different heterogeneous soft tissues, as highlighted with the red arrows.

### 4.3 Ablation Study

#### 4.3.1 Adversarial regularization

We compare our models trained with different  $\beta$  scores to evaluate the impact of adversarial regularization. As it is shown in Fig. 3, when  $\beta = 0.01$ , the generated images are realistic and similar to the target slices. When  $\beta = 0$ , there is no adversarial regularization, resulting in blurry generated images. By comparing the results with  $\beta = 0$  and  $\beta = 0.01$  in Fig. 3, it can be inferred that the adversarial regularization greatly improves the quality of the generated images.

This human qualitative assessment is aligned with LPIPS and NML results, as it is shown in Table 1. However, it is different from the observation of the SSIM and the PSNR score, which are

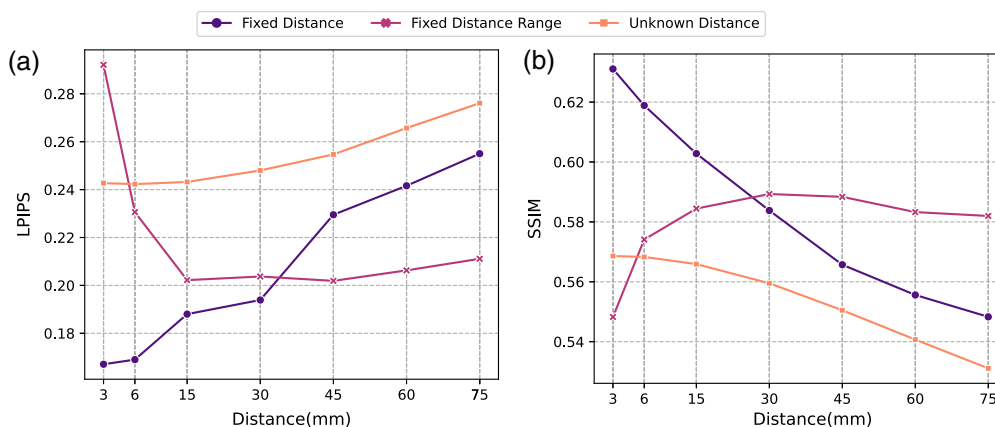
higher with  $\beta = 0$ . This observation supports that SSIM and PSNR scores may not completely reflect human perception, as mentioned in Refs. 52 and 53.

### 4.3.2 Distance impact

In our method, we aim to map an abdominal axial scan at an arbitrary vertebral level to a pre-defined target vertebral level, where the distance between the given scan and the target scan is unknown. This is also the case in the BLSA single-slice dataset. We assume that as the distance difference between the scans increases, the scans will undergo more structural changes, making the generation process more challenging. To evaluate the impact of distance on our model performance, we conduct validation experiments. Specifically, we assess the performance of models trained on scans from varying distance ranges and compare them to models trained using known distances. We also include the model trained with the abdominal region with unknown distance for comparison (model in Table 1). All the models are trained and tested with  $\beta = 0.01$ , using the in-house portal venous dataset and BPR-based target slice selection method. We evaluate the model performance with the LPIPS and SSIM scores. The results are shown in Fig. 9.

To assess the performance of the model with fixed known spacing, we do not ask the model to predict the target slice since in most of the CT volumes, the spacing in the  $z$  dimension is 3 mm. In this case, if we train with a fixed distance of 3 mm, we can only get two conditional and target slice pairs for each subject, which leads to a data deficiency problem and cannot evaluate the model performance properly. Therefore, instead of predicting the target slice, we design the system to generate corresponding slices at intervals of 3, 6, and up to 75 mm upward in the abdominal region for each given slice. For the model trained with a fixed spacing range—3, 6, and up to 75 mm—the models are trained with slices up to the specific distance from the target slice, respectively. The line unknown distance refers to the results by applying the trained model in Table 1 to a fixed distance test set.

According to Fig. 9, when the model has a fixed distance range of 3 and 6 mm, it has the worst performance among the other distance ranges. This can be explained by the data deficiency problem as mentioned before. The performance of models trained with a fixed distance is optimal when the slice distance is small but gradually drops with an increase in distance. This indicates that the model performs better with a smaller distance between the given and target slices. The models trained using the entire abdominal region (with unknown spacing) consistently performed poorly, starting from a distance of 6 mm. On the other hand, the models trained



**Fig. 9** Assessing the impact of the distance between the given slice and target slices on the model performance. Fixed distance represents the model trained with a known distance, and fixed distance range line represents the model trained with up to a distance of the given point. Unknown distance represents the inference performance of the model trained on the entire abdominal region when tested on data with a fixed distance between the input and target slices. We observe that the model's most effective distance range is between 15 and 60 mm where performance is stable in terms of both LPIPS and SSIM.

using a fixed range of 15 to 60 mm showed stable performance in terms of LPIPS and SSIM. These results indicate that the model's most effective distance range is between 15 and 60 mm, and training with a wider range can lead to decreased overall performance.

## 5 Limitations and Discussion

In this work, we improve the domain shift problem we observed in our previous publication<sup>40</sup> that when applying the model trained with the portal venous phase to the non-contrast phase BLSA data, the model has limited performance. We manage to reduce this issue by using non-contrast 3D volumetric data for training together with semi-BPR-based target slice selection for accurate target slice selection. From Figs. 4 and 5, we observe that our model can help reduce the longitudinal variance on data that are taken at obviously different vertebral levels. However, when it is applied to the data that are at similar vertebral levels, our model can introduce additional noise by predicting different heterogeneous soft tissues, as shown in Fig. 8. Predicting and synthesizing heterogeneous soft tissues, such as the colon and stomach, is challenging because these tissues' size and shape are largely dependent on individual conditions and position at the time of the CT scan, making it hard for the generative model to find a subject-specific distribution of such tissues. Hence, this remains a critical limitation of this study. Solving the heterogeneous soft tissue, shape, and size generation problem can be the future work direction. Exploring solutions for the generation of heterogeneous soft tissue and preserving shape and boundary information can also be the future work direction.

In addition, we validate our method's most effective distance range. Models trained with data that are up to 60 mm away from the target slice have comparable performance to that of the model trained with up to 15 mm away from the target slice. Furthermore, the model trained within the 60 mm range performs markedly better than the models presented in Table 1, which is trained using slices from the entire abdominal region. Therefore, for the model to be most effective, it would be preferable to collect data within a range of no more than 60 mm, or roughly around  $\pm 3$  vertebral level<sup>54</sup> in future data collection.

## 6 Conclusion

Herein, we present our C-SliceGen model, which utilizes an arbitrary 2D axial abdominal CT slice as input and generates a subject-specific slice at a desired vertebral level. Our model can effectively capture changes in the organs across different vertebral levels and generate images that are realistic and structurally similar. In addition, we demonstrate our model's effectiveness in harmonizing longitudinal body composition variance caused by positional differences among different visits in the BLSA single-slice CT dataset. Specifically, in subjects with scans taken at different vertebral levels, our model effectively harmonizes positional variation, resulting in a significantly lower CV compared to the original images ( $p < 0.01$ , Wilcoxon signed-rank test). Overall, this approach offers a promising solution for managing imperfect single-slice CT abdominal data in longitudinal analysis.

---

### Disclosures

The authors of the paper are directly employed by institutes or companies provided in this paper. No conflicts of interest exist in the submission of this paper. During the preparation of this work, the authors used ChatGPT to improve readability and language. After using this tool, the authors reviewed and edited the content as needed and take full responsibility for the content of the publication.

### Code and Data Availability

The in-house and BLSA data are unavailable to the public due to the sensitive nature of the research. The BTCV data are publicly online at BTCV website: <https://www.synapse.org/#/Synapse:syn3193805/wiki/217789>.

## Acknowledgments

This research was supported by the National Science Foundation (NSF) through CAREER award numbers 1452485 and 2040462 and the National Institutes of Health (NIH) under award numbers R01EB017230, R01EB006136, R01NS09529, T32EB001628, NIH R01DK135597 (Huo), 5UL1TR002243-04, 1R01MH121620-01, and T32GM007347; by ViSE/VICTR VR3029; and by the National Center for Research Resources (Grant No. UL1RR024975-01) and is now at the National Center for Advancing Translational Sciences (Grant No. 2UL1TR000445-06). This research was conducted with the support from the Intramural Research Program of the National Institute on Aging of the NIH. The content is solely the responsibility of the authors and does not necessarily represent the official views of the NIH. The identified datasets used for the analysis described were obtained from the Research Derivative (RD), database of clinical and related data. The in-house imaging dataset(s) used for the analysis described were obtained from ImageVU, a research repository of medical imaging data and image-related metadata. ImageVU and RD are supported by the VICTR CTSA award (ULTR000445 from NCATS/NIH) and Vanderbilt University Medical Center institutional funding. ImageVU pilot work was also funded by PCORI (Contract No. CDRN-1306-04869).

## References

1. R. Kuriyan, "Body composition techniques," *Indian J. Med. Res.* **148**(5), 648 (2018).
2. A. De Schutter et al., "Body composition in coronary heart disease: how does body mass index correlate with body fatness?" *Ochsner J.* **11**(3), 220–225 (2011).
3. S. M. Ribeiro and J. J. Kehayias, "Sarcopenia and the analysis of body composition," *Adv. Nutr.* **5**(3), 260–267 (2014).
4. J. D. Solanki et al., "Body composition in type 2 diabetes: change in quality and not just quantity that matters," *Int. J. Prev. Med.* **6**, 122 (2015).
5. A. Andreoli et al., "Body composition in clinical practice," *Eur. J. Radiol.* **85**(8), 1461–1468 (2016).
6. L. Ferrucci, "The Baltimore longitudinal study of aging (BLSA): a 50-year-long journey and plans for the future," *J. Gerontol. Series A: Biol. Sci. Med. Sci.* **63**(12), 1416–1419 (2008).
7. X. Yu et al., "Longitudinal variability analysis on low-dose abdominal CT with deep learning-based segmentation," *Proc. SPIE* **12464**, 1246423 (2023).
8. Q. Yang et al., "Label efficient segmentation of single slice thigh ct with two-stage pseudo labels," *J. Med. Imaging* **9**(5), 052405–052405 (2022).
9. Q. Yang et al., "Single slice thigh CT muscle group segmentation with domain adaptation and self-training," *J. Med. Imaging* **10**(4), 044001–044001 (2023).
10. A. Z. Moore et al., "Correlates of life course physical activity in participants of the baltimore longitudinal study of aging," in *Aging Cell*, p. e14078 (2024).
11. Q. Yang et al., "Quantification of muscle, bones, and fat on single slice thigh CT," *Proc. SPIE* **12032**, 422–429 (2022).
12. X. Yu et al., "Accelerating 2D abdominal organ segmentation with active learning," *Proc. SPIE* **12032**, 893–899 (2022).
13. F. Hu et al., "Image harmonization: a review of statistical and deep learning methods for removing batch effects and evaluation metrics for effective harmonization," *NeuroImage* **274**, 120125 (2023).
14. J.-P. Fortin et al., "Harmonization of multi-site diffusion tensor imaging data," *Neuroimage* **161**, 149–170 (2017).
15. H. Horng et al., "Generalized combat harmonization methods for radiomic features with multi-modal distributions and multiple batch effects," *Sci. Rep.* **12**(1), 4493 (2022).
16. R. Pomponio et al., "Harmonization of large MRI datasets for the analysis of brain imaging patterns throughout the lifespan," *NeuroImage* **208**, 116450 (2020).
17. J. C. Beer et al., "Longitudinal ComBat: a method for harmonizing longitudinal multi-scanner imaging data," *NeuroImage* **220**, 117129 (2020).
18. A. A. Chen et al., "Mitigating site effects in covariance for machine learning in neuroimaging data," *Hum. Brain Mapp.* **43**(4), 1179–1195 (2022).
19. A. Avalos-Pacheco, D. Rossell, and R. S. Savage, "Heterogeneous large datasets integration using Bayesian factor regression," *Bayesian Anal.* **17**(1), 33–66 (2022).
20. I. Goodfellow et al., "Generative adversarial nets," in *Adv. Neural Inf. Process. Syst.*, Z. Ghahramani et al., Eds., Vol. 27, Curran Associates, Inc. (2014).
21. T. Luan et al., "High fidelity 3D hand shape reconstruction via scalable graph frequency decomposition," in *Proc. IEEE/CVF Conf. Comput. Vision and Pattern Recognit.*, pp. 16795–16804 (2023).
22. A. Q. Nichol and P. Dhariwal, "Improved denoising diffusion probabilistic models," in *Int. Conf. Mach. Learn.*, pp. 8162–8171 (2021).

23. T. Luan et al., “Pc-hmr: Pose calibration for 3D human mesh recovery from 2D images/videos,” in *Proc. AAAI Conf. Artif. Intell.*, Vol. 35, pp. 2269–2276 (2021).
24. S. Pan et al., “2D medical image synthesis using transformer-based denoising diffusion probabilistic model,” *Physi. Med. Biol.* **68**(10), 105004 (2023).
25. S. Pan et al., “Synthetic CT generation from MRI using 3D transformer-based denoising diffusion model,” *Med. Phys.* (2022).
26. L. Shen et al., “Patient-specific reconstruction of volumetric computed tomography images from a single projection view via deep learning,” *Nat. Biomed. Eng.* **3**(11), 880–888 (2019).
27. S. Bond-Taylor et al., “Deep generative modelling: a comparative review of VAEs, GANs, normalizing flows, energy-based and autoregressive models,” *IEEE Trans. Pattern Anal. Mach. Intell.* **44**, 7327–7347 (2021).
28. D. P. Kingma and M. Welling, “Auto-encoding variational Bayes,” arXiv:1312.6114 (2013).
29. A. B. L. Larsen et al., “Autoencoding beyond pixels using a learned similarity metric,” in *Int. Conf. Mach. Learn.*, PMLR, pp. 1558–1566 (2016).
30. M. Mirza and S. Osindero, “Conditional generative adversarial nets,” arXiv:1411.1784 (2014).
31. K. Sohn, H. Lee, and X. Yan, “Learning structured output representation using deep conditional generative models,” in *Adv. Neural Inf. Process. Syst.* **28** (2015).
32. R. De Bem et al., “A conditional deep generative model of people in natural images,” in *IEEE Winter Conf. Appl. of Comput. Vision (WACV)*, IEEE, pp. 1449–1458 (2019).
33. P. Henderson, C. H. Lampert, and B. Bickel, “Unsupervised video prediction from a single frame by estimating 3D dynamic scene structure,” arXiv:2106.09051 (2021).
34. Y. Tang et al., “Pancreas CT segmentation by predictive phenotyping,” *Lect. Notes Comput. Sci.* **12901**, 25–35 (2021).
35. B. Landman et al., “MICCAI multi-atlas labeling beyond the cranial vault—workshop and challenge,” in *Proc. MICCAI Multi-Atlas Labeling Beyond Cranial Vault-Workshop Challenge*, Vol. 5, p. 12 (2015).
36. Z. Wang et al., “Image quality assessment: from error visibility to structural similarity,” *IEEE Trans. Image Process.* **13**(4), 600–612 (2004).
37. A. Hore and D. Ziou, “Image quality metrics: PSNR vs. SSIM,” in *20th Int. Conf. Pattern Recognit.*, IEEE, pp. 2366–2369 (2010).
38. R. Zhang et al., “The unreasonable effectiveness of deep features as a perceptual metric,” in *CVPR* (2018).
39. C. E. Shannon, “A mathematical theory of communication,” *Bell Syst. Tech. J.* **27**(3), 379–423 (1948).
40. X. Yu et al., “Reducing positional variance in cross-sectional abdominal CT slices with deep conditional generative models,” *Lect. Notes Comput. Sci.* **13437**, 202–212 (2022).
41. A. Creswell et al., “Generative adversarial networks: an overview,” *IEEE Signal Process Mag.* **35**(1), 53–65 (2018).
42. I. Goodfellow et al., “Generative adversarial networks,” *Commun. ACM* **63**(11), 139–144 (2020).
43. I. Gulrajani et al., “Improved training of Wasserstein GANs,” in *Adv. Neural Inf. Process. Syst.* **30** (2017).
44. Y. Tang et al., “Body part regression with self-supervision,” *IEEE Trans. Med. Imaging* **40**(5), 1499–1507 (2021).
45. F. Maes et al., “Medical image registration using mutual information,” in *Proc. IEEE* **91**(10), 1699–1722 (2015).
46. M. M. J. Wittens et al., “Inter-and intra-scanner variability of automated brain volumetry on three magnetic resonance imaging systems in Alzheimer’s disease and controls,” *Front Aging Neurosci.* **13**, 746982 (2021).
47. R. Gao et al., “Lung cancer risk estimation with incomplete data: a joint missing imputation perspective,” *Lect. Notes Comput. Sci.* **12905**, 647–656 (2021).
48. S. C.-X. Li and B. Marlin, “Learning from irregularly-sampled time series: a missing data perspective,” in *Int. Conf. Mach. Learn.*, PMLR, pp. 5937–5946 (2020).
49. D. L. Duren et al., “Body composition methods: comparisons and interpretation,” *J. Diabetes Sci. Technol.* **2**(6), 1139–1146 (2008).
50. J. C. Bezdek, R. Ehrlich, and W. Full, “FCM: the fuzzy c-means clustering algorithm,” *Comput. Geosci.* **10**(2–3), 191–203 (1984).
51. Y. Tang et al., “Prediction of type II diabetes onset with computed tomography and electronic medical records,” in *Multimodal Learning for Clinical Decision Support and Clinical Image-Based Procedures*, pp. 13–23, Springer (2020).
52. C. Ledig et al., “Photo-realistic single image super-resolution using a generative adversarial network,” in *Proc. IEEE Conf. Comput. Vision and Pattern Recognit.*, pp. 4681–4690 (2017).
53. Y. Almalioglu et al., “EndoL2H: deep super-resolution for capsule endoscopy,” *IEEE Trans. Med. Imaging* **39**(12), 4297–4309 (2020).
54. G. Kayalioglu, “The vertebral column and spinal meninges,” in *The Spinal Cord*, pp. 17–36, Academic Press (2009).

Biographies of the authors are not available.