ASCON: Anatomy-aware Supervised Contrastive Learning Framework for Low-dose CT Denoising

Zhihao Chen¹, Qi Gao¹, Yi Zhang², and Hongming Shan¹(⊠)

¹ Institute of Science and Technology for Brain-Inspired Intelligence, Fudan University, Shanghai, China hmshan@fudan.edu.cn

 $^{2}\,$ School of Cyber Science and Engineering, Sichuan University, Chengdu, China

Abstract. While various deep learning methods have been proposed for low-dose computed tomography (CT) denoising, most of them leverage the normal-dose CT images as the ground-truth to supervise the denoising process. These methods typically ignore the inherent correlation within a single CT image, especially the anatomical semantics of human tissues, and lack the interpretability on the denoising process. In this paper, we propose a novel Anatomy-aware Supervised CONtrastive learning framework, termed ASCON, which can explore the anatomical semantics for low-dose CT denoising while providing anatomical interpretability. The proposed ASCON consists of two novel designs: an efficient self-attention-based U-Net (ESAU-Net) and a multi-scale anatomical contrastive network (MAC-Net). First, to better capture globallocal interactions and adapt to the high-resolution input, an efficient ESAU-Net is introduced by using a channel-wise self-attention mechanism. Second, MAC-Net incorporates a patch-wise non-contrastive module to capture inherent anatomical information and a pixel-wise contrastive module to maintain intrinsic anatomical consistency. Extensive experimental results on two public low-dose CT denoising datasets demonstrate superior performance of ASCON over state-of-the-art models. Remarkably, our ASCON provides anatomical interpretability for low-dose CT denoising for the first time. Source code is available at https://github.com/hao1635/ASCON.

Keywords: CT denoising \cdot Deep learning \cdot Self-attention \cdot Contrastive learning \cdot Anatomical semantics.

1 Introduction

With the success of deep learning in the field of computer vision and image processing, many deep learning-based methods have been proposed and achieved promising results in low-dose CT (LDCT) denoising [1,4-6,9,12,23,24,26]. Typically, they employ a supervised learning setting, which involves a set of image pairs, LDCT images and their normal-dose CT (NDCT) counterparts. These methods typically use a pixel-level loss (*e.g.* mean squared error or MSE), which can cause over-smoothing problems.

To address this issue, a few studies [23,26] used a structural similarity (SSIM) loss or a perceptual loss [11]. However, they all perform in a sample-to-sample manner and ignore the inherent anatomical semantics, which could blur details in areas with low noise levels. Previous studies have shown that the level of noise in CT images varies depending on the type of tissues [17]; see an example in Fig. S1 in Supplementary Materials. Therefore, it is crucial to characterize the anatomical semantics for effectively denoising diverse tissues.

In this paper, we focus on taking advantage of the inherent anatomical semantics in LDCT denoising from a contrastive learning perspective [7, 25, 27]. To this end, we propose a novel Anatomy-aware Supervised CONtrastive learning framework (ASCON), which consists of two novel designs: an efficient selfattention-based U-Net (ESAU-Net) and a multi-scale anatomical contrastive network (MAC-Net). First, to ensure that MAC-Net can effectively extract anatomical information, diverse global-local contexts and a larger input size are necessary. However, operations on full-size CT images with self-attention are computationally unachievable due to potential GPU memory limitations [20]. To address this limitation, we propose an ESAU-Net that utilizes a channel-wise self-attention mechanism [2, 22, 28] which can efficiently capture both local and global contexts by computing cross-covariance across feature channels.

Second, to exploit inherent anatomical semantics, we present the MAC-Net that employs a disentangled U-shaped architecture [25] to produce global and local representations. Globally, a patch-wise non-contrastive module is designed to select neighboring patches with similar semantic context as positive samples and align the same patches selected in denoised CT and NDCT which share the same anatomical information, using an optimization method similar to the BYOL method [7]. This is motivated by the prior knowledge that adjacent patches often share common semantic contexts [27]. Locally, to further improve the anatomical consistency between denoised CT and NDCT, we introduce a pixel-wise contrastive module with a hard negative sampling strategy [21], which randomly selects negative samples from the pixels with high similarity around the positive sample within a certain distance. Then we use a local InfoNCE loss [18] to pull the positive pairs and push the negative pairs.

Our contributions are summarized as follows. 1) We propose a novel ASCON framework to explore inherent anatomical information in LDCT denoising, which is important to provide interpretability for LDCT denoising. 2) To better explore anatomical semantics in MAC-Net, we design an ESAU-Net, which utilizes a channel-wise self-attention mechanism to capture both local and global contexts. 3) We propose a MAC-Net that employs a disentangled U-shaped architecture and incorporates both global non-contrastive and local contrastive modules. This enables the exploitation of inherent anatomical semantics at the patch level, as well as improving anatomical consistency at the pixel level. 4) Extensive experimental results demonstrate that our ASCON outperforms other state-ofthe-art methods, and provides anatomical interpretability for LDCT denoising.



Fig. 1. Overview of our proposed ASCON. (a) Efficient self-attention-based U-Net (ESAU-Net); and (b) multi-scale anatomical contrastive network (MAC-Net).

2 Methodology

2.1 Overview of the proposed ASCON

Fig. 1 presents the overview of the proposed ASCON, which consists of two novel components: ESAU-Net and MAC-Net. First, given an LDCT image, $\boldsymbol{X} \in \mathbb{R}^{1 \times H \times W}$, where $H \times W$ denotes the image size. \boldsymbol{X} is passed through the ESAU-Net to capture both global and local contexts using a channel-wise self-attention mechanism and obtain a denoised CT image $\boldsymbol{Y}' \in \mathbb{R}^{1 \times H \times W}$.

Then, to explore inherent anatomical semantics and remain inherent anatomical consistency, the denoised CT Y' and NDCT Y are passed to the MAC-Net to compute a global MSE loss \mathcal{L}_{global} in a patch-wise non-contrastive module and a local infoNCE loss \mathcal{L}_{local} in a pixel-wise contrastive module. During training, we use an alternate learning strategy to optimize ESAU-Net and MAC-Net separately, which is similar to GAN-based methods [10]. Please refer to Alg. 1 in Supplementary Materials for a detailed optimization.

2.2 Efficient Self-attention-based U-Net

To better leverage anatomical semantic information in MAC-Net and adapt to the high-resolution input, we design the ESAU-Net that can capture both local and global contexts during denoising. Different from previous works that only

use self-attention in the coarsest level [20], we incorporate a channel-wise selfattention mechanism [2, 28] at each up-sampling and down-sampling level in the U-Net [22] and add an identity mapping in each level, as shown in Fig. 1(a).

Specifically, in each level, given the feature map F_{l-1} as the input, we first apply a 1×1 convolution and a 3×3 depth-wise convolution to aggregate channelwise contents and generate query (Q), key (K), and value (V) followed by a reshape operation, where $Q \in \mathbb{R}^{C \times HW}$, $K \in \mathbb{R}^{C \times HW}$, and $V \in \mathbb{R}^{C \times HW}$ (see Fig. 1(a)). Then, a channel-wise attention map $A \in \mathbb{R}^{C \times C}$ is generated through a dot-product operation by the reshaped query and key, which is more efficient than the regular attention map of size $HW \times HW$ [3], especially for high-resolution input. Overall, the process is defined as

Attention
$$(\mathbf{F}) = w(\mathbf{V}^T \mathbf{A}) = w(\mathbf{V}^T \cdot \text{Softmax}(\mathbf{K}\mathbf{Q}^T/\alpha)),$$
 (1)

where $w(\cdot)$ first reshapes the matrix back to the original size $C \times H \times W$ and then performs 1×1 convolution; α is a learnable parameter to scale the magnitude of the dot product of \mathbf{K} and \mathbf{Q} . We use multi-head attention similar to the standard multi-head self-attention mechanism [3]. The output of the channelwise self-attention is represented as: $\mathbf{F}'_{l-1} = \text{Attention}(\mathbf{F}_{l-1}) + \mathbf{F}_{l-1}$. Finally, the output \mathbf{F}_l of each level is defined as: $\mathbf{F}_l = \text{Conv}(\mathbf{F}'_{l-1}) + \text{Iden}(\mathbf{F}_{l-1})$, where $\text{Conv}(\cdot)$ is a two-layer convolution and $\text{Iden}(\cdot)$ is an identity mapping using a 1×1 convolution; refer to Fig. S2(a) for the details of ESAU-Net.

2.3 Multi-scale Anatomical Contrastive Network

Overview of MAC-Net. The goal of our MAC-Net is to exploit anatomical semantics and maintain anatomical embedding consistency, First, a disentangled U-shaped architecture [22] is utilized to learn global representation $F_{\rm g} \in \mathbb{R}^{512 \times \frac{H}{16} \times \frac{W}{16}}$ after four down-sampling layers, and learn local representation $F_{\rm g} \in \mathbb{R}^{64 \times H \times W}$ by removing the last output layer. And we cut the connection between the coarsest feature and its upper level to make $F_{\rm g}$ and $F_{\rm l}$ more independent [25] (see Fig. S2(b)). The online network and the target network, both using the same architecture above, handle denoised CT Y' and NDCT Y, respectively, with $F'_{\rm g}$ and $F'_{\rm l}$ generated by the online network, and $F_{\rm g}$ and $F_{\rm l}$ generated by the target network (see Fig. 1(b)). The parameters of the target network, following the previous works [7,8]. Next, a patch-wise non-contrastive module uses $F_{\rm g}$ and $F'_{\rm l}$ to compute a global MSE loss $\mathcal{L}_{\rm global}$, while a pixel-wise contrastive module uses $F_{\rm l}$ and $F'_{\rm l}$ to compute a local infoNCE loss $\mathcal{L}_{\rm local}$. Let us describe these two loss functions specifically.

Patch-wise non-contrastive module. To better learn anatomical representations, we introduce a patch-wise non-contrastive module, also shown in Fig. 1(b). Specifically, for each pixel $f_g^{(i)} \in \mathbb{R}^{512}$ in the F_g where $i \in \{1, 2, \ldots, \frac{HW}{256}\}$ is the index of the pixel location, it can be considered as a patch due to the expanded receptive field achieved through a sequence of convolutions and downsampling operations [19]. To identify positive patch indices, we adopt a neighboring positive matching strategy [27], assuming that a semantically similar patch $f_{\rm g}^{(j)}$ exists in the vicinity of the query patch $f_{\rm g}^{(i)}$, as neighboring patches often share a semantic context with the query. We empirically consider a set of 8 neighboring patches. To sample patches with similar semantics around the query patch $f_{\rm g}^{(i)}$, we measure the semantic closeness between the query patch $f_{\rm g}^{(i)}$ and its neighboring patches $f_{\rm g}^{(j)}$ using the cosine similarity, which is formulated as

$$s(i,j) = (\mathbf{f}_{g}^{(i)})^{\top} (\mathbf{f}_{g}^{(j)}) / \|\mathbf{f}_{g}^{(i)}\|_{2} \|\mathbf{f}_{g}^{(j)}\|_{2}.$$
 (2)

We then select the top-4 positive patches $\{f_{g}^{(j)}\}_{j\in\mathcal{P}^{(i)}}$ based on s(i, j), where $\mathcal{P}^{(i)}$ is a set of selected patches (*i.e.*, $|\mathcal{P}^{(i)}| = 4$). To obtain patch-level features $g^{(i)} \in \mathbb{R}^{512}$ for each patch $f_{g}^{(i)}$ and its positive neighbors, we aggregate their features using global average pooling (GAP) in the patch dimension. For the local representation of $f'_{g}^{(i)}$, we select positive patches as same as $\mathcal{P}^{(i)}$, *i.e.*, $\{f'_{g}^{(j)}\}_{j\in\mathcal{P}^{(i)}}$. Formally,

$$\boldsymbol{g}^{(i)} := \text{GAP}(\boldsymbol{f}_{g}^{(i)}, \{\boldsymbol{f}_{g}^{(j)}\}_{j \in \mathcal{P}^{(i)}}), \quad \boldsymbol{g}'^{(i)} := \text{GAP}(\boldsymbol{f}'_{g}^{(i)}, \{\boldsymbol{f}'_{g}^{(j)}\}_{j \in \mathcal{P}^{(i)}}).$$
(3)

From the patch-level features, the online network outputs a projection $\mathbf{z}_{g}^{\prime (i)} = p_{g}^{\prime}(\mathbf{g}^{\prime (i)})$ and a prediction $q^{\prime}(\mathbf{z}_{g}^{\prime (i)})$ while target network outputs the target projection $\mathbf{z}_{g}^{(i)} = p_{g}(\mathbf{g}^{(i)})$. The projection and prediction are both multi-layer perceptron (MLP). Finally, we compute the global MSE loss between the normalized prediction and target projection [7],

$$\mathcal{L}_{\text{global}} = \sum_{i \in \mathcal{N}_{\text{pos}}^{\text{g}}} \left\| q'(\mathbf{z}_{\text{g}}^{\prime(i)}) - \mathbf{z}_{\text{g}}^{(i)} \right\|_{2}^{2} = \sum_{i \in \mathcal{N}_{\text{pos}}^{\text{g}}} 2 - 2 \cdot \frac{\langle q'(\mathbf{z}_{\text{g}}^{(i)}), \mathbf{z}_{\text{g}}^{(i)} \rangle}{\|q'(\mathbf{z}_{\text{g}}^{(i)})\|_{2} \cdot \|\mathbf{z}_{\text{g}}^{(i)}\|_{2}},$$
(4)

where $\mathcal{N}_{\text{pos}}^{\text{g}}$ is the indices set of positive samples in the patch-level embedding. **Pixel-wise contrastive module.** In this module, we aim to improve anatomical consistency between the denoised CT and NDCT using a local InfoNCE loss [18] (see Fig. 1(b)). First, for a query $f_1^{\prime(i)} \in \mathbb{R}^{64}$ in the F_1^{\prime} and its positive sample $f_1^{(i)} \in \mathbb{R}^{64}$ in the F_1 $(i \in \{1, 2, \ldots, HW\}$ is the location index), we use a hard negative sampling strategy [21] to select "diffcult" negative samples with high probability, which enforces the model to learn more from the finegrained details. Specifically, candidate negative samples are randomly sampled from F_1 as long as their distance from $f_1^{(i)}$ is less than m pixels (m=7). We also use cosine similarity in Eq. (2) to select a set of semantically closest pixels, *i.e.* $\{f_1^{(j)}\}_{j \in \mathcal{N}_{neg}^{(i)}}$. Then we concatenate $f_1^{\prime(i)}, f_1^{(i)}$, and $\{f_1^{(j)}\}_{j \in \mathcal{N}_{neg}^{(i)}}$ and map them to a K-dimensional vector (K=256) through a two-layer MLP, obtaining The local InfoNCE loss in the pixel level is defined as

$$\mathcal{L}_{\text{local}} = -\sum_{i \in \mathcal{N}_{\text{pos}}^{\text{l}}} \log \frac{\exp\left(\boldsymbol{v}_{1}^{\prime(i)} \cdot \boldsymbol{v}_{1}^{(i)} / \tau\right)}{\exp\left(\boldsymbol{v}_{1}^{\prime(i)} \cdot \boldsymbol{v}_{1}^{(i)} / \tau\right) + \sum_{j=1}^{|\mathcal{N}_{\text{nes}}^{(i)}|} \exp\left(\boldsymbol{v}_{1}^{\prime(i)} \cdot \boldsymbol{v}_{1}^{(j)} / \tau\right)},$$
(5)

where $\mathcal{N}_{\text{pos}}^{\text{l}}$ is the indices set of positive samples in the pixel level. $\boldsymbol{v}_{1}^{(i)}, \boldsymbol{v}_{l}^{(i)}$, and $\boldsymbol{v}_{l}^{(j)} \in \mathbb{R}^{256}$ are the query, positive, and negative sample in $\boldsymbol{z}_{l}^{(i)}$, respectively.

2.4 Total Loss Function

The final loss is defined as $\mathcal{L} = \mathcal{L}_{global} + \mathcal{L}_{local} + \lambda \mathcal{L}_{pixel}$, where \mathcal{L}_{pixel} consists of two common supervised losses: MSE and SSIM, defined as $\mathcal{L}_{pixel} = \mathcal{L}_{MSE} + \mathcal{L}_{SSIM}$. λ is empirically set to 10.

3 Experiments

3.1 Dataset and Implementation Details

We use two publicly available low-dose CT datasets released by the NIH AAPM-Mayo Clinic Low-Dose CT Grand Challenge in 2016 [15] and lately released in 2020 [16], denoted as Mayo-2016 and Mayo-2020, respectively. There is no overlap between the two datasets. Mayo-2016 includes normal-dose abdominal CT images of 10 anonymous patients and corresponding simulated quarter-dose CT images. Mayo-2020 provides the abdominal CT image data of 100 patients with 25% of the normal dose, and we randomly choose 20 patients for our experiments.

For the Mayo-2016, we choose 4800 pairs of 512×512 images from 8 patients for the training and 1136 pairs from the rest 2 patients as the test set. For the Mayo-2020, we employ 9600 image pairs with a size of 256×256 from randomly selected 16 patients for training and 580 pairs of 512×512 images from rest 4 patients for testing. The use of large-size training is to adapt our MAC-Net to exploit inherent semantic information. The default sampling hyper-parameters for Mayo-2016 are $|\mathcal{N}_{\rm pos}^{\rm l}| = 32$, $|\mathcal{N}_{\rm pos}^{\rm g}| = 512$, $|\mathcal{N}_{\rm neg}^{(i)}| = 24$, while $|\mathcal{N}_{\rm pos}^{\rm l}| = 16$, $|\mathcal{N}_{\rm pos}^{\rm g}| = 256$, $|\mathcal{N}_{\rm neg}^{(i)}| = 24$ for Mayo-2020. We use a binary function to filter the background while selecting queries in MAC-Net. For the training strategy, we employ a window of [-1000, 2000] HU. We train our network for 100 epochs on 2 NVIDIA GeForce RTX 3090, and use the AdamW optimizer [14] with the momentum parameters $\beta_1 = 0.9$, $\beta_2 = 0.99$ and the weight decay of 1.0×10^{-9} . We initialize the learning rate as 1.0×10^{-4} , gradually reduced to 1.0×10^{-6} with the cosine annealing [13]. Since MAC-Net is only implemented during training, the testing time of ASCON is close to most of the compared methods.

3.2 Performance Comparisons

Quantitative evaluations. We use three widely-used metrics such as peak signal-to-noise ratio (PSNR), root-mean-square error (RMSE), and SSIM. Table 1 presents the testing results on Mayo-2016 and Mayo-2020 datasets. We compare our methods with 5 state-of-the-art methods, including RED-CNN [1], WGAN-VGG [26], EDCNN [12], DU-GAN [9], and CNCL [6]. Table 1 shows that our ESAU-Net with MAC-Net achieves the best performance on both the Mayo-2016 and the Mayo-2020 datasets. Compared to the ESAU-Net, ASCON further improves the PSNR by up to 0.54 dB on Mayo-2020, which demonstrates the effectiveness of the proposed MAC-Net and the importance of the inherent

Methods	Mayo-2016			Mayo-2020			
	$PSNR\uparrow$	RMSE↓	$SSIM\uparrow$	$PSNR\uparrow$	RMSE↓	$SSIM\uparrow$	
U-Net [22]	$44.13_{\pm 1.19}$	$0.64{\scriptstyle \pm 0.12}$	97.38 ± 1.09	47.67 ± 1.64	0.43 ± 0.09	$99.19{\scriptstyle \pm 0.23}$	
RED-CNN [1]	44.23 ± 1.26	0.62 ± 0.09	97.34 ± 0.86	48.05 ± 2.14	0.41 ± 0.11	99.28 ± 0.18	
WGAN-VGG [26]	42.49 ± 1.28	$0.76{\scriptstyle \pm 0.12}$	96.16 ± 1.30	46.88 ± 1.81	0.46 ± 0.10	98.15 ± 0.20	
EDCNN [12]	43.14 ± 1.27	$0.70_{\pm 0.11}$	96.45 ± 1.36	47.90 ± 1.27	0.41 ± 0.08	99.14 ± 0.17	
DU-GAN [9]	43.06 ± 1.22	0.71 ± 0.10	96.34 ± 1.12	47.21 ± 1.52	0.44 ± 0.10	$99.00{\scriptstyle \pm 0.21}$	
CNCL [6]	43.06 ± 1.07	0.71 ± 0.10	$96.68{\scriptstyle\pm1.11}$	45.63 ± 1.34	$0.53 \scriptstyle \pm 0.11$	$98.92{\scriptstyle \pm 0.59}$	
ESAU-Net (ours)	$44.38_{\pm 1.26}$	0.61 ± 0.09	97.47 ± 0.87	$48.31_{\pm 1.87}$	0.40 ± 0.12	$\underline{99.30}_{\pm 0.18}$	
ASCON (ours)	$44.48{\scriptstyle \pm 1.32}$	$0.60 \scriptstyle \pm 0.10$	$97.49 \scriptstyle \pm 0.86$	$48.84{\scriptstyle \pm 1.68}$	$0.37_{\pm 0.11}$	$99.32{\scriptstyle \pm 0.18}$	

Table 1. Performance comparison on the Mayo-2016 and Mayo-2020 datasets in terms of PSNR [dB], RMSE [$\times 10^{-2}$], and SSIM [%]



Fig. 2. Transverse CT images and corresponding difference images from the Mayo-2016 dataset: (a) NDCT; (b) LDCT; (c) RED-CNN [1]; (d) EDCNN [12]; (e) DU-GAN [9]; (f) ESAU-Net (ours); and (g) ASCON (ours). The display window is [-160, 240] HU.

anatomical semantics during CT denoising. We also compute the contrast-tonoise ratio (CNR) to assess the detectability of a selected area of low-contrast lesion and our ASCON achieves the best CNR in Fig. S3.

Qualitative evaluations. Fig. 2 presents qualitative results of three representative methods and our ESAU-Net with MAC-Net on Mayo-2016. Although ASCON and RED-CNN produce visually similar results in low-contrast areas after denoising. However, RED-CNN results in blurred edges between different tissues, such as the liver and blood vessels, while ASCON smoothed the noise and maintained the sharp edges. They are marked by arrows in the regions-of-interest images. We further visualize the corresponding difference images between NDCT and the generated images by our method as well as other methods as shown in the third row of Fig. 2. Note that our ASCON removes more noise components than other methods; refer to Fig. S4 for extra qualitative results on Mayo-2020. Visualization of inherent semantics. To demonstrate that our MAC-Net can exploit inherent anatomical semantics of CT images during denoising, we se-



Fig. 3. Visualization of inherent semantics; (a) NDCT; (b) clustering and t-SNE results of ASCON w/o MAC-Net; and (c) clustering and t-SNE results of ASCON.

Table 2. Ablation results of Mayo-2020 on the different types of loss functions.

Loss	$\mathrm{PSNR}\uparrow$	RMSE↓	$\mathrm{SSIM}\uparrow$
$ \begin{array}{c} \hline \\ \mathcal{L}_{\mathrm{MSE}} \\ \mathcal{L}_{\mathrm{MSE}} + \mathcal{L}_{\mathrm{Perceptual}} \\ \mathcal{L}_{\mathrm{MSE}} + \mathcal{L}_{\mathrm{SSIM}} \\ \mathcal{L}_{\mathrm{MSE}} + \mathcal{L}_{\mathrm{SSIM}} + \mathcal{L}_{\mathrm{global}} \\ \mathcal{L}_{\mathrm{MSE}} + \mathcal{L}_{\mathrm{SSIM}} + \mathcal{L}_{\mathrm{local}} \\ \mathcal{L}_{\mathrm{MSE}} + \mathcal{L}_{\mathrm{SSIM}} + \mathcal{L}_{\mathrm{local}} + \mathcal{L}_{\mathrm{global}} \end{array} $	$\begin{array}{c} 48.34{\scriptstyle\pm2.22}\\ 47.83{\scriptstyle\pm1.99}\\ 48.31{\scriptstyle\pm1.87}\\ 48.58{\scriptstyle\pm2.12}\\ 48.48{\scriptstyle\pm2.37}\\ \textbf{48.84}{\scriptstyle\pm1.68}\end{array}$	$\begin{array}{c} 0.40{\scriptstyle\pm0.11}\\ 0.42{\scriptstyle\pm0.10}\\ 0.40{\scriptstyle\pm0.12}\\ 0.39{\scriptstyle\pm0.10}\\ 0.38{\scriptstyle\pm0.11}\\ \textbf{0.37}{\scriptstyle\pm0.11} \end{array}$	$\begin{array}{r} 99.27{\scriptstyle\pm0.18}\\ 99.13{\scriptstyle\pm0.19}\\ 99.30{\scriptstyle\pm0.18}\\ 99.31{\scriptstyle\pm0.17}\\ 99.31{\scriptstyle\pm0.18}\\ \textbf{99.32}{\scriptstyle\pm0.18}\end{array}$

lect the features before the last layer in ASCON without MAC-Net and ASCON from Mayo-2016. Then we cluster these two feature maps respectively using a K-means algorithm and visualize them in the original dimension, and finally visualize the clustering representations using t-SNE, as shown in Fig. 3. Note that ASCON produces a result similar to organ semantic segmentation after clustering and the intra-class distribution is more compact, as well as the inter-class separation is more obvious. To the best of our knowledge, this is the first time that anatomical semantic information has been demonstrated in a CT denoising task, providing interpretability to the field of medical image reconstruction.

Ablation studies. We start with a ESAU-Net using MSE loss and gradually insert some loss functions and our MAC-Net. Table 2 presents the results of different loss functions. It shows that both the global non-contrastive module and local contrastive module are helpful in obtaining better metrics due to the capacity of exploiting inherent anatomical information and maintaining anatomical consistency. Then, we add our MAC-Net to two supervised models: RED-CNN [1] and U-Net [22] but it is less effective, which demonstrates the importance of our ESAU-Net that captures both local and global contexts during denoising in Table S1. In addition, we evaluate the effectiveness of the training strategies including alternate learning, neighboring positive matching and hard negative sampling in Table S2.

4 Conclusion

In this paper, we explore the anatomical semantics in LDCT denoising and take advantage of it to improve the denoising performance. To this end, we propose an Anatomy-aware Supervised CONtrastive learning framework (AS-CON), consisting of an efficient self-attention-based U-Net (ESAU-Net) and a

multi-scale anatomical contrastive network (MAC-Net), which can capture both local and global contexts during denoising and exploit inherent anatomical information. Extensive experimental results on Mayo-2016 and Mayo-2020 datasets demonstrate the superior performance of our method, and the effectiveness of our designs. We also validated that our method introduces interpretability to LDCT denoising.

References

- 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 Trans. Med. Imaging 36(12), 2524–2535 (2017)
- Chen, Z., Niu, C., Wang, G., Shan, H.: LIT-Former: Linking in-plane and throughplane transformers for simultaneous CT image denoising and deblurring. arXiv preprint arXiv:2302.10630 (2023)
- Dosovitskiy, A., Beyer, L., Kolesnikov, A., Weissenborn, D., Zhai, X., Unterthiner, T., Dehghani, M., Minderer, M., Heigold, G., Gelly, S., et al.: An image is worth 16x16 words: Transformers for image recognition at scale. arXiv preprint arXiv:2010.11929 (2020)
- Gao, Q., Li, Z., Zhang, J., Zhang, Y., Shan, H.: CoreDiff: Contextual errormodulated generalized diffusion model for low-dose CT denoising and generalization. arXiv preprint arXiv:2304.01814 (2023)
- Gao, Q., Shan, H.: CoCoDiff: a contextual conditional diffusion model for low-dose CT image denoising. In: Developments in X-Ray Tomography XIV. vol. 12242. SPIE (2022)
- Geng, M., Meng, X., Yu, J., Zhu, L., Jin, L., Jiang, Z., Qiu, B., Li, H., Kong, H., Yuan, J., et al.: Content-noise complementary learning for medical image denoising. IEEE Trans. Med. Imaging 41(2), 407–419 (2021)
- Grill, J.B., Strub, F., Altché, F., Tallec, C., Richemond, P., Buchatskaya, E., Doersch, C., Avila Pires, B., Guo, Z., Gheshlaghi Azar, M., et al.: Bootstrap your own latent-a new approach to self-supervised learning. Proc. Adv. Neural Inf. Process. Syst. 33, 21271–21284 (2020)
- He, K., Fan, H., Wu, Y., Xie, S., Girshick, R.: Momentum contrast for unsupervised visual representation learning. In: Proc. IEEE Conf. Comput. Vis. Pattern Recognit. pp. 9729–9738 (2020)
- 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 Trans. Instrum. Meas. 71, 1–12 (2021)
- Isola, P., Zhu, J.Y., Zhou, T., Efros, A.A.: Image-to-image translation with conditional adversarial networks. In: Proc. IEEE Conf. Comput. Vis. Pattern Recognit. pp. 1125–1134 (2017)
- Johnson, J., Alahi, A., Fei-Fei, L.: Perceptual losses for real-time style transfer and super-resolution. In: Proc. European Conf. Comput. Vis. pp. 694–711. Springer (2016)
- Liang, T., Jin, Y., Li, Y., Wang, T.: EDCNN: Edge enhancement-based densely connected network with compound loss for low-dose CT denoising. In: 2020 15th IEEE Int. Conf. Signal Process. vol. 1, pp. 193–198. IEEE (2020)
- Loshchilov, I., Hutter, F.: SGDR: Stochastic gradient descent with warm restarts. arXiv preprint arXiv:1608.03983 (2016)

- 10 Z. Chen et al.
- Loshchilov, I., Hutter, F.: Decoupled weight decay regularization. arXiv preprint arXiv:1711.05101 (2017)
- McCollough, C.H., Bartley, A.C., Carter, R.E., Chen, B., Drees, T.A., Edwards, P., Holmes III, D.R., Huang, A.E., Khan, F., Leng, S., et al.: Low-dose CT for the detection and classification of metastatic liver lesions: results of the 2016 low dose CT grand challenge. Med. Phys. 44(10), e339–e352 (2017)
- Moen, T.R., Chen, B., Holmes III, D.R., Duan, X., Yu, Z., Yu, L., Leng, S., Fletcher, J.G., McCollough, C.H.: Low-dose CT image and projection dataset. Med. Phys. 48(2), 902–911 (2021)
- Mussmann, B.R., Mørup, S.D., Skov, P.M., Foley, S., Brenøe, A.S., Eldahl, F., Jørgensen, G.M., Precht, H.: Organ-based tube current modulation in chest CT. a comparison of three vendors. Radiography 27(1), 1–7 (2021)
- Oord, A.v.d., Li, Y., Vinyals, O.: Representation learning with contrastive predictive coding. arXiv preprint arXiv:1807.03748 (2018)
- Park, T., Efros, A.A., Zhang, R., Zhu, J.Y.: Contrastive learning for unpaired image-to-image translation. In: Proc. European Conf. Comput. Vis., Glasgow, UK, August 23–28, 2020, Proceedings, Part IX 16. pp. 319–345. Springer (2020)
- Petit, O., Thome, N., Rambour, C., Themyr, L., Collins, T., Soler, L.: U-net transformer: Self and cross attention for medical image segmentation. In: Mach.Learn. Med. Imag. pp. 267–276. Springer (2021)
- Robinson, J., Chuang, C.Y., Sra, S., Jegelka, S.: Contrastive learning with hard negative samples. arXiv preprint arXiv:2010.04592 (2020)
- Ronneberger, O., Fischer, P., Brox, T.: U-net: Convolutional networks for biomedical image segmentation. In: Med. Image Comput. Comput. Assist. Interv. pp. 234–241. Springer (2015)
- Shan, H., Padole, A., Homayounieh, F., Kruger, U., Khera, R.D., Nitiwarangkul, C., Kalra, M.K., Wang, G.: Competitive performance of a modularized deep neural network compared to commercial algorithms for low-dose CT image reconstruction. Nat. Mach. Intell. 1(6), 269–276 (2019)
- Shan, H., Zhang, Y., Yang, Q., Kruger, U., Kalra, M.K., Sun, L., Cong, W., Wang, G.: 3-D convolutional encoder-decoder network for low-dose CT via transfer learning from a 2-D trained network. IEEE Trans. Med. Imaging 37(6), 1522–1534 (2018)
- Yan, K., Cai, J., Jin, D., Miao, S., Guo, D., Harrison, A.P., Tang, Y., Xiao, J., Lu, J., Lu, L.: SAM: Self-supervised learning of pixel-wise anatomical embeddings in radiological images. IEEE Trans. Med. Imaging 41(10), 2658–2669 (2022)
- Yang, Q., Yan, P., Zhang, Y., Yu, H., Shi, Y., Mou, X., Kalra, M.K., Zhang, Y., Sun, L., Wang, G.: Low-dose CT image denoising using a generative adversarial network with Wasserstein distance and perceptual loss. IEEE Trans. Med. Imaging 37(6), 1348–1357 (2018)
- Yun, S., Lee, H., Kim, J., Shin, J.: Patch-level representation learning for selfsupervised vision transformers. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 8354–8363 (2022)
- Zamir, S.W., Arora, A., Khan, S., Hayat, M., Khan, F.S., Yang, M.H.: Restormer: Efficient transformer for high-resolution image restoration. In: Proc. European Conf. Comput. Vis. pp. 5728–5739 (2022)

Supplementary Materials



Fig. S1. The mean and standard deviation of different tissues in an example of an LDCT image. The standard deviation (std) of an ROI in the liver is 63.23 HU, while the std of an ROI in the muscle is 44.73 HU.



Fig. S2. Overview of (a) ESAU-Net and (b) disentangled U-shape architecture.



Fig. S3. Contrast-to-noise ratio (CNR) of the low-contrast lesion from Mayo-2016 dataset. We select the yellow rectangle of every image to visualize (see in (a)). ROIs of the red rectangle and the blue rectangle are the lesion and the background. (b) NDCT; (c) LDCT; (d) RED-CNN; (e) WGAN-VGG; (f) CNCL; (g) ESAU-Net (ours); and (h) ASCON (ours). Note that our proposed method achieves the best performance of CNR.



Fig. S4. Transverse CT images and corresponding difference images from the Mayo-2020 dataset: (a) NDCT; (b) LDCT; (c) RED-CNN; (d) WGAN-VGG; (e) CNCL; (f) ESAU-Net (ours); and (g) ASCON (ours). Zoomed ROI of the rectangle is shown below the full-size one. The display window is [-160, 240] HU.

Algorithm 1 Optimization of alternate learning

1: Input: LDCT X, NDCT Y 2: Networks in ASCON: ESAU-Net and MAC-Net 3: Optimizers: Optimizer_E (ESAU-Net) and Optimizer_M (MAC-Net) 4: 5: $\mathbf{Y}' = \text{ESAU-Net.forward}(\mathbf{X})$ 6: 7: # Optimization of MAC-Net 8: Set requires grad (MAC-Net, True) 9: MAC-Net .forward(\mathbf{Y}' .detach(), \mathbf{Y}) 10: Optimizer_M.zero_grad() 11: Compute $\mathcal{L}_{contra} = \mathcal{L}_{local} + \mathcal{L}_{global}$ 12: \mathcal{L}_{contra} .backward () 13: Optimizer_M.step () 14:15: # Optimization of ESAU-Net 16: Set_requires_grad (MAC-Net, False) 17: Optimizer_E.zero_grad () 18: Compute $\overline{\mathcal{L}} = \mathcal{L}_{pixel} + \mathcal{L}_{contra}$ 19: \mathcal{L} .backward () 20: Optimizer E.step ()

Table S1. Ablation results on different supervised models with MAC-Net. Note that the performance of RED-CNN and Unet decreases when adding a MAC-Net. It verifies the importance of our ESAU-Net that captures both local and global contexts during denoising.

	Mayo	-2016 Da	taset	Μ	ayo-2020) Datas	set
Method	PSNR↑	RMSE↓	$SSIM\uparrow$	PSN	R↑ RMS	SE↓ SS	SIM↑
RED-CNN	44.23	0.62	97.34	48.0	05 0.4	41 99	9.28
	± 1.26	± 0.09	± 0.86	$\pm 2.$	14 ± 0	$11 \pm$	0.18
RED-CNN w/ MAC-Net	43.55	0.67	97.11	47.8	85 0.4	12 9	9.27
	± 1.16	± 0.11	± 0.85	$\pm 2.$	08 ± 0	$10 \pm$	0.21
Unet	44.13	0.64	97.38	47.0	67 0. 4	13 99	9.19
	± 1.19	± 0.12	± 1.09	$\pm 1.$	64 ± 0	$.09 \pm$	0.23
Unet w/ MAC-Net	44.04	0.65	97.35	47.0	66 0.4	13 9	9.20
,	± 1.15	± 0.11	± 0.98	$\pm 1.$	82 ± 0	$10 \pm$	0.21
ESAU-Net	44.38	0.61	97.47	48.3	31 0.4	10 9	9.30
	± 1.26	± 0.09	± 0.87	$\pm 1.$	$87 \pm 0.$	$12 \pm$	0.18
ASCON (proposed)	44.48	0.60	97.49	48.	84 0.3	37 99	9.32
	± 1.32	± 0.10	± 0.86	±1.	68 ± 0	.11 ±	0.18

Table S2. Ablation results of Mayo-2020 on the different training strategies in terms of PSNR, RMSE, and SSIM.

Method	$\mathrm{PSNR}\uparrow$	$\mathrm{RMSE}{\downarrow}$	$\mathrm{SSIM}\uparrow$
w/o alternate training	48.57 ± 1.82	$0.38 {\pm} 0.10$	$99.28 {\pm} 0.19$
w/o neighboring positive matching strategy	48.51 ± 1.61	$0.39 {\pm} 0.10$	$99.31 {\pm} 0.17$
w/o hard negative sampling	48.41 ± 1.73	$0.40{\pm}0.09$	$99.29 {\pm} 0.19$
ASCON (proposed)	$\textbf{48.84}{\pm}\textbf{1.68}$	$\boldsymbol{0.37{\pm}0.11}$	$99.32{\pm}0.18$