

LUCIDA: Low-dose Universal-tissue CT Image Domain Adaptation For Medical Segmentation

Yixin Chen¹[0000–0002–2727–6387], Xiangxi Meng², Yan Wang³, Shuang Zeng⁴,
Xi Liu⁵, and Zhaoheng Xie^{1*}

¹ Institute of Medical Technology and National Biomedical Imaging Center, Peking University, Beijing 100191, China 2311110791@stu.pku.edu.cn

² Department of Nuclear medicine, Peking University Cancer Hospital, Beijing 100142, China

³ School of Instrumentation and Optoelectronic Engineering, State Key Laboratory of Software Development Environment, Beihang University, Beijing 102488, China

⁴ Biomedical of Engineering, College of Future Technology, Peking University, Beijing 100191, China

⁵ Central Research Institute, United Imaging Healthcare, Co., Ltd., Shanghai 201807, China

Abstract. Accurate segmentation in low-dose CT scans remains a challenge in medical imaging, primarily due to the high annotation costs. This study introduces LUCIDA, a Low-dose Universal-tissue CT Image Domain Adaptation model. LUCIDA operates under an unsupervised framework, eliminating the need for LDCT annotations. A novelty of LUCIDA lies in its integration of the Weighted Segmentation Reconstruction (WSR) module with a Fourier-based UNet (F-UNet), which not only establishes a linear relationship between prediction maps and ROI-based reconstructed images but also enhances segmentation accuracy through frequency domain adaptation of LDCT images. LUCIDA improves the accuracy of prediction maps, facilitating a new domain adaptation framework. Through extensive evaluation experiments, LUCIDA has demonstrated its effectiveness in accurately recognizing a wide array of tissues, significantly outperforming conventional methods. Additionally, we present the LUCIDA Ensemble model, which achieves performance comparable to supervised learning models in organ segmentation, capable of recognizing up to 112 tissue types.

Keywords: Low-Dose Computed Tomography · Universal Model · Segmentation · Domain Adaptation

1 Introduction

In Positron Emission Tomography/Computed Tomography (PET/CT) and other disease screening applications, accurate segmentation from Low-Dose CT (LDCT) is pivotal for precise disease diagnosis and formulating effective treatment plans.

* Corresponding author: Zhaoheng Xie (xiezhaozheng@pku.edu.cn)

Despite advancements in imaging technology, LDCT images lack a universal segmentation model that can identify a broad spectrum of human tissues. Manual annotation of all tissues in LDCT images requires substantial costs and resources. Therefore, leveraging extensively annotated NDCT datasets (source domain) for application in LDCT (target domain) represents a cost-effective unsupervised domain adaptation (UDA) task.

Existing UDA methods like CycleGAN [25], SIFA [3], FVP [19], and AIGAN [6] have shown promise in UDA and LDCT denoising. However, GAN-based UDA methods [3, 6, 25] often suffer from poor stability, adversely affecting the adaptation performance in high-accuracy segmentation. High-accuracy segmentation has essentially been achieved in NDCT, with existing methods averaging over 90% Dice score.

In response to these challenges, we introduce Low-dose Universal-tissue CT Image Domain Adaptation (LUCIDA). LUCIDA operates under an unsupervised protocol, capable of recognizing 112 tissue types in LDCT without the need for annotations. The core feature of LUCIDA is the Weighted Segmentation Reconstruction (WSR) module, establishing a linear relationship between prediction maps and ROI-based reconstructed images. By enhancing the quality of reconstructed images, LUCIDA improves the accuracy of prediction maps, constituting a novel UDA framework. Specifically, LUCIDA employs a Fourier-based UNet (F-UNet) for frequency domain adaptation of LDCT images, optimizing F-UNet parameters by minimizing the error between reconstructed and input images. Moreover, we supplement the NDCT dataset with seven additional categories: subcutaneous fat, visceral fat, skull, filler, radius, sternum, and other muscles, resulting in a comprehensive dense label dataset. The training code is already available publicly, and both the well-trained checkpoint and enriched annotations will be released soon.

2 Related Works

2.1 Low-dose CT Denoising

A typical drawback of low-dose CT is image noise due to a weaker X-ray flux. Numerous denoising methods have been developed to address this ill-posed problem and improve low-dose CT image quality. The objective of denoising methods is to align LDCT with NDCT data, which can be seen as a domain adaptation in LDCT. Deep learning (DL) methods have shown superior performance in this task [1, 4, 5, 8–11, 14, 23].

Several DL-based methods adopt mean square error (MSE) loss function to train the model, such as the basic convolutional neural networks [5], UNet-based denoising network [14], residual encoder-decoder convolutional neural networks (RED-CNN) [4], and deep cascade residual networks (DCRN) [9]. Besides, there are many GAN-based denoising methods, such as dual-domain UNet-based GAN [11], Wasserstein GAN [1, 8], and CycleGAN [23]. Attention-encoding Integrated Generative Adversarial Network (AIGAN) [4] achieves the state-of-the-art performance on commonly used metrics (PSNR, SSIM and FSIM) among

UNet [16], CycleGAN [23], CycleWGAN [24], REF-CNN [4]. Moreover, diffusion models have been adopted in LDCT denoising task. Diffusion model overcomes the limitations of GANs by generating more diverse and less blurry images [13]. Dn-Dp [15] leverages diffusion priors and requires only normal-dose CT images for training.

In this paper, we evaluate denoising methods by the performance of downstream task. To the best of our knowledge, there is limited research focusing on the improvement of downstream task performance through domain adaptation in the context of LDCT denoising. This gap underscores the novelty of our approach. This dual focus on denoising and task performance could offer a new direction for future research in medical imaging.

2.2 Domain Adaptation

LDCT denoising task is a subtask of medical Domain Adaptation. In addition to common domain adaptation (DA) methods like CycleGAN [23] and AIGAN [6], other noteworthy approaches in the medical domain include MSGAN [18] and SIFA [3]. MSGAN [18] achieves higher-quality medical image synthesis through dense multi-scale feature connections. SIFA [2] attempts to align feature-level and used the same encoder for image transformation and segmentation tasks, resulting in a synergistic effect between image alignment and feature alignment.

Spectral analysis via Fourier transform plays a crucial role in domain adaptation. Recent studies [21, 22] have highlighted that the amplitude component of an image’s Fourier transform is more reflective of its style, whereas the phase component predominantly carries semantic information. Methods like FDA [22] adjust the amplitude of low-frequency components in source images to match those in the target images.

3 Methodology

Given an NDCT dataset (source domain) (X_s, Y_s) , where X_s and Y_s are the NDCT images and corresponding annotations. The target domain is LDCT dataset without annotations (X_t). We denote the segmentation model trained on NDCT dataset as S , which is commonly referred to as the “source model” in the UDA context.

3.1 Dense Label Annotation And Source Model Training

We employ the TotalSeg dataset [20] as NDCT dataset, which includes 104 annotated tissue masks. However, it does not achieve dense annotation of all tissue structures for each voxel, as shown in Figure 1. In segmentation tasks, “dense labels” refer to annotations for every voxel in the image, assigning each one to a specific class or category.

Expanding upon the Y_s , we annotated labels for subcutaneous fat, visceral fat, skull, filler, radius, sternum, and additional muscle group. We adopt the




Fig. 1. The overview of LUCIDA framework

AI-based pre-segmentation annotation method protocol [20]. Furthermore, we implemented incremental label annotation using a rule-based label fusion.

A preliminary segmentation model was trained after manually annotating expanded labels for the initial 5 cases. Its predictions were merged with the existing labels based on rules, followed by a radiologist’s manual refinement if necessary. The retraining of the segmentation model was conducted following the review and refinement of 5 cases, 20 cases, and 100 cases, respectively.

The rule-based label fusion is defined as:

$$y^*(v) = \begin{cases} y(v) & \text{if } y(v) \neq \text{bg} \\ \hat{y}(v) & \text{if } y(v) = \text{bg} \ \& \ \hat{y}(v) \neq \text{bg}, \\ \text{bg} & \text{if } y(v) = \text{bg} \ \& \ \hat{y}(v) = \text{bg} \end{cases} \quad (1)$$

where y is the original annotation and $\hat{y}$ is the predicted annotation. bg is abbreviation for “background”, and v refers to voxel index within an annotation $y(v)$. According to Equation 1, we combine y and $\hat{y}$ to obtain y^* . The y^* then undergoes manual review and refinement to yield the final annotation result with dense label.

After supplementing the annotation to dense labels, we employ binary cross entropy (BCE) and the Dice Similarity Coefficient (DSC) loss to train the NDCT-Segmentor (source model) S .

3.2 Weighted Segmentation Reconstruction

The LUCIDA framework includes a simple yet intriguing and effective Weighted Segmentation Reconstruction (WSR) module. This is a linear perceptron composed of trainable weights and biases. As shown in Step 2 of Figure 1, the source

model S_s identifies 112 types of tissues in the NDCT input x_s , accomplishing dense segmentation of the CT image and yielding 112 prediction maps.

LUCIDA uses a simple linear mapping to reconstruct the CT image using the prediction maps, achieving a ROI-based reconstruction process that strictly relies on anatomical regions. In this framework, the quality of the reconstructed image highly depends on the prediction maps, in other words, the performance of the segmentation model. Therefore, we could optimize the segmentation performance by improving reconstruction process. This provides a new optimization approach for unsupervised segmentation domain adaptation.

The optimization of WSR is based on mean absolute error (MAE):

$$\arg \min_{w,b} L_{\text{WSR}} = \arg \min_{w,b} |x_s - r_s| = \arg \min_{w,b} |x_s - \sum_{i=0}^{112} (p_i w_i + b_i)|, \quad (2)$$

where r_s represents the reconstructed image of source domain input x_s . p_i represents the prediction map of the i^{th} tissue, and w_i , b_i are linear parameters for p_i . We adopt MAE rather than mean squared error (MSE) because MAE is robust to outliers. Using MAE, WSR could tolerate minor density variations within ROIs without severely penalizing the reconstruction.

3.3 Fourier-based Domain Adaptation

Current research indicates that frequency-based domain adaptation is more effective for CT imaging [19, 21, 22]. For a target image x_t , its discrete Fourier transform is defined as :

$$\mathbb{F}(x_t)(u, q, p) = \sum_{h=0}^{H-1} \sum_{w=0}^{W-1} \sum_{z=0}^{Z-1} x_t(h, w, z) e^{-(j \frac{2\pi u h}{H} + j \frac{2\pi q w}{W} + j \frac{2\pi p z}{Z})} \quad (3)$$

Its amplitude and phase components are $\mathbb{F}_A(x_t)$ and $\mathbb{F}_P(x_t)$, respectively. The inverse Fourier transform is denoted as $\mathbb{F}^{-1}$. The Fourier transformation and inverse can be calculated via the Fast Fourier Transform (FFT). We denote a Residual U-Net as U , whose purpose is to process the amplitude and phase components of a target image to match the pattern of the source domain. The process can be defined as :

$$\mathbb{F}_A^*(x_t), \mathbb{F}_P^*(x_t) = U(\mathbb{F}_A(x_t), \mathbb{F}_P(x_t)), \quad (4)$$

where $\mathbb{F}_A^*(x_t)$ and $\mathbb{F}_P^*(x_t)$ are processed amplitude and phase components. Therefore, the processed image can be defined as $x_t^* = \mathbb{F}^{-1}(\mathbb{F}_A^*(x_t), \mathbb{F}_P^*(x_t))$. The parameters of U are denoted as θ , so we could rewrite the above process, $x_t^* = U(x_t|\theta)$.

Due to the domain shift between the target image x_t and the source domain, performance degradation occurs when the source model is applied to the target domain. This impairment is reflected in the accuracy of the prediction map, which in turn directly impacts the quality of the target reconstructed image r_t .

By enhancing the quality of the reconstructed image, we can mitigate the impact of domain shift on the prediction map, thereby facilitating domain adaptation. We could use the following function to optimize U :

$$\arg \min_{\theta} |x_t - r_t| = \arg \min_{\theta} |x_t - \text{WSR}(S_s(U(x_t|\theta)))| \quad (5)$$

4 Experiments And Results

NDCT source Dataset. The TotalSeg dataset [20] comprises full-dose CT images meticulously chosen from multiple clinical routines, covering 104 distinct anatomical structures. In line with protocols [12, 20], the dataset is divided into 1081 training cases, 57 validation cases, and 65 cases for final testing. In this study, we performed dense label annotation on 1081 training cases, adding seven additional labels: subcutaneous fat, visceral fat, skull, filler, radius, sternum, and other muscles.

LDCT target Dataset. The AutoPET dataset [7], comprising PET/CT data with the CT component being low-dose CT scans. Diagnostic CT scans of the neck, thorax, abdomen, and pelvis were acquired with 90 seconds post-intravenous contrast agent injection. The dataset includes 1,014 training studies from 900 patients and 200 test studies, with PET/CT scans from Siemens Biograph and GE Discovery 690 scanners.

Evaluation Protocol. The LDCT dataset comprises over 1000 low-dose CT images from PET/CT scans. Annotating these images represents a significant manual cost. In the context of the large-scale dataset, we employ a characteristic of UDA to assess the performance of various approaches: the upper limit of UDA performance is supervised learning [2, 25]. The MOOSE (Multi-Organ Objective Segmentation) model [17] was trained on 50 LDCT images from PET/CT images to segment 40 non-cerebral structures, with expert manual segmentation. We evaluate UDA methods using the Dice Similarity Coefficient (DSC), comparing them to the MOOSE model, which segments fewer tissues (40 vs. 112 types). In comparisons, we merge related tissues (e.g., left and right kidneys into 'kidneys') to match the segmented categories.

Implements. We utilized a uniform preprocessing pipeline for source model training and domain adaptation. CT images were resampled to a spacing of $1.5 \times 1.5 \times 1.5 \text{ mm}^3$ and then randomly cropped into patches of size $192 \times 192 \times 192$. The Hounsfield units (HU) of the CT images were truncated for values less than -1000 and more excellent than 1000. This was followed by normalization to a range of 0 to 1 for input into the LUCIDA framework. The Adam optimizer was used for optimization, with a learning rate of $3e-5$ and parameter L2 regularization set at $3e-5$. The batch size was set to 4, and the training was conducted on four 80G-cached A800 GPUs. During the inference phase, the CT images were subjected to the same resampling process and segmented using a sliding window approach with a 0.5 overlap.

4.1 Results

Table 1 presents a comprehensive evaluation of various UDA methods applied to the target domain of LDCT datasets, focusing on segmentation performance as measured by the DSC. The UNet architecture, when unadapted, achieved a mean DSC of 90.8% across all categories, serving as a baseline for comparison. The introduction of UDA methods generally improved performance across all categories. Notably, the application of CycleGAN, AIGAN, and SIFA increased the mean DSC to 91.5%, 91.6%, and 91.6%, respectively. The novel method LUCIDA, without GAN-based adaptation, significantly outperformed other methods, achieving a mean DSC of 93.5%, indicating its effectiveness in domain adaptation. STUNet showed a remarkable baseline performance with a mean DSC of 93.3%. The introduction of UDA methods varied in effectiveness, with SIFA and FVP slightly improving the performance to 93.4% and 93.5%, respectively. LUCIDA continued to exhibit the highest performance enhancement, achieving a mean DSC of 94.6%.

Table 1. Performance comparison between UDA methods on target domain LDCT dataset. Mean DSC is evaluated. Abd. represents ‘abdomen’

Method	Adaptation	GAN	Param. (M)	Abd. Organs	Muscles	Fat	Skeleton	Vessels	Heart	Lungs	All
UNet	-	-	78.5	89.6	93.7	94.3	90.8	79.5	93.3	94.5	90.8
	CycleGAN	✓	108.7	90.2	94.1	94.6	91.3	80.4	94.7	95.3	91.5
	AIGAN	✓	164.8	90.4	94.2	95.1	91.5	80.0	94.3	95.6	91.6
	SIFA	✓	146.7	90.3	94.3	94.9	91.2	80.4	94.4	95.7	91.6
	Dn-Dp	-	1,476.0	90.1	93.9	94.9	91.1	81.4	94.8	95.8	91.7
	FVP	-	108.7	89.8	93.9	94.4	90.7	79.7	93.4	94.6	90.9
	LUCIDA	-	109.5	92.9	96.8	95.9	91.2	83.6	96.3	97.6	93.5
Swin- UNETR	-	-	248.1	90.5	94.5	94.7	91.1	80.1	94.1	95.1	91.4
	CycleGAN	✓	278.3	90.2	94.2	94.4	90.9	79.8	93.6	94.9	91.1
	AIGAN	✓	334.4	90.7	94.8	94.9	91.5	80.3	94.3	95.1	91.7
	SIFA	✓	316.3	90.6	94.6	94.8	91.2	80.1	94.2	95.0	91.5
	Dn-Dp	-	1,476.0	90.7	94.6	94.7	91.4	80.1	94.3	95.4	91.6
	FVP	-	248.1	90.4	94.4	94.6	91.2	79.9	94.3	95.5	91.5
	LUCIDA	-	279.1	92.4	96.2	95.6	90.9	83.2	94.8	95.9	92.7
STUNet	-	-	440.5	91.7	97.7	95.7	91.0	84.0	95.6	97.4	93.3
	CycleGAN	✓	470.7	90.3	96.8	94.1	89.2	83.5	94.0	96.4	92.0
	AIGAN	✓	526.8	92.3	97.1	94.5	90.8	83.4	95.3	97.7	93.0
	SIFA	✓	508.7	92.5	97.6	95.5	91.0	84.3	95.7	97.5	93.4
	Dn-Dp	-	1,476.0	91.9	96.7	95.2	90.5	84.8	95.1	96.9	93.0
	FVP	-	440.5	92.4	97.4	95.9	90.8	85.2	95.7	97.3	93.5
	LUCIDA	-	471.5	93.8	97.3	97.0	92.9	87.7	96.0	97.6	94.6

These results underscore the effectiveness of LUCIDA in enhancing segmentation performance across anatomical structures in LDCT images without the need for GAN-based adaptation. The consistent improvement across different architectures and categories suggests the potential of LUCIDA as a versatile and robust approach for domain adaptation in medical image analysis.

The results show limitations of GAN-based UDA methods, especially when the original model performs well on the target domain. For the UNet model, GAN-based improvements were minor (CycleGAN to 91.5%, AIGAN and SIFA to 91.6%). For STUNet, GAN methods decreased performance (to 92.0%, 93.0%, and 93.4%). This suggests GAN training’s instability might affect precision. Conversely, non-GAN models like FVP and LUCIDA consistently enhanced performance across architectures.

4.2 Tissue Analysis

Figure 2 illustrates our comparative study of 21 distinct organ tissues using five different models, including two baseline models and three variations of the LUCIDA model. This comprehensive analysis offers insights into the models’ performance across various organ tissues, ranging from the most easily recognizable to those notoriously challenging to identify.




Fig. 2. DSC performance comparison across 21 representative tissue types. LUCIDA-E represents the average ensemble results from predictions made by LUCIDA-UNet and LUCIDA-STUNet.

The results find that organs such as the spleen, kidneys, liver, muscles, and fat are the easiest to recognize. The source models without adaptation exhibit outstanding performance on these organs. A notable domain gap effect was observed in the segmentation of the stomach, pancreas, lungs, esophagus, trachea, small bowel, duodenum, and colon. However, the application of LUCIDA, particularly LUCIDA-E, remarkably bridged this gap. LUCIDA-E achieved an impressive average Dice Similarity Coefficient (DSC) of over 94% for these organs. The gallbladder, adrenal glands, vessels, duodenum, and bladder represent some

of the most challenging tissues for segmentation due to their smaller size and complex structural composition. LUCIDA-E facilitated an average performance exceeding 90% DSC in low-dose CT scans.

5 Conclusion

In this study, we introduced LUCIDA, a novel approach in the UDA realm for segmenting low-dose CT images without the need for LDCT annotations. Our comprehensive evaluation shows that LUCIDA significantly outperforms traditional UDA methods, marking a major advancement in medical imaging and offering an effective tool for clinicians and radiologists.

Acknowledgments. This work was supported by the Natural Science Foundation of China (62394311, 62394310). Zhaoeng Xie was supported by the Start-Up Funds of Peking University Health Science Center under Grant BMU2023YJ027 and Grant BMU2023GJJXK004.

Disclosure of Interests. The authors have no competing interests to declare that are relevant to the content of this article.

References

1. Unpaired image denoising via Wasserstein GAN in low-dose CT image with multi-perceptual loss and fidelity loss, author=Yin, Zhixian and Xia, Kewen and He, Ziping and Zhang, Jiangnan and Wang, Sijie and Zu, Baokai. *Symmetry* **13**(1), 126 (2021)
2. Chen, C., Dou, Q., Chen, H., Qin, J., Heng, P.A.: Synergistic image and feature adaptation: Towards cross-modality domain adaptation for medical image segmentation. In: *Proceedings of the AAAI Conference on Artificial Intelligence*. vol. 33, pp. 865–872 (2019)
3. Chen, C., Dou, Q., Chen, H., Qin, J., Heng, P.A.: Unsupervised Bidirectional Cross-Modality Adaptation via Deeply Synergistic Image and Feature Alignment for Medical Image Segmentation. *arXiv preprint arXiv:2002.02255* (2020)
4. Chen, H., Zhang, Y., Kalra, M.K., Lin, F., Chen, Y., Liao, P., Zhou, J., Wang, G.: Low-dose CT with a residual encoder-decoder convolutional neural network. *IEEE transactions on medical imaging* **36**(12), 2524–2535 (2017)
5. Chen, H., Zhang, Y., Zhang, W., Liao, P., Li, K., Zhou, J., Wang, G.: Low-dose CT via convolutional neural network. *Biomedical optics express* **8**(2), 679–694 (2017)
6. Fu, Y., Dong, S., Niu, M., Xue, L., Guo, H., Huang, Y., Xu, Y., Yu, T., Shi, K., Yang, Q., et al.: AIGAN: Attention–encoding Integrated Generative Adversarial Network for the reconstruction of low-dose CT and low-dose PET images. *Medical Image Analysis* **86**, 102787 (2023)
7. Gatidis, S., Hepp, T., Früh, M., La Fougère, C., Nikolaou, K., Pfannenberger, C., Schölkopf, B., Küstner, T., Cyran, C., Rubin, D.: A whole-body FDG-PET/CT Dataset with manually annotated Tumor Lesions. *Scientific Data* **9**(1), 601 (2022)
8. Hu, Z., Jiang, C., Sun, F., Zhang, Q., Ge, Y., Yang, Y., Liu, X., Zheng, H., Liang, D.: Artifact correction in low-dose dental CT imaging using Wasserstein generative adversarial networks. *Medical physics* **46**(4), 1686–1696 (2019)

9. Huang, Z., Chen, Z., Quan, G., Du, Y., Yang, Y., Liu, X., Zheng, H., Liang, D., Hu, Z.: Deep Cascade Residual Networks (DCRNs): Optimizing an Encoder–Decoder Convolutional Neural Network for Low-Dose CT Imaging. *IEEE Transactions on Radiation and Plasma Medical Sciences* **6**(8), 829–840 (2022)
10. Huang, Z., Liu, Z., He, P., Ren, Y., Li, S., Lei, Y., Luo, D., Liang, D., Shao, D., Hu, Z., et al.: Segmentation-guided denoising network for low-dose CT imaging. *Computer Methods and Programs in Biomedicine* **227**, 107199 (2022)
11. Huang, Z., Zhang, J., Zhang, Y., Shan, H.: DU-GAN: Generative adversarial networks with dual-domain U-Net-based discriminators for low-dose CT denoising. *IEEE Transactions on Instrumentation and Measurement* **71**, 1–12 (2021)
12. Huang, Z., Wang, H., Deng, Z., Ye, J., Su, Y., Sun, H., He, J., Gu, Y., Gu, L., Zhang, S., et al.: Stu-net: Scalable and transferable medical image segmentation models empowered by large-scale supervised pre-training. *arXiv preprint arXiv:2304.06716* (2023)
13. Khader, F., Müller-Franzes, G., Tayebi Arasteh, S., Han, T., Haarbuerger, C., Schulze-Hagen, M., Schad, P., Engelhardt, S., Baekler, B., Foersch, S., et al.: Denoising diffusion probabilistic models for 3D medical image generation. *Scientific Reports* **13**(1), 7303 (2023)
14. Li, H., Mueller, K.: Low-dose CT streak artifacts removal using deep residual neural network. *Proc. Fully Three-Dimensional Image Reconstruction Radiol. Nucl. Med.(Fully3D)* pp. 191–194 (2017)
15. Liu, X., Xie, Y., Cheng, J., Diao, S., Tan, S., Liang, X.: Diffusion Probabilistic Priors for Zero-Shot Low-Dose CT Image Denoising (2023)
16. Ronneberger, O., Fischer, P., Brox, T.: U-net: Convolutional networks for biomedical image segmentation. In: *Medical Image Computing and Computer-Assisted Intervention–MICCAI 2015: 18th International Conference, Munich, Germany, October 5–9, 2015, Proceedings, Part III* 18. pp. 234–241. Springer (2015)
17. Sundar, L.K.S., Yu, J., Muzik, O., Kulterer, O.C., Fueger, B., Kifjak, D., Nakuz, T., Shin, H.M., Sima, A.K., Kitzmantl, D., Badawi, R.D., Nardo, L., Cherry, S.R., Spencer, B.A., Hacker, M., Beyer, T.: Fully Automated, Semantic Segmentation of Whole-Body 18F-FDG PET/CT Images Based on Data-Centric Artificial Intelligence. *Journal of Nuclear Medicine* **63**(12), 1941–1948 (2022). <https://doi.org/10.2967/jnumed.122.264063>, <https://jnm.snmjournals.org/content/63/12/1941>
18. Wang, Y., Chen, Y., Wang, W., Zhu, H.: MSGAN: Multi-Stage Generative Adversarial Networks for Cross-Modality Domain Adaptation. In: *2022 44th Annual International Conference of the IEEE Engineering in Medicine & Biology Society (EMBC)*. pp. 520–524 (2022). <https://doi.org/10.1109/EMBC48229.2022.9871048>
19. Wang, Y., Cheng, J., Chen, Y., Shao, S., Zhu, L., Wu, Z., Liu, T., Zhu, H.: FVP: Fourier Visual Prompting for Source-Free Unsupervised Domain Adaptation of Medical Image Segmentation. *arXiv preprint arXiv:2304.13672* (2023)
20. Wasserthal, J., Meyer, M., Breit, H.C., Cyriac, J., Yang, S., Segeroth, M.: TotalSegmentator: robust segmentation of 104 anatomical structures in CT images. *arXiv preprint arXiv:2208.05868* (2022)
21. Xu, Q., Zhang, R., Zhang, Y., Wang, Y., Tian, Q.: A fourier-based framework for domain generalization
22. Yang, Y., Soatto, S.: Fda: Fourier domain adaptation for semantic segmentation. In: *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*. pp. 4085–4095 (2020)

23. Yin, Z., Xia, K., Wang, S., He, Z., Zhang, J., Zu, B.: Unpaired low-dose CT denoising via an improved cycle-consistent adversarial network with attention ensemble. *The Visual Computer* **39**(10), 4423–4444 (2023)
24. Zhou, L., Schaefferkoetter, J.D., Tham, I.W., Huang, G., Yan, J.: Supervised learning with cyclegan for low-dose FDG PET image denoising. *Medical image analysis* **65**, 101770 (2020)
25. Zhu, J.Y., Park, T., Isola, P., Efros, A.A.: Unpaired image-to-image translation using cycle-consistent adversarial networks pp. 2223–2232 (2017)