DosePlug: Observation-Driven Dose Modeling for Robust Low-Dose Reconstruction and Segmentation

Do Hyun Ki[1], Yeong Jong Lee[1], and Seok Bong Yoo[*]

Department of Artificial Intelligence Convergence, Chonnam National University, Gwangju, Korea
sbyoo@jnu.ac.kr

Abstract. Low-dose computed tomography and low-dose positron emission tomography reduce examination time and radiation exposure, crucial for radiation-sensitive patients and repeated exams. However, dose reduction yields photon-starved measurements that amplify noise and structured artifacts, undermining reliable lesion segmentation. Although deep learning advances low-dose reconstruction, most models lack explicit dose modeling. Some rely on dose metadata, often missing or poorly aligned with actual degradation, leaving residual dose-dependent distortions. To address these limitations, we define an observation-driven dose (OD-dose) as a latent coordinate of task-relevant corruption severity and estimate it by fusing periodicity degradation cues and mutual information cues extracted from sinograms. The estimated OD-dose modulates intermediate features via dose-aware feature-wise polynomial modulation, enabling nonlinear, task-aligned dose-adaptive restoration. The proposed module is a lightweight, plug-and-play component that can be attached to diverse reconstruction backbones. Experimental results confirm that DosePlug improves reconstruction and lesion segmentation on LDCT and LDPET benchmarks, while maintaining stable gains on real-world datasets. The code is available at https://github.com/kikidohyun/DosePlug.git.

Keywords: Low-dose medical image reconstruction and segmentation · Observation-driven dose estimation · Dose-aware feature-wise polynomial modulation.

1 Introduction

Computed tomography (CT) provides high-resolution anatomy and positron emission tomography (PET) captures functional activity, but both rely on ionizing radiation, motivating the use of low-dose CT (LDCT) and low-dose PET (LDPET) to reduce exposure and shorten examinations [3,17]. Fig. 1(a) shows low-dose CT and PET images acquired by reducing exposure parameters. In practice, LDCT reduces x-ray exposure by adjusting scan protocol parameters such as the tube

[1] Equal contribution, * Corresponding author.

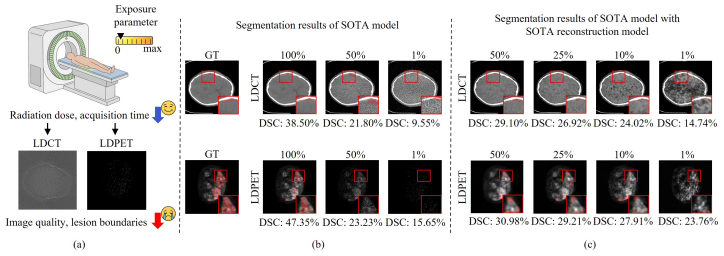

Fig. 1. (a) Illustration of LDCT and LDPET acquisition. Qualitative examples from a state-of-the-art (SOTA) segmentation model [19] and the average DSC trends across dose ratios are shown in (b) without reconstruction and in (c) with SOTA reconstruction backbones [1,24].

current-time product (mAs), while LDPET reduces counts via lower injected activity or shorter acquisition time per bed position [2,12]. Dose reduction lowers CT photon counts and PET coincidence counts, degrading counting statistics and elevating relative noise, artifacts, and contrast loss that hinder lesion segmentation [5,6,10]. Fig. 1(b) shows these degradation effects and the resulting DSC drop under dose reduction, using the SOTA segmentation model SAMIHS [19] on the CTICH [9] and AutoPET [8] datasets, which were simulated using the method detailed in the experiment section.

Recent deep learning reconstruction backbones improve visual quality under low-dose conditions, yet many apply similar processing across dose levels without explicit dose modeling, leaving residual distortions that weaken boundary and contrast cues for lesion segmentation [1,11,18,22,23,24]. Fig. 1(c) highlights this gap with VisNet [24] for CT and RED [1] for PET. Even after reconstruction, DSC gains remain limited, indicating that these residual distortions persist.

While some methods rely on physical dose metadata (e.g., protocol parameters) [14], such metadata can become inaccessible after data preprocessing and may not directly reflect degradation relevant to reconstruction and segmentation [2,20]. Therefore, we define an observation-driven dose (OD-dose), estimated from sinograms where dose-dependent distortions originate and are most directly preserved, to directly quantify image degradation. OD-dose is not a physical dose measure, but can be viewed as a latent coordinate of task-relevant corruption severity inferred directly from sinograms under Poisson perturbations.

DosePlug inserts a lightweight dose-conditioning module into reconstruction backbones, estimates OD-dose from sinogram cues, and calibrates intermediate features to enable OD-dose-adaptive restoration and more robust segmentation. The main contributions are as follows: (i) We define OD-dose to characterize dose-dependent degradation linked to task-relevant corruption severity, and estimate it from sinograms by fusing mutual information (MI) and periodicity degradation cues. (ii) We propose dose-aware feature-wise polynomial modulation (FwPM), applying an OD-dose-conditioned second-order polynomial reparameterization for stable nonlinear modulation, mitigating noise amplification as OD-dose decreases.

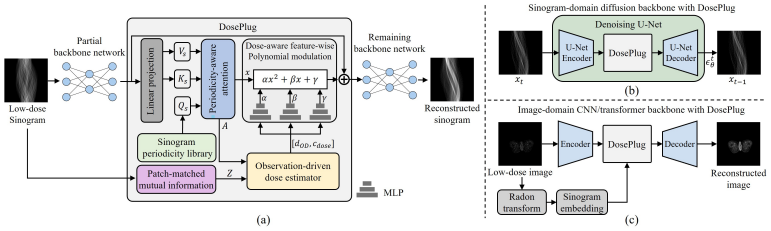

Fig. 2. (a) Overall architecture of the proposed DosePlug. Inserting DosePlug into diverse reconstruction backbones in (b) sinogram domains and (c) image domains.

(iii) We propose periodicity-aware attention that restores corrupted sinogram periodicity via a sinogram periodicity library (SPL) and provides periodicity degradation cues for OD-dose estimation.

2 Methods

DosePlug is a plug-and-play module preserving the reconstruction backbone (Fig. 2(a)). Given intermediate backbone features and the input sinogram, it estimates OD-dose and applies OD-dose-conditioned polynomial modulation. DosePlug supports diverse backbone families and domains. For sinogram-domain diffusion backbones [1], DosePlug is inserted at the denoising U-Net bottleneck (Fig. 2(b)). For image-domain CNN and transformer backbones [24] (Fig. 2(c)), the input image is converted to a sinogram via a real-time Radon transform, embedded by a CNN-based sinogram embedding layer, and fed into DosePlug for periodicity-aware attention and OD-dose-conditioned backbone feature modulation.

In CT/PET, each sinogram bin counts detected events, modeled by Poisson statistics [7,16]. Under this model, the mean count intensity λ governs noise, and the relative noise level scales approximately as $1/\sqrt{\lambda}$, decreasing as physical dose increases. Using this, the OD-dose is modeled from the input sinogram x_s as $d_{\mathrm{OD}} \triangleq a\,\sigma(x_s)^{-2}$, where $\sigma(x_s)$ denotes a noise-scale estimate computed as the standard deviation of a high-frequency/band-pass residual, and $a > 0$ is a calibration constant set on the training set to match the numerical scale of the target dose coordinate. Since λ is proportional to the normalized dose ratio, d_{OD} is monotonically related to physical dose (e.g., mAs), serving as a task-aligned latent degradation coordinate rather than a mere noise proxy.

2.1 Periodicity-aware Attention

As dose decreases, sinogram corruption becomes structured and disrupts periodic patterns, and the degree of periodicity breakdown provides an informative cue for OD-dose estimation. To restore these periodic structures while extracting an OD-dose cue, we introduce periodicity-aware attention. In Fig. 3(a), SPL is a set of learnable 2D periodic patterns initialized from multi-scale $n \times n$ patches, where

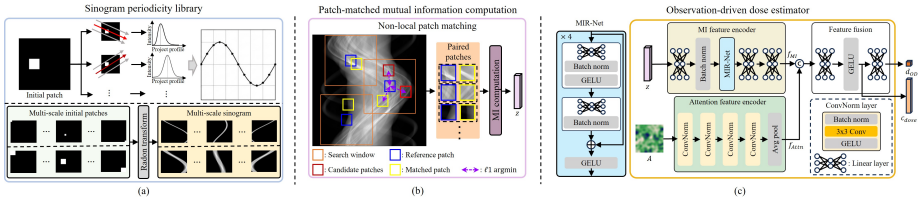

Fig. 3. Illustration of (a) the multi-scale initialization of the sinogram periodicity library, (b) window-based non-local patch matching for patch-matched MI computation, and (c) the observation-driven dose estimator.

smaller patches capture fine oscillations and larger patches capture coarse trends. During reconstruction, backbone features are projected to keys K_s and values V_s, and SPL patterns are projected to queries Q_s. Periodicity-aware attention uses a cross-attention mechanism to inject periodic cues and produce SPL-guided features x. Additionally, to explicitly capture the structural correlation, the pre-softmax query-to-key similarity defines a periodicity alignment map $A = \frac{Q_s K_s^\top}{\sqrt{d_s}}$ (where d_s is the query/key embedding dimension), which is retained as a cue for OD-dose estimation as it directly reflects the degree of periodicity breakdown.

2.2 Patch-matched Mutual Information Computation

MI from an input sinogram x_s can provide an internal cue for OD-dose estimation because dose reduction alters noise statistics and disrupts repeated structures. However, MI is sensitive to patch pairing. As shown in Fig. 3(b), we improve robustness via window-based non-local patch matching. Reference patches $\mathcal{R}_{i,j}$ of size $s \times s$ are sampled with stride s, where (i, j) denotes the top-left coordinate. For each reference patch $\mathcal{R}_{i,j}$, candidate patches within a 160×160 search window are collected as $\mathcal{C}_{i,j}$, and the matched patch $\mathcal{M}_{i,j}$ is selected from $\mathcal{C}_{i,j}$ by minimizing the ℓ_1 distance. The paired patch $\rho_{i,j}$ is composed of $(\mathcal{R}_{i,j}, \mathcal{M}_{i,j})$, and MI is computed only on these paired patches to reduce content-induced confounding. For each paired patch $\rho_{i,j}$, the MI value $\mathrm{MI}_{i,j}$ is obtained using joint and marginal histograms over the discrete bins, and is defined as

$$\mathrm{MI}_{i,j} = \sum_{x \in X} \sum_{y \in Y} p_{\rho_{i,j}}(x, y) \log \left(\frac{p_{\rho_{i,j}}(x, y)}{p_{\rho_{i,j}}(x)\, p_{\rho_{i,j}}(y) + \epsilon} \right), \tag{1}$$

where X and Y denote the discrete intensity sets of the reference and matched patches, $p_{\rho_{i,j}}(x, y)$ is the joint distribution estimated from the paired patch $\rho_{i,j}$ with marginals $p_{\rho_{i,j}}(x)$ and $p_{\rho_{i,j}}(y)$, and $\epsilon = 10^{-6}$ ensures stability. $\mathrm{MI}_{i,j}$ is computed at multiple locations and aggregated to form an MI descriptor z. To compress the dynamic range and stabilize aggregation, each $\mathrm{MI}_{i,j}$ is transformed as $\log(1 + \mathrm{MI}_{i,j})$. The descriptor z is then defined as $z = \left[\log(1 + \mathrm{MI}_{i,j})\right]_{(i,j) \in \mathcal{I}}$, where $\mathcal{I}$ denotes the index set of reference patch locations.

2.3 Observation-driven Dose Estimator

As shown in Fig. 3(c), our OD-dose estimator fuses two complementary cues: the model-agnostic MI descriptor z (prone to content bias) and the backbone-aware periodicity alignment map A (sensitive to feature fluctuations). This fusion yields a more reliable estimate than using either alone.

The estimator consists of two lightweight encoders and a fusion head. The MI encoder maps the descriptor z to an MI embedding f_{MI} using an MIRNet-based MI feature encoder. The attention encoder maps A to an attention response embedding f_{Attn} using an attention feature encoder. The two embeddings are concatenated and fed into a fusion block to regress the scalar OD-dose d_{OD}. The intermediate fused feature before the final regression layer is retained as a compact OD-dose conditioning vector c_{dose}, which summarizes the fused evidence from z and A. Both d_{OD} and c_{dose} are used for downstream dose-aware modulation.

2.4 Dose-aware Feature-wise Polynomial Modulation

Dose induces a conditional distribution shift in intermediate representations. We introduce FwPM, which formulates dose conditioning as a per-channel polynomial reparameterization driven by the low-dose degradation level. Using the predicted OD-dose d_{OD} and the conditioning vector c_{dose}, FwPM calibrates the SPL-guided features x via a dose-indexed second-order polynomial reparameterization, enabling nonlinear modulation across the degradation continuum. Because dose-conditioned feature modulation cannot be adequately expressed by a purely first-order parameterization across the degradation continuum, FwPM adopts a second-order term to model the resulting nonlinear response. The per-channel coefficients α, β, and γ respectively control the quadratic term, linear term, and offset, and are predicted by MLPs taking $[c_{\mathrm{dose}}, d_{\mathrm{OD}}]$ as input. Given the SPL-guided features x, the calibrated features $\hat{x}$ are computed as

$$\hat{x} = \alpha \odot x^2 + \beta \odot x + \gamma, \tag{2}$$

where $\odot$ denotes channel-wise multiplication. The calibrated features $\hat{x}$ are added to the input features of DosePlug through a residual pathway and re-injected into the original backbone stream. The subsequent backbone layers then operate on this result to produce the final reconstruction output.

2.5 Training Objective

DosePlug is trained by optimizing the reconstruction objective of the chosen backbone together with an OD-dose loss $\mathcal{L}_{\mathrm{OD}}$. Since the module is inserted into a reconstruction backbone, the reconstruction loss $\mathcal{L}_{\mathrm{recon}}$ follows the original backbone design. To stabilize learning over a wide dose range, $\mathcal{L}_{\mathrm{OD}}$ penalizes dose errors in the log scale using mean squared error (MSE):

$$\mathcal{L}_{\mathrm{OD}} = \left\| \log(1 + d_{\mathrm{OD}}) - \log(1 + \mathrm{gt}_d) \right\|_2^2, \tag{3}$$

Table 1. Quantitative comparison of segmentation performance across reconstruction backbones, w/ and w/o DosePlug, using two representative segmentation methods.

Reconstruction backbone	Setting	Segmentation method	1%		5%		10%		25%		50%	
			DSC	HD95	DSC	HD95	DSC	HD95	DSC	HD95	DSC	HD95
AutoPET Dataset												
Input (low-dose)	-	SCUNet++	6.53	80.97	10.95	69.62	15.96	65.62	16.66	59.50	23.90	58.67
		SAMIHS	15.65	76.61	17.46	70.21	20.92	59.23	21.62	55.23	23.23	51.07
VisNet	Standard	SCUNet++	13.48	64.78	19.78	60.66	21.17	58.19	22.33	54.36	26.02	51.99
		SAMIHS	21.11	58.63	25.56	56.31	26.49	53.16	27.14	50.98	30.51	46.85
	w/ Ours	SCUNet++	16.83	57.66	23.62	56.89	24.58	53.72	25.10	50.18	29.81	45.29
		SAMIHS	26.69	55.49	29.33	53.10	30.15	49.26	30.44	46.87	33.17	41.18
RED	Standard	SCUNet++	14.10	61.94	20.73	59.86	21.71	57.22	24.14	53.14	28.56	49.81
		SAMIHS	23.76	46.51	25.67	44.45	27.91	41.86	29.21	40.04	30.98	39.39
	w/ Ours	SCUNet++	17.92	52.34	23.85	50.13	25.01	47.08	27.32	46.73	30.40	44.14
		SAMIHS	27.92	38.96	29.13	35.56	31.32	33.06	32.89	32.04	33.46	30.99
CTICH Dataset												
Input (low-dose)	-	SCUNet++	6.23	70.96	7.36	64.30	11.12	61.33	13.38	60.19	16.14	55.17
		SAMIHS	9.55	55.80	15.63	51.20	18.20	47.08	19.08	45.63	21.80	39.55
VisNet	Standard	SCUNet++	13.43	50.84	15.16	47.85	19.44	39.27	20.90	30.44	22.16	25.84
		SAMIHS	14.74	49.53	21.18	46.06	24.02	36.57	26.92	29.48	29.10	21.42
	w/ Ours	SCUNet++	16.11	46.62	18.20	41.59	23.89	32.18	24.07	27.35	26.30	22.57
		SAMIHS	17.93	48.15	26.21	40.85	26.56	32.62	29.48	25.24	33.82	17.92
RED	Standard	SCUNet++	10.25	56.33	11.70	52.57	15.04	45.13	16.37	36.43	20.85	33.47
		SAMIHS	12.64	44.29	19.53	39.18	23.00	36.07	23.29	34.75	24.67	31.07
	w/ Ours	SCUNet++	13.79	47.03	15.42	46.15	19.62	44.27	21.28	30.20	24.96	26.85
		SAMIHS	17.94	40.81	22.54	34.55	25.73	32.14	26.24	28.80	27.51	24.33

where $\|\cdot\|_2^2$ denotes the squared ℓ_2 norm and gt_d denotes the ground truth dose ratio. Jointly optimizing $\mathcal{L}_{\mathrm{recon}}$ and $\mathcal{L}_{\mathrm{OD}}$ shapes OD-dose to support dose-adaptive reconstruction without requiring dose metadata at inference.

3 Experiments

3.1 Experimental Setup

Dataset. CTICH [9] includes head scans from 82 traumatic brain injury patients, and AutoPET [8] includes whole-body PET from 900 patients; both datasets provide annotated lesion masks. For real-world low-dose evaluation, we use UDPET [21] (500 patients with multiple dose ratios) and the Mayo Clinic LDCT (hereafter, Mayo LDCT) dataset [13], which provides paired single low-dose and full-dose CT scans from 299 patients. Lacking lesion masks, these real-world low-dose datasets are exclusively used to evaluate reconstruction performance. Since CTICH and AutoPET do not provide paired low-dose scans, low-dose inputs are synthetically generated from clean slices by injecting noise in the sinogram domain using a projection noise method [15,25]. Data are split 8:2 for training and testing. DSC, the 95th-percentile Hausdorff distance (HD95), peak signal-to-noise ratio (PSNR), region-of-interest PSNR (rPSNR), and structural similarity index measure (SSIM) are used for evaluation.

Implementation Details. The experiments were conducted on an NVIDIA RTX 3090 GPU. DosePlug was fine-tuned end-to-end with the reconstruction backbone for 10 epochs using the Adam optimizer with a learning rate of 1×10^{-4}. Training was performed jointly using samples from all dose levels. The MI patch size was set to $s = 32$, and the SPL patch sizes were set to $(12, 15)$.

Table 2. Quantitative comparison of reconstruction performance at multiple dose levels on the AutoPET and CTICH datasets, w/ and w/o DosePlug.

Reconstruction backbone	Setting	1%		5%		10%		25%		50%	
		PSNR ↑	rPSNR ↑	PSNR ↑	rPSNR ↑	PSNR ↑	rPSNR ↑	PSNR ↑	rPSNR ↑	PSNR ↑	rPSNR ↑
AutoPET Dataset [8]											
Input (low-dose)	-	19.65	17.16	20.29	17.31	21.43	17.82	24.18	19.12	27.24	20.64
VisNet [24]	Standard	24.34	17.57	27.31	19.47	27.87	20.86	29.08	21.83	30.75	22.12
	w/ Ours	28.68	20.52	30.73	22.36	31.07	23.19	33.52	23.94	34.72	24.80
RED [1]	Standard	25.56	18.05	28.14	20.48	29.85	21.72	31.34	22.02	31.93	22.16
	w/ Ours	27.32	21.91	31.04	23.07	32.41	24.34	33.98	25.16	35.02	25.50
CTICH Dataset [9]											
Input (low-dose)	-	11.94	6.73	12.49	6.77	12.99	6.91	13.93	7.48	14.87	8.22
VisNet [24]	Standard	23.23	15.22	24.36	16.57	25.21	17.67	28.05	18.85	31.75	20.30
	w/ Ours	27.54	17.35	28.90	18.82	29.88	20.40	32.04	22.00	34.48	22.86
RED [1]	Standard	15.59	9.50	18.03	11.43	20.28	14.24	23.82	15.81	25.11	17.39
	w/ Ours	19.32	12.62	22.00	14.96	23.37	17.18	25.79	17.76	27.94	18.55

Table 3. (a) Ablation study of DosePlug, (b) Sensitivity analysis of the MI patch size s, (c) Sensitivity analysis of the patch size n, and (d) Computational complexity of DosePlug.

Align map	MI	FwPM order	DSC ↑ / PSNR ↑ AutoPET	DSC ↑ / PSNR ↑ CTICH
✓	✓	2	29.96/31.74	26.80/30.57
✓	✓	1	28.03/29.91	25.95/28.68
✓		2	27.81/28.50	25.10/27.38
	✓	2	27.94/28.68	25.67/27.56
			26.16/27.87	23.19/26.52

(a)

s	DSC ↑	PSNR ↑
AutoPET		
16	29.01	30.99
32	29.96	31.74
64	29.83	31.28
CTICH		
16	25.07	29.55
32	26.80	30.57
64	26.51	30.18

(b)

n	DSC ↑	PSNR ↑
AutoPET		
12	28.70	30.62
15	29.16	30.89
(12, 15)	29.96	31.74
CTICH		
12	25.41	29.41
15	25.75	29.10
(12, 15)	26.80	30.57

(c)

VisNet	
Params (K) ↓	338.22
FLOPs (G) ↓	124.91
Latency (ms) ↓	300.75
DosePlug (ours)	
Params (K) ↓	13.10
FLOPs (G) ↓	1.25
Latency (ms) ↓	2.02

(d)

3.2 Experimental Results

Tables 1 and 2 summarize the quantitative results on AutoPET and CTICH across dose ratios from 1% to 50%. Table 1 evaluates lesion segmentation after reconstruction using two representative segmentation methods, SCUNet++ [4] and SAMIHS [19], on two reconstruction backbones, RED [1] and VisNet [24]. DosePlug consistently increases DSC and reduces HD95 for both segmentation methods on both backbones. Table 2 reports reconstruction quality at multiple dose levels. DosePlug improves PSNR and rPSNR across doses on both datasets and backbones. The rPSNR metric was selected in collaboration with nuclear medicine clinicians to better reflect clinically relevant reconstruction fidelity. Fig. 4 presents qualitative results with and without DosePlug. In Fig. 4(a), models trained on CTICH and AutoPET are evaluated at 1% dose on their respective test sets, and DSC and PSNR are reported. DosePlug improves both PSNR and DSC, in line with the quantitative results. In particular, for CT, we observed a tendency to reduce streak artifacts. In Fig. 4(b), cross-dataset generalization is evaluated by training on CTICH and testing on Mayo LDCT, and by training on AutoPET and testing on UDPET, with PSNR and SSIM reported. The results on real-world datasets indicate that the estimated OD-dose remains reliable beyond synthetic simulation and supports robust reconstruction

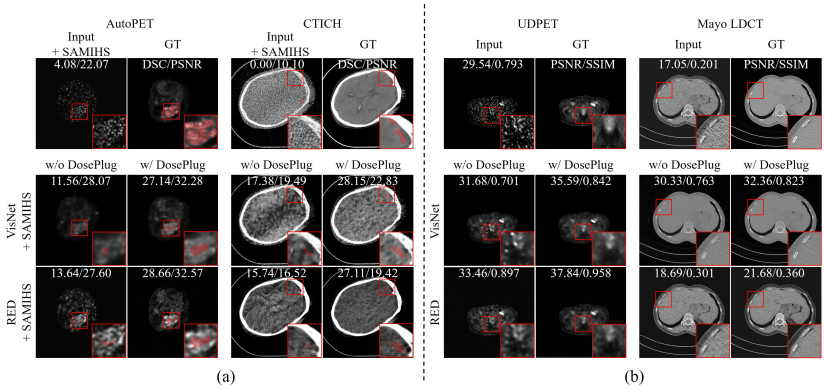

Fig. 4. Qualitative comparisons with and without DosePlug. (a) CTICH and AutoPET datasets at 1% dose. (b) Cross-dataset results on real-world low-dose UDPET (1% dose) and Mayo LDCT.

calibration under real acquisition conditions, which suggests potential benefits for downstream segmentation in real-world low-dose imaging settings.

OD-dose Estimation Accuracy. We evaluate OD-dose estimation accuracy using MSE on the scalar dose ratio (1-100), rather than a log-transformed scale, obtaining 0.43 on AutoPET, 1.80 on CTICH, and 0.82 on UDPET (real-world low-dose), indicating consistent capture of dataset-level relative degradation without metadata at inference.

Ablation Study. Ablations insert DosePlug into VisNet and evaluate SAMIHS segmentation, with DSC and PSNR averaged over five dose levels. Table 3(a) analyzes the contribution of the two internal cues. Using both the alignment map and MI computation yields the best performance, while removing either cue degrades reconstruction and segmentation metrics. The polynomial order study shows that second-order FwPM outperforms first-order under the full-cue setting. Table 3(b) indicates that an MI patch size of $s = 32$ yields the best performance, and Table 3(c) shows that combining multi-scale SPL patch sizes $(12, 15)$ outperforms using a single patch size. Table 3(d) reports the computational complexity of DosePlug. Relative to VisNet, DosePlug adds 3.87% parameters, 1.00% FLOPs, and 0.67% latency, while improving DSC and PSNR at the 1% dose level by 26.74% and 16.80%, respectively. The accuracy-to-complexity trade-off was qualitatively endorsed by hospital nuclear medicine clinicians.

4 Conclusion

We introduced DosePlug, a lightweight plug-in that enables low-dose adaptation without dose metadata at inference. By inferring a task-relevant degradation coordinate from sinogram periodicity and MI cues, DosePlug applies polynomial feature calibration. Experiments across CT/PET backbones show consistent gains

with minimal overhead, and the improvements persist on real-world low-dose datasets. However, OD-dose estimation may become less stable under extreme domain shifts unrelated to dose, such as scanner-dependent artifacts. Future work will extend DosePlug to volumetric 3D reconstruction, evaluate real-world low-dose segmentation, and further study robustness to such non-dose domain shifts.

Acknowledgments. This work was supported by the IITP grant funded by the Korea Government (MSIT) (No. RS-2024-00437718, RS-2023-00256629, RS-2022-00156287).

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

References

1. Ai, X., Huang, B., Chen, F., Shi, L., Li, B., Wang, S., Liu, Q.: Red: Residual estimation diffusion for low-dose pet sinogram reconstruction. Medical Image Analysis **102**, 103558 (2025). https://doi.org/10.1016/j.media.2025.103558
2. Boellaard, R., Delgado-Bolton, R., Oyen, W.J., Giammarile, F., Tatsch, K., Eschner, W., Verzijlbergen, F.J., Barrington, S.F., Pike, L.C., Weber, W.A., et al.: Fdg pet/ct: Eanm procedure guidelines for tumour imaging: version 2.0. European journal of nuclear medicine and molecular imaging **42**(2), 328–354 (2015)
3. Brix, G., Lechel, U., Glatting, G., Ziegler, S.I., Münzing, W., Müller, S.P., Beyer, T.: Radiation exposure of patients undergoing whole-body dual-modality 18f-fdg pet/ct examinations. Journal of nuclear medicine **46**(4), 608–613 (2005)
4. Chen, Y., Zou, B., Guo, Z., Huang, Y., Huang, Y., Qin, F., Li, Q., Wang, C.: Scunet++: Swin-unet and cnn bottleneck hybrid architecture with multi-fusion dense skip connection for pulmonary embolism ct image segmentation. In: Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision (WACV). pp. 7759–7767. IEEE (2024). https://doi.org/10.1109/WACV57701.2024.00758
5. Chen, Z., Gao, Q., Zhang, Y., Shan, H.: Ascon: Anatomy-aware supervised contrastive learning framework for low-dose ct denoising. In: International conference on medical image computing and computer-assisted intervention. pp. 355–365. Springer (2023)
6. Cui, J., Xie, Y., Gong, K., Kim, K., Yang, J., Larson, P., Hope, T., Behr, S., Seo, Y., Liu, H., et al.: Pet denoising and uncertainty estimation based on nvae model. In: 2021 IEEE Nuclear Science Symposium and Medical Imaging Conference (NSS/MIC). pp. 1–3. IEEE (2021)
7. Elbakri, I.A., Fessler, J.A.: Statistical image reconstruction for polyenergetic x-ray computed tomography. IEEE transactions on medical imaging **21**(2), 89–99 (2002)
8. Gatidis, S., Kuestner, T.: A whole-body fdg-pet/ct dataset with manually annotated tumor lesions (fdg-pet-ct-lesions). The Cancer Imaging Archive (TCIA) (2022). https://doi.org/10.7937/gkr0-xv29, dataset
9. Hssayeni, M., Croock, M., Salman, A., Al-khafaji, H., Yahya, Z., Ghoraani, B.: Computed tomography images for intracranial hemorrhage detection and segmentation. Data **5**(1), 14 (2020). https://doi.org/10.3390/data5010014

10. Jing, J., Wang, T., Yu, H., Lu, Z., Zhang, Y.: Inter-slice consistency for unpaired low-dose ct denoising using boosted contrastive learning. In: International Conference on Medical Image Computing and Computer-Assisted Intervention. pp. 238–247. Springer (2023)
11. Kim, M.H., Yoo, S.B.: Data poisoning attack defense and evolutionary domain adaptation for federated medical image segmentation. In: Proceedings of the Thirty-Fourth International Joint Conference on Artificial Intelligence. pp. 1305–1313 (2025)
12. McCollough, C.H., Primak, A.N., Braun, N., Kofler, J., Yu, L., Christner, J.: Strategies for reducing radiation dose in ct. Radiologic Clinics of North America **47**(1), 27 (2009)
13. 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. Medical physics **48**(2), 902–911 (2021)
14. Niu, X., Ye, Z., Han, X., Huang, Y., Sun, H., Wu, H., Lu, L.: Mdaa-diff: Ct-guided multi-dose adaptive attention diffusion model for pet denoising. In: International Conference on Medical Image Computing and Computer-Assisted Intervention. pp. 333–343. Springer (2025)
15. Sanaat, A., Arabi, H., Mainta, I., Garibotto, V., Zaidi, H.: Projection space implementation of deep learning–guided low-dose brain pet imaging improves performance over implementation in image space. Journal of Nuclear Medicine **61**(9), 1388–1396 (2020)
16. Shepp, L.A., Vardi, Y.: Maximum likelihood reconstruction for emission tomography. IEEE transactions on medical imaging **1**(2), 113–122 (2007)
17. Townsend, D.W., Beyer, T., Blodgett, T.M.: Pet/ct scanners: a hardware approach to image fusion. In: Seminars in nuclear medicine. vol. 33, pp. 193–204. Elsevier (2003)
18. Wang, R., Chen, Z., Song, Z., Fang, W., Zhang, J., Tu, D., Tang, Y., Xu, M., Ye, X., Lu, L., et al.: Anatomy-aware low-dose ct denoising via pretrained vision models and semantic-guided contrastive learning. In: International Conference on Medical Image Computing and Computer-Assisted Intervention. pp. 13–23. Springer (2025)
19. Wang, Y., Chen, K., Yuan, W., Tang, Z., Meng, C., Bai, X.: Samihs: Adaptation of segment anything model for intracranial hemorrhage segmentation. In: 2024 IEEE International Symposium on Biomedical Imaging (ISBI). pp. 1–5. IEEE (2024). https://doi.org/10.1109/ISBI56570.2024.10635673
20. Xia, W., Lu, Z., Huang, Y., Liu, Y., Chen, H., Zhou, J., Zhang, Y.: Ct reconstruction with pdf: Parameter-dependent framework for data from multiple geometries and dose levels. IEEE Transactions on Medical Imaging **40**(11), 3065–3076 (2021)
21. Xue, S., Wang, H., Chen, Y., Liu, F., Zhu, H., Viscione, M., Guo, R., Rominger, A., Li, B., Shi, K.: Udpet: Ultra-low dose pet imaging challenge dataset. In: Medical Image Computing and Computer Assisted Intervention – MICCAI 2025. Lecture Notes in Computer Science, vol. 15972, pp. 616–623. Springer (2025). https://doi.org/10.1007/978-3-032-05169-1_59
22. Yoon, J.H., Lee, Y.J., Yoo, S.B.: La-seg: Disentangled sinogram pattern-guided transformer for lesion segmentation in limited-angle computed tomography. Computers in Biology and Medicine **196**, 110728 (2025)
23. Yoon, J.H., Lee, Y.J., Yoo, S.B.: Transsino: Prior sinogram pattern-based transformer for limited-angle ct image segmentation. In: International Conference on Medical Image Computing and Computer-Assisted Intervention. pp. 586–596. Springer (2025)

24. Yue, H., Kuang, H., Ma, L., Liu, J., Li, J., Cheng, J., Wang, J.: Visnet: A human visual system inspired lightweight dual-path network for medical images denoising. In: Medical Image Computing and Computer Assisted Intervention – MICCAI 2025. Lecture Notes in Computer Science, vol. 15972, pp. 666–676. Springer Nature Switzerland (Sep 2025). https://doi.org/10.1007/978-3-032-05169-1_64
25. Zeng, D., Huang, J., Bian, Z., Niu, S., Zhang, H., Feng, Q., Liang, Z., Ma, J.: A simple low-dose x-ray ct simulation from high-dose scan. IEEE transactions on nuclear science **62**(5), 2226–2233 (2015)