

Learning to Distill Global Representation for Sparse-View CT

Zilong Li¹ Chenglong Ma² Jie Chen¹ Junping Zhang¹ Hongming Shan^{2,3*}

¹ Shanghai Key Lab of Intelligent Information Processing, School of Computer Science, Fudan University, Shanghai 200433, China

² Institute of Science and Technology for Brain-inspired Intelligence and MOE Frontiers Center for Brain Science, Fudan University, Shanghai 200433, China

³ Shanghai Center for Brain Science and Brain-inspired Technology, Shanghai 200031, China

{zilongli21, clma22}@m.fudan.edu.cn, {chenj19, jpzhang, hmshan}@fudan.edu.cn

Abstract

Sparse-view computed tomography (CT)—using a small number of projections for tomographic reconstruction—enables much lower radiation dose to patients and accelerated data acquisition. The reconstructed images, however, suffer from strong artifacts, greatly limiting their diagnostic value. Current trends for sparse-view CT turn to the raw data for better information recovery. The resultant dual-domain methods, nonetheless, suffer from secondary artifacts, especially in ultra-sparse view scenarios, and their generalization to other scanners/protocols is greatly limited. A crucial question arises: have the image post-processing methods reached the limit? Our answer is not yet. In this paper, we stick to image post-processing methods due to great flexibility and propose global representation (GloRe) distillation framework for sparse-view CT, termed GloReDi. First, we propose to learn GloRe with Fourier convolution, so each element in GloRe has an image-wide receptive field. Second, unlike methods that only use the full-view images for supervision, we propose to distill GloRe from intermediate-view reconstructed images that are readily available but not explored in previous literature. The success of GloRe distillation is attributed to two key components: representation directional distillation to align the GloRe directions, and band-pass-specific contrastive distillation to gain clinically important details. Extensive experiments demonstrate the superiority of the proposed GloReDi over the state-of-the-art methods, including dual-domain ones. The source code is available at <https://github.com/longzilicart/GloReDi>.

1. Introduction

X-ray CT is one of the major modalities widely used in clinical screening and diagnosis. Despite the benefits,

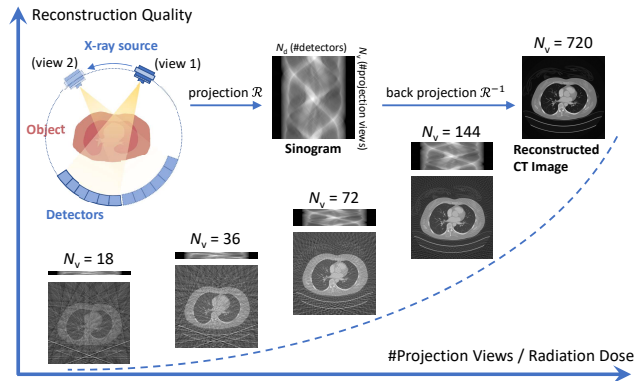


Figure 1. Sparse-view CT only uses a few projection views for tomographic reconstruction, thus providing fast and low-dose CT scanning. However, the reconstruction results of the conventional method suffer from severe *global* streak artifacts.

there have been growing concerns that X-ray radiation exposure could increase the risk of cancer induction [47]. Following the “As Low As Reasonably Achievable” principle in the medical community [36], efforts have been made to lower the radiation dose while maintaining imaging quality [20]. Sparse-view CT is one of the effective solutions which reduces the radiation by sampling part of the projection data, *a.k.a* sinogram, for image reconstruction, as shown in Fig. 1. However, analytical reconstruction algorithms such as filtered back projection (FBP) produce inferior image quality with globally severe streak artifacts, significantly compromising its diagnosis value. *How to effectively reconstruct the sparse-view CT* remains challenging and, hence, is gaining increasing attention in the computer vision and medical imaging communities.

Image-domain methods regard sparse-view CT reconstruction as an image post-processing task on the FBP-reconstructed images and have achieved exciting performance on streak artifacts removal and structure preserva-

tion [5, 22, 32, 58]. However, due to the limited receptive field, these methods have difficulty in modeling global information efficiently, thus leading to suboptimal results. Current dominant methods turn to the help of sinogram by restoring the sinogram and performing CT image post-processing simultaneously [17, 29, 31, 46]. Specifically, in addition to the image domain methods, deep neural networks are also applied to interpolate the missing data in sinogram so that the reconstructed images can be recovered in a global manner [10, 25].

Despite the common success of dual-domain methods, they suffer from unsatisfactory and unstable performance due to the following reasons. *First*, the processing of sinogram data is sensitive and even subtle changes may remarkably affect the reconstructed images and introduce stubborn secondary artifacts that are hard to be eliminated. *Second*, in ultra sparse-view scenarios, inpainting methods are unable to accurately restore the excessive missing raw data; e.g., inpainting sparse CT data of 18 views to the full CT data of 720 views only retains about 2.5% of the valid information, in which the involvement of sinogram processing makes the learning more difficult and compromise the performance [29, 46]. *Third*, it is often impractical to access sinogram data given the privacy and commercial concerns, and if accessible, the requirement of raw data greatly limits the generalization to other scanners/protocols.

Motivation. Given the great flexibility of the image domain methods, we raise a critical question in this paper: *have the image domain (post-processing) methods reached the limits?* Our answer is not yet, given the following observations. *First*, existing image-domain methods typically suffer from *limited* receptive field, which fails to extract and recover the global information efficiently [43], given that streak artifacts are spread globally on the reconstructed images as shown in Fig. 1. *Second*, existing image-domain methods only utilize the full-view images as the unique supervision on the result [22, 58]. They ignore the importance of representation learning, especially global representation (GloRe), making the artifact removal and detail recovery entangled, and leading to sub-optimal results due to the significant gap between sparse- and full-view. *Third*, the intermediate-view reconstructed images can be readily available during data preparation, but to the best of our knowledge, they are surprisingly neglected in the past literature. We argue that they can provide extra information and build bridges for sparse- and full-view CT reconstruction.

To address the abovementioned issues, we propose a novel image-domain method for sparse-view CT reconstruction. *First*, to address the limited receptive field inherent in conventional convolutional neural networks, we propose to learn GloRe with fast Fourier convolution (FFC) [7], so each element in GloRe has an image-wide receptive field. This global nature allows artifacts and information,

spread over the entire image, to be better modeled while also easing the alignment of the representations from different views. *Second*, to leverage extra supervision from intermediate-view CT images, we propose a novel distillation framework to learn better GloRe, termed GloReDi, which contains a parallel teacher network to distill knowledge from the readily intermediate-view images to provide high-quality and appropriate guidance for learning the GloRe of sparse-view CT images. Specifically, we first leverage intermediate-view reconstructed images to train a teacher network, which is then used to guide the learning of the student model (*i.e.* the sparse-view CT). The distillation scheme benefits GloRe in two folds: (1) representation directional distillation that aligns the directions between the student and teacher GloRe, which provides appropriate supervision considering the massive information loss due to the domain gaps between CT images with different views; and (2) band-pass-specific contrastive distillation that utilizes contrastive learning solely on the band-pass components to help distill the specific clinical value of each CT image without compromising the reconstruction accuracy.

Contributions. In summary, our contributions are listed as follows. *First*, we propose the global representation (GloRe) learning for sparse-view CT with Fourier convolution, which, to the best of our knowledge, is the first study to emphasize on the representation learning for image post-processing in sparse-view CT. *Second*, we propose a novel GloRe distillation framework, which can leverage the extra supervision from intermediate-view reconstructed images for high-quality information recovery and reconstruction. *Third*, we present representation directional distillation and band-pass-specific contrastive distillation for distilling GloRe to align the representation directions and gain clinically important details. *Last*, extensive experimental results demonstrate the superiority of GloReDi over the state-of-the-art sparse-view CT reconstruction methods in terms of quantitative metrics and visual comparison.

2. Detailed Related Work

2.1. Deep-learning-based Sparse-view CT

Numerous model-based iterative reconstruction (IR) methods have been developed to solve the sparse-view CT problem, mainly utilizing the sparsity of the total variation (TV) of the image [4, 12, 23, 33, 37, 51, 52]. However, most of them suffer from over-smoothed results, hand-crafted tuning for each image, and extremely high computation cost [46], limiting their clinical application. In contrast, deep learning methods are often faster and more accessible, thus arousing attention in this field. Among them, most image-domain methods regard sparse-view CT reconstruction as an image post-processing task. RedCNN [5], FBPCNN [22], and DDNet [58] are proposed to recon-




Figure 2. Overview of the proposed GloReDi for sparse-view CT reconstruction. GloReDi facilitates efficient information recovery by learning to distill the global representation from intermediate-view images.

struct the sparse-view CT by the convolutional neural network (CNN). These methods achieve competitive performance on streak artifact removal and structure preservation compared to conventional ones, but most of them fail to capture the global context and underperform in global reconstruction and artifact removal. Since streak artifacts result directly from the incomplete projection views, various methods are proposed to remove the artifacts by interpolating sinogram data [10, 25]. Further, dual-domain methods were becoming popular for their superior reconstruction performance by combining the knowledge of both domains. DuDoNet [31] first proposed a novel Radon inversion layer linking the gradient between the image and sinogram domain networks. Furthermore, various techniques have been proposed to enhance dual domain approaches through network design [1, 6] and unrolling architecture [57]. Recently, Transformer [45] has been introduced to dual-domain methods for its capability of capturing long-range dependencies, achieving superior performance [29, 46, 49]. However, the problems of the irreversible secondary artifacts and additional computational costs are not well addressed, and the requirement of raw data greatly limits their generalizability to other CT scanners/protocols. In this work, we challenge reconstructing sparse-view CT without raw data while still achieving state-of-the-art results.

2.2. Knowledge Distillation

Knowledge distillation transfers rich knowledge from teacher models to lightweight student models in the label or feature domains, mainly applied in model compression and acceleration [15, 41]. Among various approaches, feature-based knowledge distillation is extensively applied in image restoration tasks such as image super-resolution [14, 16, 27, 38]. It minimizes the distance between feature representations to learn richer information from the teacher model compared to softened labels [48]. Further, contrastive representation distillation is proposed to exploit the structural characteristics among different samples by harnessing the

discriminative representations, thereby enhancing differentiation within data [8, 19, 44]. This work distills knowledge from the intermediate-view images to enhance the sparse-view CT reconstruction.

2.3. Frequency Methods in Deep Learning

Frequency methods have been widely used in digital image processing and machine learning [2]. Wavelet convolution [32] and fast Fourier convolution [7] have been proposed to provide a larger/global receptive field, achieving great success in image restoration tasks [43]. Frequency methods have also yielded promising results in domain adaption and image translation, mainly by taking advantage of domain-invariant spectrum components [3, 18, 54–56]. Previous research also shows that particular features and details can be better extracted in the frequency domain, leading to improvements in camouflaged object detection [59], face forgery detection [21, 28], and face editing [9]. Frequency-based networks have also made notable strides in CT reconstruction. Previous studies [1, 26] employ multi-wavelet CNN [32] in both sinogram and image domain to suppress artifacts, substituting the pooling operator with wavelet transforms to achieve a larger receptive field. However, since both the Haar filter-based wavelet transform and convolution are applied on the spatial feature map, these methods struggle to capture global representation effectively. Recent research has demonstrated that employing convolution in the Fourier domain can efficiently eliminate the global artifacts [30, 34]. In this work, we utilize a frequency-based network for global representation learning, enhancing clinically critical features by band-pass-specific contrastive distillation.

3. Method

3.1. Problem Definition

Assume we have a two-dimensional (2D) CT slice of size $N \times N$, $I \in \mathbb{R}^{N \times N}$, where each pixel contains the attenua-

tion coefficient of the corresponding human body. The raw data or sinogram from the CT scanning, $\mathbf{S} \in \mathbb{R}^{N_v \times N_d}$, can be obtained via the Radon transform [39], where N_v and N_d denote the number of projection views and the number of the detectors, respectively. Inversely, the image reconstruction process can be expressed as $\mathbf{I} = \mathcal{R}^{-1}(\mathbf{S})$, where \mathcal{R}^{-1} denotes the inverse Radon transform such as FBP. Ideally, when N_v is sufficiently large, FBP can produce pleasing image quality. However, when N_v is fairly small, the image reconstruction of so-called sparse-view CT becomes an undetermined problem.

Here, we use \mathbf{I}_F and \mathbf{I}_S to represent the reconstructed images from full view (*i.e.* ground-truth) and sparse view, respectively. The image-domain methods for sparse-view CT are to improve the image quality of \mathbf{I}_S towards \mathbf{I}_F through a neural network \mathcal{F} . That is, the output of the network, $\hat{\mathbf{I}}_S = \mathcal{F}(\mathbf{I}_S)$, is expected to be of comparable quality to the full-view ground-truth. In addition, we also use the images, \mathbf{I}_T , reconstructed from an intermediate number of projection views that are between sparse and full views, as the teacher supervision, forming a student-teacher framework where the teacher encoder with \mathbf{I}_T supervises the training of student encoder with \mathbf{I}_S .

3.2. Overview of Our GloReDi

Fig. 2 presents the proposed GloReDi, which mainly consists of four parts: two encoders to learn global representations (GloRe) from sparse- and intermediate-view reconstructed images, \mathbf{I}_S , and \mathbf{I}_T , respectively, and two decoders to produce the final processed images from GloRe. In addition to learning GloRe with Fourier convolution, we have two novel modules to perform distillation from the teacher network: (1) representation directional distillation to align the directions between the student and teacher GloRes; and (2) band-pass-specific contrastive distillation to distill the clinically important features for better detail recovery.

Next, we detail each of the proposed components.

3.3. Global Representation Learning

The key technique behind the proposed network is the global representation learning with fast Fourier convolution (FFC) [7]. Concretely, FFC first applies real Fourier transform to get the frequency feature map and then performs convolution on frequency components before finally back transforming the frequency features as shown in Fig. 2. Therefore, GloRe from sparse-view images \mathbf{I}_S is learned by a Fourier-based encoder network as follows:

$$\mathbf{Z}_S = \text{encoder}_S(\mathbf{I}_S), \quad (1)$$

where $\mathbf{Z}_S \in \mathbb{R}^{N_w \times N_h \times N_c}$ denotes GloRe extracted from the sparse-view reconstructed image \mathbf{I}_S through the (student) encoder encoder_S and $N_w \times N_h \times N_c$ denotes the

width, height, and channel size of GloRe. Compared to the vanilla convolution with a limited receptive field, FFC works in the frequency domain, so each element in GloRe contains global information of the sparse-view image by using an image-wide receptive field, representing patterns of artifacts and the image contents simultaneously. Furthermore, such global nature can aid in modeling the artifact and information distributed throughout the image while easing the alignment between GloRe of different views.

3.4. Global Representation Distillation

In the context of sparse-view CT, images with an intermediate view \mathbf{I}_T can be easily obtained from the full-view sinogram as extra supervision. To get the teacher representation from \mathbf{I}_T , a parallel teacher encoder encoder_T , using the same architecture as encoder_S but with different weights, takes \mathbf{I}_T as input and yields teacher GloRe \mathbf{Z}_T of the same shape with \mathbf{Z}_S as follows:

$$\mathbf{Z}_T = \text{encoder}_T(\mathbf{I}_T). \quad (2)$$

\mathbf{Z}_T can be regarded as a high-quality pseudo target for student GloRe in the same latent space since richer information can be extracted from \mathbf{I}_T . Exploring distillation from \mathbf{I}_T is obviously easier than that of \mathbf{I}_F given the significant gap between \mathbf{I}_S and \mathbf{I}_F .

Intuitively, \mathbf{Z}_S should be similar to \mathbf{Z}_T if the artifact-free target information is precisely encoded. Therefore, a natural way for distillation is to use \mathbf{Z}_T as the supervision for \mathbf{Z}_S by minimizing their Euclidean distance. However, this could harm the reconstruction performance since the student encoder has to partially recover \mathbf{I}_T . *Finding an effective solution to distilling the information in \mathbf{Z}_T for sparse-view CT reconstruction* is one of the highlights in this work. Our idea is to separate the distillation process. First, we present representation directional distillation to align the representation directions so the images of different views are along the same direction. Second, we propose a band-pass-specific contrastive distillation to enhance clinically critical details.

We detail these two distillation techniques as follows.

3.4.1 Representation Directional Distillation

Given that \mathbf{I}_S and \mathbf{I}_T share the same ground truth \mathbf{I}_F , but differ mainly in artifact patterns, student and teacher GloRe should at least demonstrate a consistent direction for reconstruction with high confidence. Therefore, We define the directional distillation loss for \mathbf{Z}_S as follows:

$$\mathcal{L}_{\text{rdd}} = \frac{1}{N_w N_h} \sum_{i,j} \left(1 - \text{CosSim}([\mathbf{Z}_S]_{i,j}, [\mathbf{Z}_T]_{i,j}) \right), \quad (3)$$

where CosSim denotes the cosine similarity, defined as $\text{CosSim}(\mathbf{a}, \mathbf{b}) = \mathbf{a}^T \mathbf{b} / \max(\|\mathbf{a}\|_2 \|\mathbf{b}\|_2, \epsilon)$ with ϵ set to 1×10^{-8} to avoid division by zero.

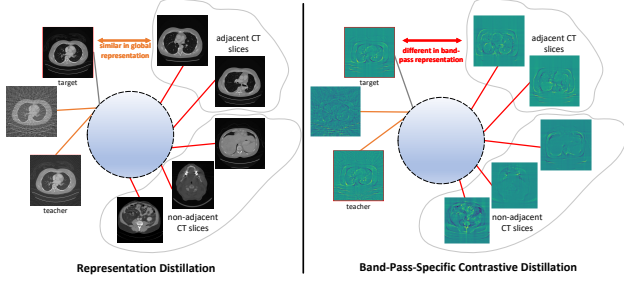


Figure 3. The difference between representation distillation (left) and band-pass-specific contrastive distillation (right). Global representation distillation fails to distinguish clinical details due to high similarity among similar body parts, *e.g.*, adjacent lung slices. Band-pass-specific contrastive distillation further projects the representation to the DCT domain, enabling distillation of specific clinical details on corresponding DCT components.

3.4.2 Band-Pass-Specific Contrastive Distillation

The directional distillation aligns the directions of student representation \mathbf{Z}_S and the teacher representation \mathbf{Z}_T in the view-independent space, which, however, does not guarantee a high-quality reconstruction of the image. What distinguishes sparse-view CT reconstruction task from classic image restoration one (*e.g.* denoising, deraining, *etc.*) is that the former is more demanding in terms of fine-grained details for diseases identification, organs delineation, and low-contrast lesions detection, *etc* [11]. In other words, it places more emphasis on learning the image-specific representation.

Contrastive representation learning, as a self-supervised approach, encourages the model to maximize the difference among representations of different samples so that model can emphasize those category-specific representations. However, directly implementing these losses would lead to imprecise reconstruction results because most CT images are inherently similar in intensity distribution, especially for images of the same body parts, as shown in Fig. 3. Recent methods have achieved great success learning on specific DCT spectrum [18, 50, 54], due to its strong energy compaction property, making it simple to select specific frequency components [40]. Therefore, we distilled the image-specific details by applying contrastive distillation on specific discrete cosine transform (DCT) spectrums. On the one hand, detailed knowledge can be better distilled from specific frequency components of teacher GloRe. On the other hand, contrastive against different CT slices can further enhance image-specific details on corresponding DCT components.

Discrete cosine transform. Discrete cosine transform (DCT) is widely used for data compression in low-frequency components. The basis function of 2D DCT is

defined as follows:

$$B_{w,h}^{i,j} = \cos\left(\frac{\pi w}{N_w}\left(i + \frac{1}{2}\right)\right) \cos\left(\frac{\pi h}{N_h}\left(j + \frac{1}{2}\right)\right), \quad (4)$$

where N_w and N_h represent the width and height of the input, respectively. Then, 2D DCT of the latent feature \mathbf{Z} can be written as:

$$f_{w,h}(\mathbf{Z}) = \sum_{i=0}^{N_w-1} \sum_{j=0}^{N_h-1} [\mathbf{Z}]_{i,j} B_{w,h}^{i,j}, \quad (5)$$

where $w \in \{0, 1, \dots, N_w - 1\}$, $h \in \{0, 1, \dots, N_h - 1\}$. The $f(\mathbf{Z}) \in \mathbb{R}^{N_w \times N_h}$ represents the 2D DCT frequency spectrum of \mathbf{Z} . Thus, frequency components can be easily selected by a mask $\mathbf{M} \in \{0, 1\}^{N_w \times N_h}$.

Band-pass-specific contrastive distillation. Crucially, the low-frequency components of GloRe contribute most to the reconstruction accuracy, while high-frequency may contain inherent noise and artifacts of CT images. We therefore use a band-pass supervised contrastive loss for the distillation, where a band-pass mask \mathbf{M} is incorporated to filter out the mid-frequency part in the frequency domain, containing main structure as well as the necessary detail information. The proposed band-pass supervised contrastive loss can be written as:

$$\mathbf{z} = \text{flatten}(\text{crop}(f_{w,h}(\mathbf{Z}), \mathbf{M})), \quad (6)$$

where `crop` operator will keep the value at locations where $\mathbf{M} = 1$ while the `flatten` operator flattens the 2D spectrum map into 1D representation.

It is to be observed that the clinical details in the CT images vary from case by case and are each very critical for diagnosis, requiring more sophisticated instance-level distinction. To this end, we consider all representations from different CT slices as negative samples to maximize image-specific features while updating the memory bank \mathcal{M} with the widely applied First-In-First-Out mechanism per iteration from the previous mini-batches. The proposed band-pass-specific contrastive distillation loss for sample \mathbf{z}_S can be written as follows:

$$\mathcal{L}_{\text{bcd}} = -\log \frac{\exp(\mathbf{z}_S^T \mathbf{z}_T / \tau)}{\sum_{\mathbf{z} \in \mathcal{M}} \exp(\mathbf{z}_S^T \mathbf{z} / \tau)}, \quad (7)$$

where \mathbf{z}_S , \mathbf{z}_T , and \mathbf{z} are frequency representations of \mathbf{Z}_S , \mathbf{Z}_T , and negative samples from memory bank \mathcal{M} of size N_{mem} , respectively. τ is the temperature term used to adjust the sensitivity of negative samples.

3.5. Loss Function and Training Procedure

Teacher reconstruction. To constrain the GloRe extracted from \mathbf{I}_S and \mathbf{I}_T in the same latent space, a naive solution is to share the same decoder in the training phase, which, however, brings instability. Therefore, we adopt the widely

Algorithm 1 The training process of GloReDi.

Require: encoder_S , encoder_T , decoder_S , decoder_T ; Input images: I_S , I_T , I_F ; momentum m and Memory bank

- 1: **for** $iter = 0$ to Max_Iter **do**
- 2: Update decoder_T using Eq. (8)
- 3: $\hat{I}_T, Z_T = \text{decoder}_T(\text{encoder}_T(I_T))$
- 4: Optimize encoder_T using Eq. (9)
- 5: $\hat{I}_S, Z_S = \text{decoder}_S(\text{encoder}_S(I_S))$
- 6: Optimize $\text{encoder}_S + \text{decoder}_S$ using Eq. (11)
- 7: Update memory bank
- 8: **end for**

Return: $\text{encoder}_S + \text{decoder}_S$

used Exponential Moving Average (EMA) with momentum m to update the parameters θ of the teacher decoder by the student decoder [13]:

$$\theta_{\text{decoder}_T} = m\theta_{\text{decoder}_T} + (1 - m)\theta_{\text{decoder}_S}. \quad (8)$$

Then, the global representation Z_T learned by the teacher is fed into the decoder to obtain the teacher reconstruction result \hat{I}_T , and we use ℓ_1 loss to measure the pixel-wise difference between \hat{I}_T and the full-view ground-truth I_F and train the teacher:

$$\mathcal{L}_{\text{pixel}_T} = \|\text{decoder}_T(Z_T) - I_F\|_1. \quad (9)$$

Student reconstruction. The pixel-wise error measurement for the student is defined in a similar way:

$$\mathcal{L}_{\text{pixel}_S} = \|\text{decoder}_S(Z_S) - I_F\|_1. \quad (10)$$

Finally, a compound loss function combining pixel-wise error and two distillation losses, respectively defined in Eq. (3) and Eq. (7) is introduced for student network training to achieve high-quality sparse-view CT reconstruction:

$$\mathcal{L}_{\text{Stu}} = \mathcal{L}_{\text{pixel}_S} + \alpha\mathcal{L}_{\text{rdd}} + \beta\mathcal{L}_{\text{bcd}}, \quad (11)$$

GloReDi is trained following the process summarized in Algorithm 1, where the student and teacher encoder are trained iteratively.

4. Experimental Results

4.1. Experimental Setup

Dataset. We use DeepLesion dataset [53] and “2016 NIH-AAPM-Mayo Clinic Low-Dose CT Grand Challenge” AAPM dataset [35] to demonstrate the effectiveness of the proposed GloReDi. The DeepLesion dataset is the largest multi-lesion real-world CT dataset made available to the

public, from which 40,000 images of 303 patients are selected as the training set while 1000 images of another 18 patients are selected as the test set. AAPM dataset contains routine dose CT data from 10 patients, where a total of 5,410 slices from 9 patients are chosen for training and 526 slices from the remaining 1 patient for testing. All images are resized to 256×256 . We simulate the forward and back projection using fan-beam geometry under 120 kVp and 500 mA with TorchRadon toolbox [42]. The distance from X-ray source to the rotation center is 59.5cm, and the number of detectors is set to 672. Sparse-view CT images are generated from $N_v = 18, 36, 72, 144$ projection views uniformly sampled from full 720 views covering $[0, 2\pi]$. To simulate the photon noise presented in real-world CT, an intensity of 1×10^6 Poisson noise is added to the sinograms.

Implementation details. Our model is implemented in PyTorch. We use Adam optimizer [24] with $(\beta_1, \beta_2) = (0.5, 0.999)$ to train the model. The learning rate starts from 1×10^{-3} and is halved for every 40 epochs. For each method, four models of different N_v are trained separately on 4 NVIDIA RTX 3090 GPUs for 120 epochs with a batch size of 8. For hyperparameters in GloReDi, we set α and β in Eq. (11) empirically to 0.1 and 0.0002 to fit the scale between losses. The momentum m and temperature τ are set to 0.9 and 1.0, respectively, while the size L of the memory bank is set to 300 to balance the performance and computational cost. We set intermediate-view as $N_v \times 2$ by default. More details, including the network architecture, can be found in the Supplementary Material.

Evaluation metrics. For quantitative evaluation, we use peak signal-to-noise ratio (PSNR), structural similarity (SSIM) [60] and root mean square error (RMSE); all of them are widely adopted for image quality assessment.

4.2. Comparison with State-of-the-Art Methods

We compare GloReDi with the following state-of-the-art methods: DDNet [58], FBPCConvNet [22], DuDoNet [31], DDPTrans [29], and DuDoTrans [46]. In addition, we name the network without the distillation from I_T as the frequency encoder and decoder network (Fred-Net), optimized with only the pixel-wise loss in Eq. (10). FBP directly applies Radon transform to reconstruct the image from the sparse-view sinogram. DDNet and FBPCConvNet are image-domain deep-learning methods taking FBP results as input. DuDoNet is a state-of-the-art dual-domain method that restores the CT image using two U-Nets in sinogram and image domains. Furthermore, DDPTrans and DuDoTrans adopt Transformer on both domains to restore the image by utilizing the long-range dependency. We try our best to reproduce these methods and then train and test them on the same dataset for a fair comparison.

Quantitative comparison. Table 1 presents the quantita-

Methods	$N_v = 18$			$N_v = 36$			$N_v = 72$			$N_v = 144$		
	PSNR	SSIM	RMSE	PSNR	SSIM	RMSE	PSNR	SSIM	RMSE	PSNR	SSIM	RMSE
FBP	22.22	35.36	0.0795	25.49	47.49	0.0543	30.88	63.81	0.0293	37.12	82.94	0.0144
DDNet [58]	34.57	91.94	0.0187	38.24	94.91	0.0126	41.66	97.27	0.0083	46.75	98.79	0.0048
FBPConvNet [22]	35.95	93.62	0.0164	39.79	96.13	0.0106	43.76	97.47	0.0066	48.46	98.60	0.0039
DuDoNet [31]	35.69	93.96	0.0169	40.36	96.94	0.0103	44.86	98.39	0.0059	49.33	99.23	0.0036
DDPTrans [29]	35.11	93.48	0.0181	38.68	95.99	0.0121	43.56	98.16	0.0069	48.72	99.22	0.0038
DuDoTrans [46]	36.08	93.28	0.0161	40.75	96.67	0.0095	45.16	98.44	0.0057	49.96	99.28	0.0034
Fred-Net (ours)	38.08	95.20	0.0129	40.86	96.81	0.0093	44.43	98.18	0.0063	48.45	99.08	0.0039
GloReDi (ours)	38.65	95.87	0.0120	41.25	97.05	0.0090	45.18	98.43	0.0057	48.96	99.21	0.0037

Table 1. Quantitative evaluation [PSNR (db), SSIM (%) and RMSE] for state-of-the-art methods on DeepLesion dataset.

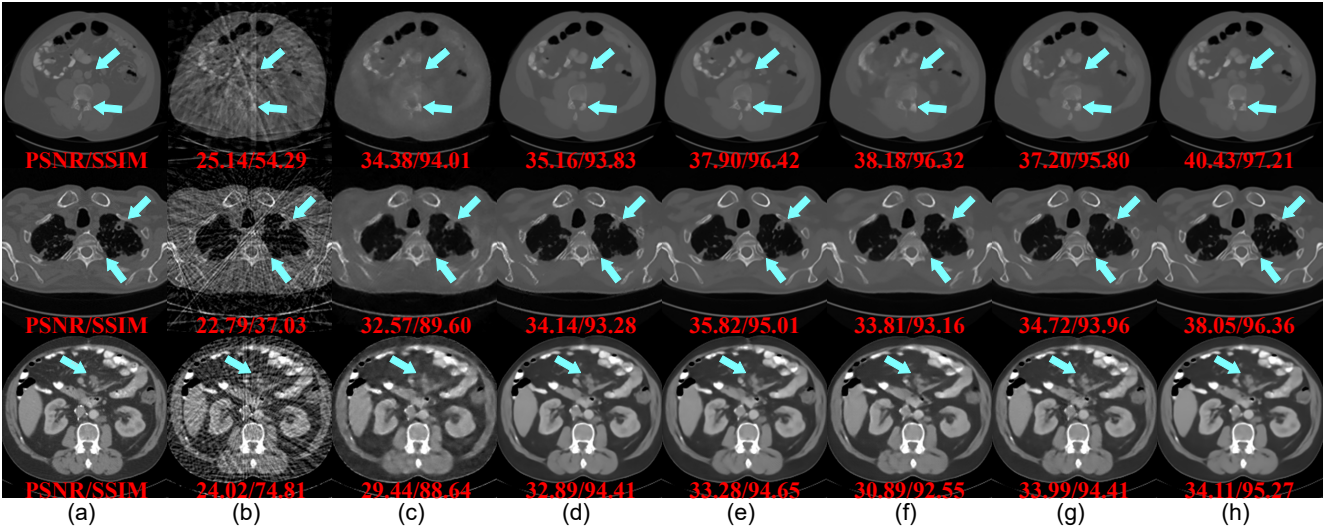


Figure 4. Visual comparison of state-of-the-art methods on DeepLesion dataset: (a) Ground Truth, (b) FBP, (c) DDNet, (d) FBPConvNet, (e) DuDoNet, (f) DDPTrans, (g) DuDoTrans, and (h) GloReDi. From top to bottom: the results under $N_v = 18, 36, 72$; display window is set to $[-1000, 2000]$ HU for the first two rows and $[-200, 300]$ HU for the third row.

itive comparison. Generally, dual-domain methods perform better than single-domain methods, especially in the case of $N_v = 72, 144$ thanks to the sinogram inpainting network. However, we highlight in ultra sparse-view scenarios where $N_v = 18$, that dual-domain methods like DuDoNet perform worse than classical image-domain-only methods such as FBPConvNet. This can be attributed to the secondary artifacts introduced by the unsuccessful sinogram inpainting; *e.g.*, when $N_v = 18$ in our setting, restoring the sinogram can be recognized as a $40\times$ super-resolution task, which is rather challenging. DuDoTrans is superior to DuDoNet to some extent, indicating the importance of long-range dependency in reconstruction. DDPTrans fails to surpass DuDoNet because of the limited parameter number though most transformers are memory-consuming.

Compared with the previous works, the proposed GloReDi achieves competitive results by learning to distill GloRe for high-quality reconstruction. Particularly in ultra-sparse scenarios, *e.g.*, when $N_v = 18$, GloReDi has 2.57db and 2.70db improvement on PSNR over DuDoTrans and

FBPConvNet, respectively. As N_v increases, the advantage of GloReDi decreases, mainly due to a shift towards artifacts removal tasks rather than reconstruction. We highlight that the fewer projection views indicate less radiation dose and more speedup. We notice that global representation distillation further improves the performance of GloReDi from 0.51db to 0.89db compared to Fred-Net. The results show that the proposed GloReDi successfully inspire the potential of image post-processing methods and is significantly superior to the state-of-the-art methods under various settings.

Visual comparison. Fig. 4 presents the visualization results of three sparse-view CT. DDNet and FBPConvNet cannot remove the artifact thoroughly, while the reconstruction results are over-smooth in the boundary of organs and bones. Such effects can be easily observed in the second row of Figs. 4(c) and 4(d) on the spine where the anatomical structure is complicated. Although dual-domain methods perform better in general, their accuracy in details is even worse than FBPConvNet in comparison with Figs. 4(d) and 4(e)-(g). Among all these methods, Glo-

Methods	$N_v = 18$			$N_v = 36$			$N_v = 72$			$N_v = 144$		
	PSNR	SSIM	RMSE	PSNR	SSIM	RMSE	PSNR	SSIM	RMSE	PSNR	SSIM	RMSE
FBP	22.73	35.06	0.0732	26.27	47.52	0.0486	31.36	64.56	0.0270	37.85	84.57	0.0128
DDNet [58]	34.29	89.68	0.0194	36.95	93.03	0.0142	40.41	96.04	0.0096	44.35	98.05	0.0061
FBPConvNet [22]	35.73	92.84	0.0165	37.95	93.73	0.0127	42.92	97.19	0.0072	47.35	98.76	0.0043
DuDoNet [31]	34.82	93.00	0.0196	39.89	95.99	0.0102	44.06	98.02	0.0062	48.39	99.11	0.0038
DDPTrans [29]	34.47	91.95	0.0191	38.13	94.83	0.0125	42.77	97.65	0.0073	47.84	99.06	0.0041
DuDoTrans [46]	35.85	93.07	0.0162	40.04	96.02	0.0100	44.20	98.07	0.0062	49.02	99.21	0.0036
Fred-Net (ours)	37.20	93.90	0.0140	41.46	96.84	0.0085	43.64	97.84	0.0066	47.50	99.03	0.0042
GloReDi (ours)	37.91	94.58	0.0128	41.57	96.70	0.0084	44.24	98.03	0.0062	48.10	99.08	0.0040

Table 2. Quantitative evaluation [PSNR (db), SSIM (%) and RMSE] for state-of-the-art methods transferred to AAPM dataset.




Figure 5. Visual comparison of state-of-the-art methods on AAPM dataset: (a) Ground Truth, (b) FBP, (c) DDNet, (d) FBPConvNet, (e) DuDoNet, (f) DDPTrans, (g) DuDoTrans, and (h) GloReDi. From top to bottom: $N_v = 18, 36$; the display window is $[-1000, 2000]$ HU.

ReDi best reconstructs the structures and details, especially in ultra-sparse scenarios. In the first two rows of Fig. 4(h), only GloReDi precisely recovers those clinically important details such as soft tissue and corrupted pulmonary alveoli.

4.3. Ablation Study

We first evaluate the effectiveness of each component in GloReDi. We use Fred-Net as the baseline to add components. The configurations involved in Table 3 are mainly four groups: (1) the baseline Fred-Net using pixel-wise loss to train student model without distillation components (a); (2) the baseline without Fourier convolutions (b–c); (3) the baseline with distillation from intermediate-view reconstructed images (d–g); and (4) the dual-domain version with raw data processed by sinogram-domain sub-network of DuDoNet [31] (h). Unless noted otherwise, the pixel-wise loss is involved in training.

Ablations on global representations. The results between a) and b) in Table 3 confirm that learning global representations with the image-wide receptive field by Fourier convolution is beneficial for sparse CT reconstruction. Also, we notice that in Tables 1 and 2, our baseline Fred-Net still outperforms state-of-the-art methods when $N_v = 18$ and 36. Interestingly, we found that c) performs even worse than b) in Table 3. This suggests that global representation is essential for bridging the gap between images of different views. Generally, learning global representations is pow-

	config.	PSNR
a)	Fred-Net	38.09
b)	Fred-Net w/o Fourier	36.52
c)	Fred-Net w/o Fourier + $\mathcal{L}_{rdd} + \mathcal{L}_{bcd}$	36.30
d)	Fred-Net + data augmentation	34.41
e)	Fred-Net + \mathcal{L}_{rdd}	38.39
f)	Fred-Net + \mathcal{L}_{bcd}	38.42
g)	Fred-Net + $\mathcal{L}_{rdd} + \mathcal{L}_{bcd} = \mathbf{GloReDi}$	38.65
h)	Fred-Net + raw data	38.44

Table 3. Quantitative evaluation of different configurations.

erful and provides a new perspective for sparse-view CT reconstruction, which also answers that the limit of image post-processing methods is still far beyond reach.

Ablations on global representation distillation. Through a comparison between Fred-Net and the one trained with intermediate-view images as data augmentation (a vs. d) in Table 3, we found that directly applying data augmentation had a negative impact, resulting in a 3.67db PSNR drop. This indicates that the network failed to benefit from denser-view images directly because of the significant domain gap among images of different views. By comparing a), e) and f) in Table 3, we can see that both e) representation directional distillation and f) band-pass-specific contrastive distillation can improve the performance of a), showing that end-to-end supervision cannot thoroughly unlock the potential of the GloRe. By training with both distillation loss, g) further boosts the performance and outperforms all other

config.	$N_v \times 2$	$N_v \times 3$	$N_v \times 4$	720
$N_v = 18$	38.38	38.29	38.28	38.02
$N_v = 72$	44.84	44.46	44.72	44.38

Table 4. PSNR evaluation of GloReDi trained with different intermediate-view images for 60 epochs. The first and second rows show the results of sparse-view images with $N_v = 18, 72$ distilled from images of $N_v \times 2, 3, 4$ and 720 (full view), respectively.

configurations. By comparing **a**), **g**) and **h**), we found unsurprisingly that using raw data does improve the performance of the image-domain network. However, our method GloReDi yields even better performance by effectively distilling GloRe from intermediate-view image data.

Ablations on intermediate views for distillation. Selecting a suitable intermediate view for distillation is crucial in realizing the full potential of GloReDi. Images reconstructed from a denser view can provide richer information to the teacher GloRe while introducing a more significant domain gap. As presented in Table 4, we found that models distilled from $N_v \times 2$ views exhibit the best performance. Interestingly, models distilled from the full-view image yield the worst performance, primarily because of the considerable domain gap between the input data. Please refer to Supplementary Material for more ablation studies on *framework designs, configurations of residual blocks, and distillation loss*.

4.4. Transfer to Other Dataset

To test the generalizability and robustness of GloReDi along with other state-of-the-art methods, we finetune each model for another ten epochs on AAPM dataset to bridge the domain gap following the same setting. In Table 2, the proposed GloReDi shows excellent transferability over all other state-of-the-art methods in (ultra) sparse-view scenarios ($N_v = 18, 36, 72$) and still achieves performance comparable with the dual-domain methods when $N_v = 72$. Specifically, when $N_v = 18$, we notice GloReDi has 2.06db and 2.18db improvements on PSNR compared to DuDotrans and FBPCovNet, respectively, indicating that our method is significantly superior in ultra sparse-view scenarios. Fig. 5 shows images with $N_v = 18$ and 36, which are corrupted severely by FBP as shown in Fig. 5(b), thus losing its clinical value. Previous image-domain methods remove the artifact to some extent but lose most of the details, as shown in Figs. 5(c) and (d). Although the state-of-the-art dual-domain methods gain better results in general, we notice that clinical details such as lesion and bone, as pointed out by the blue arrows, are not accurately reconstructed in Figs. 5(e)-(g) while severe secondary artifacts are introduced, especially when $N_v = 18$. In contrast, the proposed GloReDi excels at reconstructing the details and greatly improves the clinical value of sparse-view CT.

5. Discussion and Conclusion

In this paper, we challenge the sparse-view CT reconstruction in the absence of sinogram data and demonstrate that, with appropriate design, the image post-processing method can outperform the state-of-the-art dual-domain method. We propose to learn GloRe with Fourier convolution for image-wide receptive field and present directional and band-pass-specific contrastive distillation for high-quality reconstruction. Our method is particularly advantageous in ultra-sparse scenarios.

We emphasize that our contributions can be easily extended to dual-domain methods when sinogram data are available. On the one hand, learning to distill GloRe can also benefit the image-domain network in a dual-domain framework. On the other hand, the proposed method can also be used to enhance the sinogram recovery by distilling the sinogram with intermediate views. Future work can concentrate on extending GloReDi to dual domain methods and other CT reconstruction tasks, such as limited-views CT reconstruction and metal artifact reduction. Although the teacher network is trained alongside the student network in the current version, it is valuable to explore the utilization of pre-trained teacher models, enabling GloReDi to progressively distill the GloRe from a series of teacher models.

We also acknowledge some limitations. First, the teacher network will introduce additional computational costs during the training phase. It is possible to study how to distill knowledge from a pretrained teacher network. Second, although there are clear improvements in quantitative and visual results, it is desirable to have feedback from radiologists for clinical practice. The generalizability and robustness of the trained networks in clinical applications require further investigation. Third, we only considered 2D cases in the paper due to memory respect. If GPU memory permits, one can extend GloReDi to 3D by replacing 2D convolutions with 3D ones by applying high-dimension kernels for FFCs [7]. Lastly, GloReDi failed to surpass the dual-domain methods in denser view scenarios when sinogram can provide effective information for performance improvement. Studying how to improve the performance given an arbitrary N_v would be of great value for future research.

Acknowledgement This work was supported in part by National Natural Science Foundation of China (Nos. 62101136 and 62176059), Shanghai Sailing Program (No. 21YF1402800), Shanghai Municipal Science and Technology Major Project (No. 2018SHZDZX01), ZJ Lab, Shanghai Municipal of Science and Technology Project (No. 20JC1419500), and Shanghai Center for Brain Science and Brain-inspired Technology.

References

- [1] Jieli Bai, Yitong Liu, and Hongwen Yang. Sparse-view CT reconstruction based on a hybrid domain model with multi-level wavelet transform. *Sensors*, 22(9):3228, 2022. 3
- [2] Gregory A Baxes. *Digital image processing: principles and applications*. John Wiley & Sons, Inc., 1994. 3
- [3] Mu Cai, Hong Zhang, Huijuan Huang, Qichuan Geng, Yixuan Li, and Gao Huang. Frequency domain image translation: More photo-realistic, better identity-preserving. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 13930–13940, 2021. 3
- [4] Guang-Hong Chen, Jie Tang, and Leng Shuai. Prior image constrained compressed sensing (PICCS): A method to accurately reconstruct dynamic CT images from highly under-sampled projection data sets. *Medical Physics*, 35(2):660–663, 2008. 2
- [5] Hu Chen, Yi Zhang, Mannudeep K. Kalra, Feng Lin, Yang Chen, Peixi Liao, Jiliu Zhou, and Ge Wang. Low-dose CT with a residual encoder-decoder convolutional neural network. *IEEE Transactions on Medical Imaging*, 36(12):2524–2535, 2017. 2
- [6] Theodor Chelsterean-Boghiu, Felix C. Hofmann, Manuel Schultheiss, Franz Pfeiffer, Daniela Pfeiffer, and Tobias Lasser. WNet: A data-driven dual-domain denoising model for sparse-view computed tomography with a trainable reconstruction layer. *IEEE Transactions on Computational Imaging*, 9:120–132, 2022. 3
- [7] Lu Chi, Borui Jiang, and Yadong Mu. Fast Fourier convolution. In *Advances in Neural Information Processing Systems*, volume 33, pages 4479–4488, 2020. 2, 3, 4, 9
- [8] Hao Fu, Shaojun Zhou, Qihong Yang, Junjie Tang, Guiquan Liu, Kaikui Liu, and Xiaolong Li. Lrc-bert: latent-representation contrastive knowledge distillation for natural language understanding. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 35, pages 12830–12838, 2021. 3
- [9] Yue Gao, Fangyun Wei, Jianmin Bao, Shuyang Gu, Dong Chen, Fang Wen, and Zhouhui Lian. High-fidelity and arbitrary face editing. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 16115–16124, 2021. 3
- [10] Muhammad Usman Ghani and W. Clem Karl. Deep learning-based sinogram completion for low-dose CT. In *2018 IEEE 13th Image, Video, and Multidimensional Signal Processing Workshop*, pages 1–5, 2018. 2, 3
- [11] Fatemeh Haghighi, Mohammad Reza Hosseinzadeh Taher, Michael B. Gotway, and Jianming Liang. DiRA: Discriminative, restorative, and adversarial learning for self-supervised medical image analysis. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 20824–20834, 2022. 5
- [12] Ji He, Yan Yang, Yongbo Wang, Dong Zeng, Zhaoying Bian, Hao Zhang, Jian Sun, Zongben Xu, and Jianhua Ma. Optimizing a parameterized plug-and-play ADMM for iterative low-dose CT reconstruction. *IEEE Transactions on Medical Imaging*, 38(2):371–382, 2019. 2
- [13] Kaiming He, Haoqi Fan, Yuxin Wu, Saining Xie, and Ross Girshick. Momentum contrast for unsupervised visual representation learning. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 9729–9738, 2020. 6
- [14] Zibin He, Tao Dai, Jian Lu, Yong Jiang, and Shu-Tao Xia. Fakd: Feature-affinity based knowledge distillation for efficient image super-resolution. In *2020 IEEE International Conference on Image Processing (ICIP)*, pages 518–522, 2020. 3
- [15] Geoffrey Hinton, Oriol Vinyals, and Jeff Dean. Distilling the knowledge in a neural network. *arXiv preprint arXiv:1503.02531*, 2015. 3
- [16] Ming Hong, Yuan Xie, Cuihua Li, and Yanyun Qu. Distilling image dehazing with heterogeneous task imitation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 3462–3471, 2020. 3
- [17] Dianlin Hu, Jin Liu, Tianling Lv, Qianlong Zhao, Yikun Zhang, Guotao Quan, Juan Feng, Yang Chen, and Limin Luo. Hybrid-domain neural network processing for sparse-view CT reconstruction. *IEEE Transactions on Radiation and Plasma Medical Sciences*, 5(1):88–98, 2020. 2
- [18] Jiaying Huang, Dayan Guan, Aoran Xiao, and Shijian Lu. FSDR: Frequency space domain randomization for domain generalization. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 6891–6902, 2021. 3, 5
- [19] Zhizhong Huang, Jie Chen, Junping Zhang, and Hongming Shan. Learning representation for clustering via prototype scattering and positive sampling. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 2022. 3
- [20] Zhizhong Huang, Junping Zhang, Yi Zhang, and Hongming Shan. 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. 1
- [21] Shuai Jia, Chao Ma, Taiping Yao, Bangjie Yin, Shouhong Ding, and Xiaokang Yang. Exploring frequency adversarial attacks for face forgery detection. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 4103–4112, 2022. 3
- [22] Kyong Hwan Jin, Michael T McCann, Emmanuel Froustey, and Michael Unser. Deep convolutional neural network for inverse problems in imaging. *IEEE Transactions on Image Processing*, 26(9):4509–4522, 2017. 2, 6, 7, 8
- [23] Kyongsang Kim, Jong Chul Ye, William Worstell, Jinsong Ouyang, Yothin Rakvongthai, Georges El Fakhri, and Quanzheng Li. Sparse-view spectral CT reconstruction using spectral patch-based low-rank penalty. *IEEE Transactions on Medical Imaging*, 34(3):748–760, 2015. 2
- [24] Diederik P. Kingma and Jimmy Ba. Adam: A method for stochastic optimization. *arXiv preprint arXiv:1412.6980*, 2014. 6
- [25] Hoyeon Lee, Jongha Lee, Hyeonseok Kim, Byungchul Cho, and Seungryong Cho. Deep-neural-network-based sinogram synthesis for sparse-view CT image reconstruction. *IEEE Transactions on Radiation and Plasma Medical Sciences*, 3(2):109–119, 2018. 2, 3

- [26] Minjae Lee, Hyemi Kim, and Hee-Joung Kim. Sparse-view CT reconstruction based on multi-level wavelet convolution neural network. *Physica Medica*, 80:352–362, 2020. 3
- [27] Wonkyung Lee, Junghyup Lee, Dohyung Kim, and Bumsuh Ham. Learning with privileged information for efficient image super-resolution. In *Computer Vision—ECCV 2020: 16th European Conference, Glasgow, UK, August 23–28, 2020, Proceedings, Part XXIV 16*, pages 465–482, 2020. 3
- [28] Jiaming Li, Hongtao Xie, Jiahong Li, Zhongyuan Wang, and Yongdong Zhang. Frequency-aware discriminative feature learning supervised by single-center loss for face forgery detection. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 6458–6467, 2021. 3
- [29] Runrui Li, Qing Li, Hexi Wang, Saize Li, Juanjuan Zhao, Yan Qiang, and Long Wang. DDPTransformer: Dual-domain with parallel transformer network for sparse view CT image reconstruction. *IEEE Transactions on Computational Imaging*, pages 1–15, 2022. 2, 3, 6, 7, 8
- [30] Zilong Li, Qi Gao, Yaping Wu, Chuang Niu, Junping Zhang, Meiyun Wang, Ge Wang, and Hongming Shan. Quad-Net: Quad-domain network for CT metal artifact reduction. *arXiv preprint arXiv:2207.11678*, 2023. 3
- [31] Wei-An Lin, Cheng Liao, Haofu abd Peng, Xiaohang Sun, Jingdan Zhang, Jiebo Luo, Rama Chellappa, and Zhou S. Kevin. DuDoNet: Dual domain network for CT metal artifact reduction. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 10512–10521, 2019. 2, 3, 6, 7, 8
- [32] Pengju Liu, Hongzhi Zhang, K. Zhang, Liang Lin, and Wangmeng Zuo. Multi-level wavelet-cnn for image restoration. *2018 IEEE/CVF Conference on Computer Vision and Pattern Recognition Workshops (CVPRW)*, pages 886–88609, 2018. 2, 3
- [33] Yi Liu, Shangquan Hong, Quan Zhang, Hongqing Zhu, Huazhong Shu, and Zhiguo Gui. Median prior constrained TV algorithm for sparse view low-dose CT reconstruction. *Computers in Biology and Medicine*, 60:117–131, 2015. 2
- [34] Chenglong Ma, Zilong Li, Junping Zhang, Yi Zhang, and Hongming Shan. FreeSeed: Frequency-band-aware and self-guided network for sparse-view CT reconstruction. *arXiv preprint arXiv:2307.05890*, 2023. 3
- [35] C. McCollough. TU-FG-207A-04: Overview of the low dose CT grand challenge. *Medical Physics*, 43(6):3759–3760, 2016. 6
- [36] Donald L Miller and David Schauer. The alara principle in medical imaging. *Philosophy*, 44:595–600, 1983. 1
- [37] Shanzhou Niu, Yang Gao, Zhaoying Bian, Jing Huang, Wufan Chen, Gaohang Yu, Zhengrong Liang, and Jianhua Ma. Sparse-view X-ray CT reconstruction via total generalized variation regularization. *Physics in Medicine and Biology*, 59(12):2997–3017, 2014. 2
- [38] SeongUk Park and Nojun Kwak. Local-selective feature distillation for single image super-resolution. *arXiv preprint arXiv:2111.10988*, 2021. 3
- [39] J. Radon. Über die Bestimmung von Funktionen durch ihre Integralwerte längs gewisser Mannigfaltigkeiten. *Bericht über die Verhandlungen der Königlich-Säberchsischen Akademie der Wissenschaften zu Leipzig*, 69:262–277, 1917. 4
- [40] K Ramamohan Rao and Ping Yip. *Discrete cosine transform: algorithms, advantages, applications*. Academic press, 2014. 5
- [41] Adriana Romero, Nicolas Ballas, Samira Ebrahimi Kahou, Antoine Chassang, Carlo Gatta, and Yoshua Bengio. Fitnets: Hints for thin deep nets. *arXiv preprint arXiv:1412.6550*, 2014. 3
- [42] Matteo Ronchetti. TorchRadon: Fast differentiable routines for computed tomography. *arXiv preprint arXiv:2009.14788*, 2020. 6
- [43] Roman Suvorov, Elizaveta Logacheva, Anton Mashikhin, Anastasia Remizova, Arsenii Ashukha, Aleksei Silvestrov, Naejin Kong, Harshith Goka, Kiwoong Park, and Victor Lempitsky. Resolution-robust large mask inpainting with Fourier convolutions. In *2022 IEEE/CVF Winter Conference on Applications of Computer Vision*, pages 2149–2159, 2022. 2, 3
- [44] Yonglong Tian, Dilip Krishnan, and Phillip Isola. Contrastive representation distillation. *arXiv preprint arXiv:1910.10699*, 2019. 3
- [45] Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N. Gomez, Łukasz Kaiser, and Illia Polosukhin. Attention is all you need. In *Proceedings of the 31st International Conference on Neural Information Processing Systems*, pages 6000–6010, 2017. 3
- [46] Ce Wang, Kun Shang, Haimiao Zhang, Qian Li, and S. Kevin Zhou. DuDoTrans: Dual-domain transformer for sparse-view CT reconstruction. In *Machine Learning for Medical Image Reconstruction*, pages 84–94. Springer International Publishing, 2022. 2, 3, 6, 7, 8
- [47] Ge Wang, Hengyong Yu, and Bruno De Man. An outlook on X-ray CT research and development. *Medical Physics*, 35(3):1051–1064, 2008. 1
- [48] Lin Wang and Kuk-Jin Yoon. Knowledge distillation and student-teacher learning for visual intelligence: A review and new outlooks. *IEEE transactions on pattern analysis and machine intelligence*, 44(6):3048–3068, 2021. 3
- [49] Wenjun Xia, Ziyuan Yang, Zexin Lu, Zhongxian Wang, and Yi Zhang. Regformer: A local-nonlocal regularization-based model for sparse-view CT reconstruction. *IEEE Transactions on Radiation and Plasma Medical Sciences*, pages 1–1, 2023. 3
- [50] Wenbin Xie, Dehua Song, Chang Xu, Chunjing Xu, Hui Zhang, and Yunhe Wang. Learning frequency-aware dynamic network for efficient super-resolution. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 4308–4317, 2021. 5
- [51] Moran Xu, Dianlin Hu, Fulin Luo, Fenglin Liu, Shaoyu Wang, and Weiwen Wu. Limited-angle X-ray CT reconstruction using image gradient ℓ_0 -norm with dictionary learning. *IEEE Transactions on Radiation and Plasma Medical Sciences*, 5:78–87, 2021. 2
- [52] Sidky Emil Y., Chien-Min Kao, and Xiaochuan Pan. Accurate image reconstruction from few-views and limited-angle

- data in divergent-beam ct. *Journal of X-Ray Science and Technology*, 14(2):119–139, 2006. 2
- [53] Ke Yan, Xiaosong Wang, Le Lu, and Ronald M. Summers. DeepLesion: Automated mining of large-scale lesion annotations and universal lesion detection with deep learning. *Journal of Medical Imaging*, 5(3):036501, 2018. 6
- [54] Yanchao Yang and Stefano Soatto. FDA: Fourier domain adaptation for semantic segmentation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 4085–4095, 2020. 3, 5
- [55] Jaejun Yoo, Youngjung Uh, Sanghyuk Chun, Byeongkyu Kang, and Jung-Woo Ha. Photorealistic style transfer via wavelet transforms. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 9036–9045, 2019. 3
- [56] Jingyi Zhang, Jiaying Huang, Zichen Tian, and Shijian Lu. Spectral unsupervised domain adaptation for visual recognition. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 9829–9840, 2022. 3
- [57] Yi Zhang, Hu Chen, Wenjun Xia, Yang Chen, Baodong Liu, Yan Liu, Huaiqiang Sun, and Jiliu Zhou. Learn++: recurrent dual-domain reconstruction network for compressed sensing CT. *IEEE Transactions on Radiation and Plasma Medical Sciences*, 7(2):132–142, 2022. 3
- [58] Zhicheng Zhang, Xiaokun Liang, Xu Dong, Yaoqin Xie, and Guohua Cao. A sparse-view CT reconstruction method based on combination of DenseNet and deconvolution. *IEEE Transactions on Medical Imaging*, 37(6):1407–1417, 2018. 2, 6, 7, 8
- [59] Yijie Zhong, Bo Li, Lv Tang, Senyun Kuang, Shuang Wu, and Shouhong Ding. Detecting camouflaged object in frequency domain. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 4504–4513, 2022. 3
- [60] Wang Zhou, Alan Conrad Bovik, Hamid Rahim Sheikh, and Eero P Simoncelli. Image quality assessment: from error visibility to structural similarity. *IEEE Transactions on Image Processing*, 13(4), 2004. 6