

GH-NAF: Grid-Adaptive Hash-Level-Attended Neural Attenuation Fields for Discrepancy-Aware CBCT

Supplementary Material

Limitations of CBCT Reconstruction under Projection-Domain Discrepancies

In an idealized monoenergetic CT system with a perfectly calibrated geometry, each detector measurement would equal the negative log of a line integral of the linear attenuation coefficient $\mu(x)$ along a single ray. Under this assumption, the forward model is clean, shift-invariant, and fully characterized by the Beer–Lambert law [8]. In practice, however, clinical CT and CBCT systems deviate from this ideal in several systematic ways, leading to projection-domain discrepancies (i.e., structured differences between the measured log-intensity and the ideal line integral of $\mu(x)$) [3, 4].

A first major source of discrepancy is polychromatic X-ray spectra and beam hardening [1, 6, 11]. Real X-ray tubes emit a broad energy spectrum, and different materials attenuate low- and high-energy photons differently. As the beam traverses dense structures such as bone, low-energy photons are preferentially absorbed, the spectrum hardens, and the effective attenuation no longer follows a single linear $\mu(x)$. When such polychromatic measurements are inverted under a monoenergetic assumption, the forward model systematically misfits the data, producing depth-dependent errors that manifest as cupping artifacts and non-uniform intensities within the same material.

A second source is scattered radiation [2, 10, 17]. In CBCT, the large cone angle and wide detector coverage make single-scatter and multiple-scatter contributions particularly significant. Scatter adds an approximately low-frequency, view-dependent bias on top of the primary signal, violating the simple exponential attenuation model. Because scatter depends on the global 3D distribution of attenuation, it cannot be represented as a local line integral along an individual ray. Therefore, ignoring scatter in the forward model introduces a structured discrepancy that correlates with object size, shape, and composition.

Additional instrumental and geometric effects further amplify projection domain discrepancies [14]. These include detector lag and veiling glare, nonlinear and intensity-dependent detector response, imperfect gain and offset calibration that produces residual ring artifacts, and inaccuracies in the relative positioning of the X-ray source and detector. Each of these mechanisms systematically perturbs the measured projection data away from the ideal logarithmic line integrals assumed by standard reconstruction algorithms, thereby introducing additional structured error that

is not explained by the simple attenuation model.


These projection-domain mismatches matter because CT reconstruction is an ill-posed inverse problem that is extremely sensitive to forward model errors. Analytic methods such as FDK and standard filtered backprojection essentially assume that the log-projections are exact line integrals. When this assumption is violated, the reconstruction algorithm attempts to explain low-frequency biases (from beam hardening and scatter) and high-frequency inconsistencies (from miscalibration and noise) using only variations in $\mu(x)$. The result is characteristic shading, cupping artifacts, streaks around dense structures, and non-uniform attenuation within nominally homogeneous tissues. Iterative and learning-based methods inherit the same issue: if the forward operator used in optimization does not match the true measurement process, the learned attenuation field is forced to absorb the discrepancy, entangling physical structure with acquisition artifacts [9, 15].

In summary, projection discrepancy is not a small perturbation but a dominant error source in practical CT/CBCT. It arises whenever real-world acquisition physics (polychromatic spectra, scatter, detector imperfections) and system calibration differ from the simplistic assumptions built into standard forward models. Any reconstruction framework that directly learns or optimizes $\mu(x)$ from such data must therefore either explicitly model these discrepancies or accept that they will be implicitly baked into the reconstructed volume as structured artifacts and low-frequency bias.

Motivated by this observation, GH-NAF augments the standard monoenergetic Beer–Lambert forward model with a discrepancy-aware differentiable renderer that explicitly accounts for beam-hardening-induced projection mismatch. Concretely, GH-NAF introduces a learnable discrepancy density $\sigma(x)$ and a smooth correction term in the projection domain, so that large-scale shading and other low-frequency biases are predominantly captured by the discrepancy field, while $\mu(x)$ is reserved for true structural attenuation. By disentangling discrepancy from attenuation in this way, GH-NAF relieves the attenuation field from explaining projection-domain errors purely through $\mu(x)$, leading to reconstructions with more uniform intra-material contrast and fewer shading and streak artifacts even under strong projection-domain mismatch.


Details of Datasets

Representative projection examples for both the chest phantom and FIPS objects are shown in Fig. 2.




100view 125view 150view 200view

NAF




100view 125view 150view 200view

SAX



100view 125view 150view 200view

GH-NAF



720view

FDK

Figure 1. Qualitative results of NeRF-based 3D reconstruction for the chest phantom dataset.

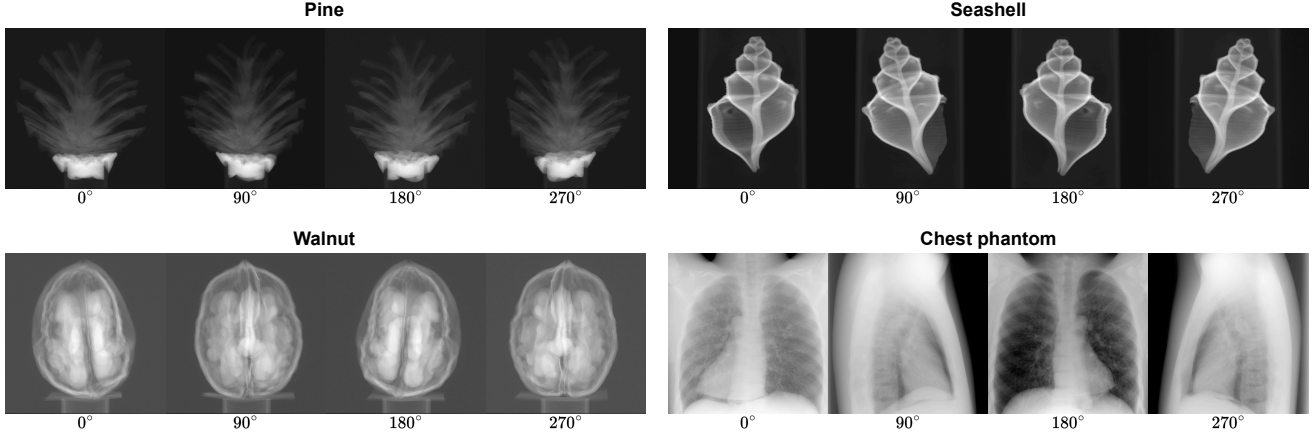


Figure 2. Real-world CBCT datasets used in our experiments. Representative projection images are shown for each case.

1. Chest Phantom Dataset

For real-data experiments with pronounced projection-domain discrepancy, we use raw projection data from a chest phantom scan acquired on a commercial mobile CBCT system. The scan comprises 720 projections over a full 360° rotation with a uniform angular increment of 0.5° . The original detector resolution is 1516×1516 pixels with an effective pixel size of 0.56 mm. The source-object distance (SOD) is 620 mm and the source-detector distance (SDD) is 920 mm. In the source-centric coordinate system, the detector center is shifted by -1.4 mm and -56.66 mm along the detector x and y axes, respectively.

Mobile CBCT systems in practice exhibit non-negligible geometric uncertainty and calibration drift across acquisitions, including small but systematic deviations in both in-plane and out-of-plane angles. To ensure a fair comparison with clinical reconstruction, we adopt as our forward model geometry the best-available system configuration used in the vendor’s reconstruction pipeline, including the refined in-plane and out-of-plane angle settings obtained by internal geometry optimization. For computational efficiency, all projections are downsampled from 1516×1516 to 758×758 pixels prior to training. All methods reconstruct a $256 \times 256 \times 256$ attenuation volume for this dataset.

2. FIPS Dataset

The FIPS dataset consists of three scanned objects (pine cone, walnut, and seashell) acquired on a benchtop CBCT system with a fixed, high-precision geometry. Each object is measured with 721 projections over 360° , corresponding to an angular step of approximately 0.5° . The acquisition uses a static, well-calibrated geometry with a source-object distance of 230.66 mm and a source-detector distance of 553.74 mm. The original detector resolution is 2368×2240 pixels with an effective pixel size of 0.05 mm. For computational efficiency, we downsample the projections from

2368×2240 to 592×560 pixels before training. As in the chest phantom experiments, all methods reconstruct a $256 \times 256 \times 256$ attenuation volume for this dataset.

Implementation Details

Our GH-NAF uses the original NAF backbone as the base architecture and only modifies the output head. For real-data experiments, the network predicts three channels per sample, corresponding to the attenuation field $\mu(x)$, the heteroscedastic variance $\beta^2(x)$, and the discrepancy density $\sigma(x)$. For synthetic discrepancy-free experiments, where no explicit discrepancy field is required, we use a two-channel head that outputs only $\mu(x)$ and $\beta^2(x)$.

For encoding, we employ a three-dimensional hashgrid encoder with 16 multi-resolution levels. Each level produces a 2-dimensional feature vector, starting from a base spatial resolution of 16 voxels. All levels share a single hash table with a capacity of 2^{24} entries, which yields a compact yet expressive multi-scale representation of the attenuation field.

Unless otherwise stated, we use fixed loss weights of $\lambda_{\text{attn}} = 0.05$ for the attention loss $\mathcal{L}_{\text{attn}}$ and $\lambda_{\text{occ}} = 0.01$ for the occlusion prior. The reconstruction and regularization term L_r is weighted by an epoch-dependent coefficient $\lambda_r^{(t)}$. Instead of a constant λ_r , we adopt a linear warm-up schedule: the effective weight starts at $0.01 \lambda_r$ in the first epoch and is gradually increased to approximately $0.1 \lambda_r$ by epoch 500. This schedule keeps the influence of L_r intentionally weak in the early stages, allowing the network to first learn a coarse but stable attenuation field from the projection data, and then progressively strengthen the reconstruction and regularization constraint as training proceeds. We follow the original discrepancy model [13], where λ_σ controls the discrepancy strength, but empirically fix $\lambda_\sigma = 2$ for all experiments. This value lies between the recommended settings ($\lambda_\sigma = 3$ for numerical and $\lambda_\sigma = 1$ for

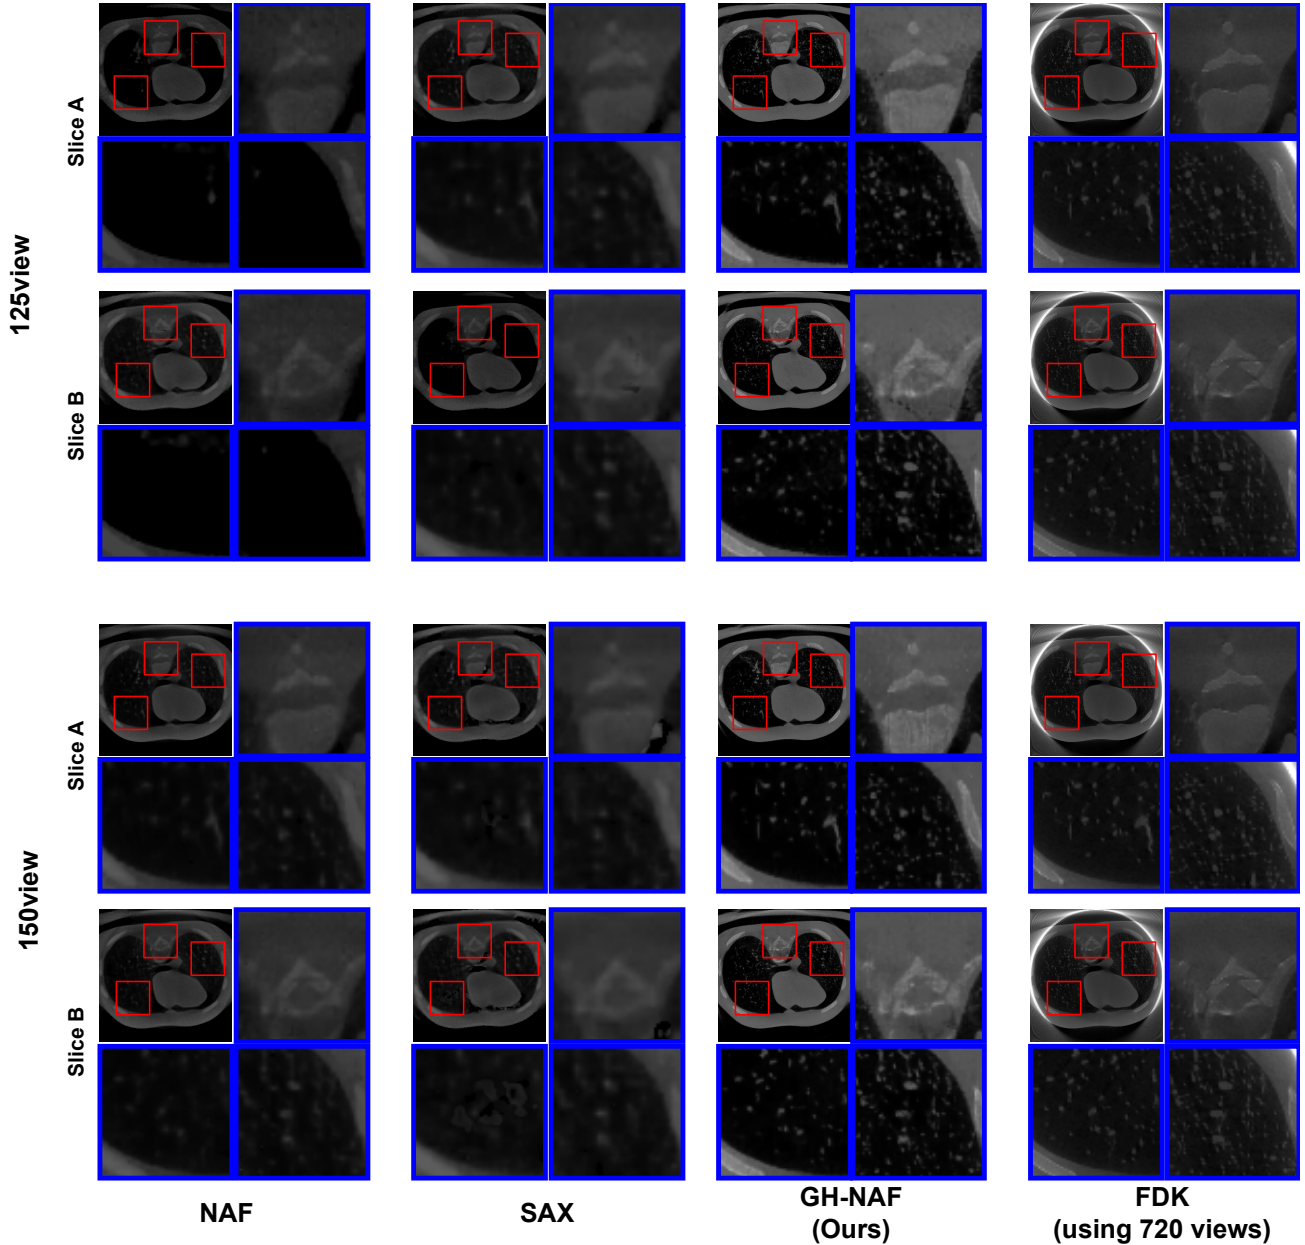


Figure 3. Qualitative detailed view of chest phantom reconstructions at 125- and 150-view.

phantom data in the original paper) and yielded the most stable reconstructions in our CBCT scenario.

To obtain precise reconstructions, we employ dataset-specific ray sampling and encoder bounds. The chest phantom data are acquired with a mobile CBCT system, where the actual imaging geometry is not perfectly fixed and exhibits view-dependent inconsistencies, as confirmed by an internal geometry optimization procedure based on the reconstructed volume. For this setting, we use a higher ray sampling budget: in each mini batch, we cast one ray per

Table 1. Tenengrad-based sharpness evaluation (averaged over three patches) of NeRF-based 3D reconstructions under different view settings on the chest phantom dataset.

Model	view (Slice A)				view (Slice B)			
	100	125	150	200	100	125	150	200
NAF	311.4	752.5	418.8	418.3	251.5	659.2	387.5	370.9
SAX	350.2	378.9	469.2	695.2	287.6	340.6	420.9	590.1
GH-NAF (ours)	750.5	813	825.5	793.9	800.2	815.1	838.3	846.0

selected detector pixel and evaluate the line integral using




Figure 4. Qualitative results of 3D reconstruction in sparse-view settings for the FIPS dataset.

1024 samples along each ray in order to better resolve the attenuation field under geometric uncertainty. The spatial bound of the hash encoder is aligned with the physical scanner geometry; in particular, we set the radius of the encoded volume to 0.62, which corresponds to the source to object distance of 620 mm. Given the higher level of projection-domain discrepancy in this scenario, we train GH-NAF on the mobile CBCT chest phantom dataset for 3000 epochs.

For the FIPS data and the synthetic minimal/free discrepancy datasets, which exhibit minimal projection domain mismatch and rely on a well-calibrated static geometry, we employ a lighter sampling budget of 320 samples along each ray. In these experiments, the hash encoder bound is set to 0.3, which tightly encloses the object support while avoiding unnecessary hash table usage outside the region of interest. Under these minimal/free discrepancy conditions, we train the model for 1500 epochs on the FIPS dataset and for 1000 epochs on the synthetic datasets, which we found sufficient for stable convergence.

Results

1. Qualitative Results on the Chest Phantom

Figs. 1 and 3 present a qualitative comparison of reconstructions on the chest phantom acquired with a mobile CBCT system. Compared with the original NAF [16] and SAX [5] backbones, GH-NAF consistently produces cleaner lung parenchyma, sharper vessel and bronchial boundaries, and visibly reduced noise and streak-like artifacts across all tested view counts. Fine vascular and parenchymal structures that appear blurred or partially washed out in NAF and SAX remain well delineated under GH-NAF, without

introducing obvious hallucinated details. In addition, regions near dense structures such as the spine and ribs exhibit less cupping and shading, indicating that GH-NAF is more robust to projection-domain discrepancy than the baseline NeRF-style models.

Table 1 reports the Tenengrad scores [7] for GH-NAF, NAF, and SAX on the same chest phantom experiments. The Tenengrad metric quantifies image sharpness by measuring the Sobel gradient magnitude; higher values indicate sharper and more well-defined edges, whereas increasing blur leads to a noticeable reduction in the score. GH-NAF consistently achieves the highest Tenengrad values across all view counts (100–200 view), outperforming both NAF and SAX. This confirms the model’s superior ability to preserve structurally meaningful edges that are typically degraded by projection discrepancies, yielding reconstructions with markedly stronger and more coherent edge responses. The advantage is particularly pronounced in sparse-view regimes (100–125 view), where GH-NAF maintains excellent edge fidelity even under limited angular sampling, complementing the qualitative observations in Figs. 1 and 3.

When contrasted with the 720-view FDK reconstruction, GH-NAF attains comparable or even superior visual quality while using only a fraction of the projections. In particular, truncation artifacts at the chest boundary are substantially suppressed, and low-frequency shading and noise-like patterns are attenuated, even at one-third or fewer of the view. As the number of view increases, GH-NAF initially benefits from denser angular sampling, yielding progressively clearer vessels and smoother soft tissue. However, this improvement does not continue indefinitely: at higher

Table 2. Ablation study on the effect of $\mathcal{L}_{\text{attn}}$ and the uncertainty term β^2 on reconstruction performance on the synthetic datasets across different view numbers.

$\mathcal{L}_{\text{attn}}$	β^2	15 view		30 view		50 view	
		PSNR	SSIM	PSNR	SSIM	PSNR	SSIM
✗	✗	33.15	0.9285	36.92	0.9655	39.59	0.9808
✓	✗	33.34	0.9304	37.17	0.9673	39.65	0.9812
✓	✓	33.56	0.9341	37.31	0.9682	39.82	0.9815

view counts, the incremental gain in visual quality becomes marginal, which is consistent with the presence of residual geometric mismatch in the mobile CBCT acquisition. These observations suggest that the proposed hash-level attention and discrepancy-aware training can effectively exploit additional projections up to a regime where projection-domain discrepancy begins to dominate, beyond which further view no longer translate into proportional qualitative gains.

2. Qualitative Results on the FIPS dataset

Fig. 4 shows qualitative reconstructions on the FIPS dataset, acquired using a benchtop CBCT system with a well-calibrated geometry and minimal projection-domain discrepancy. Using the 720-view FDK reconstruction as a reference, GH-NAF recovers sharper object boundaries and more uniform attenuation within each object than NAF and SAX, especially around high-contrast interfaces and thin structures. GH-NAF more effectively suppresses residual shadowing around dense regions and noise-like patterns in low-attenuation regions, while also better preserving fine structural details. These observations indicate that the proposed hash-level attention and discrepancy-aware training improve the multi-resolution representation even when the acquisition geometry is already accurate.

Compared to the chest phantom setting, where strong beam hardening and geometric uncertainty produce pronounced projection-domain discrepancies, the FIPS dataset represents a low-discrepancy region in which artifacts are subtler but still present. In this region, GH-NAF primarily serves as a robust regularizer on hash-level feature fusion: it reduces minor shading, stabilizes attenuation within homogeneous materials, and prevents over-smoothing at edges without introducing new artifacts. Therefore, the qualitative results on the FIPS dataset complement the chest phantom analysis, showing that GH-NAF consistently enhances structural fidelity and intra-material uniformity over NAF and SAX, rather than merely compensating for extreme mismatch between the forward model and the real acquisition process.

Table 3. Ablation on model agnosticism: comparison between the original SAX backbone and its GH-SAX extension on the synthetic datasets.

	15 view		30 view		50 view	
	PSNR	SSIM	PSNR	SSIM	PSNR	SSIM
SAX	32.69	0.9248	36.25	0.9605	38.40	0.9755
GH-SAX	32.93	0.9220	37.05	0.9662	39.74	0.9815

Ablation Study

1. Effect of Uncertainty-aware Attention Loss

Table 2 evaluates the impact of the gradient-guided attention loss $\mathcal{L}_{\text{attn}}$ and the uncertainty term β^2 . Adding $\mathcal{L}_{\text{attn}}$ alone already improves PSNR/SSIM over the baseline without attention or uncertainty across sparse-view settings. When we further enable the uncertainty-aware formulation, we obtain the best reconstruction scores at 15, 30, and 50 view. This confirms that our attention supervision is beneficial, and that weighting it by the predicted variance β^2 provides an additional, consistent gain in reconstruction quality.

2. Model-Agnostic Ablation

To verify that the proposed GH-NAF is agnostic to the underlying NeRF-style backbone, we also apply our method on top of the SAX [5] architecture instead of NAF. Concretely, we keep the original SAX network unchanged and simply attach the same hashgrid encoder and attention-guided supervision used for GH-NAF, yielding a GH-SAX variant. As shown in Table 3, GH-SAX consistently improves PSNR over the SAX baseline on the synthetic datasets across all view settings, while maintaining comparable SSIM at 15 view and achieving clear SSIM gains at 30 and 50 view. These results indicate that our hash-level attention and loss design do not rely on any NAF-specific architectural choices and can be readily integrated into other, potentially stronger NeRF/NAF backbones as they emerge.

3. Additional Analysis on Intensity Preservation and Discrepancy Disentanglement

Intensity Preservation and Material Differentiation.

Although recovering exact CT numbers under real-world projection discrepancy remains challenging, the proposed GH-NAF better preserves relative intensity variation and high-frequency structural transitions than existing methods. Fig. 5 compares the intensity distributions of representative slices across 720-view FDK and 150-view neural-field methods (NAF, SAX, and GH-NAF). Compared with NAF and SAX, which exhibit noticeable over-smoothing, GH-NAF produces sharper intensity transitions and better maintains local variation near structural boundaries. These results are consistent with our design motivation: by adaptively modulating representational capacity through learned

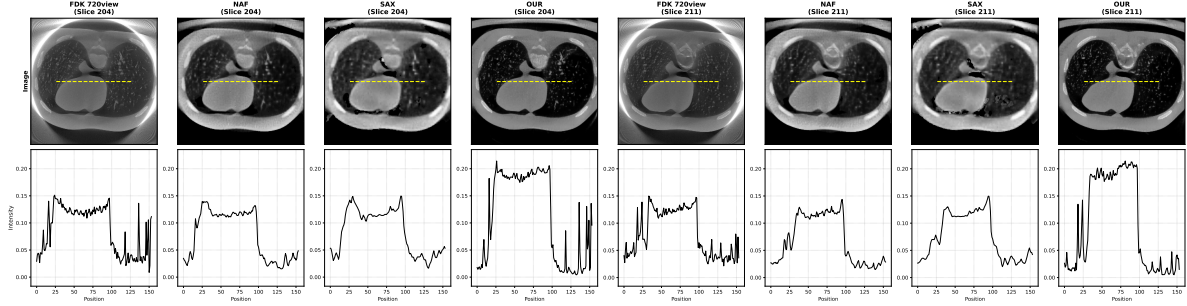


Figure 5. Comparison of intensity distributions for slices (Methods: FDK, NAF, SAX, and Ours). Our method preserves high-frequency details, avoiding the significant over-smoothing observed in NAF and SAX.

uncertainty, the proposed model alleviates the tendency of the reconstruction field to absorb mismatch as blurred intensity bias. As a result, GH-NAF shows improved capability to distinguish adjacent materials and preserve boundary-aware contrast, even when exact voxel-wise CT calibration is imperfect.




Figure 6. Phantom reconstruction comparison of GH-NAF without and with σ (discrepancy correction) under the 150-view setting.

Effect of Discrepancy Density σ . We further analyze the role of the discrepancy density σ , which is introduced to model low-frequency projection mismatch separately from the attenuation field μ , following the physical motivation of prior discrepancy-aware formulations [12, 13]. Fig. 6 compares phantom reconstructions obtained without and with σ . Without σ , the attenuation field μ must explain both anatomical attenuation and projection discrepancy, causing mismatch-related patterns to be absorbed into the reconstructed structure. This leads to stronger artifacts, degraded material separation, and less reliable boundary definition. In contrast, when σ is explicitly modeled, low-frequency discrepancy can be handled independently of the main attenuation field, resulting in cleaner reconstructions, reduced artifact contamination, and more physically plausible material transitions. These observations support our interpretation that discrepancy–attenuation disentanglement is important for stable reconstruction under real-world non-ideal acquisition conditions.




Figure 7. Low/high-level attention mass maps (levels 0–3 vs. 12–15) without and with \mathcal{L}_{attn} on an axial slice, showing that \mathcal{L}_{attn} induces spatially structured, level-dependent gating.

4. Effect of Hash-Level Attention

Fig. 7 visualizes the spatial distribution of low-level and high-level attention mass on an axial slice, comparing the cases without and with the proposed attention supervision loss \mathcal{L}_{attn} . Without \mathcal{L}_{attn} , the attention is relatively diffuse and less spatially structured, indicating that the model does not clearly distinguish homogeneous regions from structural boundaries when selecting hash levels. In contrast, with \mathcal{L}_{attn} , coarse-level attention is more strongly concentrated in relatively homogeneous regions, whereas fine-level attention becomes more focused around boundaries and anatomically detailed structures. This behavior is consistent with the intended role of GH-NAF, namely, to allocate representational capacity according to local structural complexity rather than uniformly aggregating all hash levels. Therefore, the visualization provides direct qualitative evidence that the proposed attention supervision promotes spatially adaptive hash-level selection, contributing to improved boundary preservation and reduced low-frequency bias in sparse-view reconstruction.

Limitations

Our study has several limitations that should be acknowledged. First, the mobile CBCT system used for the thoracic phantom experiments exhibits mild but non-negligible geometric mismatch and calibration drift across view, including small deviations in in-plane and out-of-plane angles that are not fully captured by a single static geometry model. Con-

sequently, the measured projections are only approximately consistent with the forward operator used during training. As the number of view increases, these discrepancies can accumulate, partially offsetting the benefits of denser angular sampling. This explains why GH-NAF does not show strictly monotonic improvement with view count on the mobile CBCT data, even though the underlying network capacity and loss design would, in principle, support higher quality reconstructions under perfectly calibrated conditions. This behavior highlights a strong coupling between the expressiveness of the learned attenuation field and the fidelity of the acquisition model: sophisticated neural reconstruction alone cannot fully compensate for systematic geometry errors and other sources of projection-domain mismatch.

Second, our discrepancy model primarily targets beam hardening and related low-frequency effects, and does not yet explicitly account for all sources of projection-domain error present in clinical systems, such as scatter, nonideal detector response, and residual ring related artifacts. These unmodeled effects may still be partially absorbed into the learned attenuation field, limiting the ultimate interpretability and robustness of the reconstructions. In future work, we plan to incorporate geometry refinement into the learning process, for example, by jointly optimizing scanner pose and calibration parameters together with the attenuation field, and to extend the discrepancy modeling to more comprehensive physical effects. By improving both geometric calibration and forward model fidelity, we expect GH-NAF to better scale with view count and to provide more reliable reconstructions in challenging real-world CBCT scenarios.

References

- [1] Monica Abella, Cristóbal Martínez, Manuel Desco, Juan José Vaquero, and Jeffrey A Fessler. Simplified statistical image reconstruction for x-ray ct with beam-hardening artifact compensation. *IEEE transactions on medical imaging*, 39(1):111–118, 2019. 1
- [2] Awais Ashfaq and Jonas Adler. A modified fuzzy c means algorithm for shading correction in craniofacial cbct images. In *CMBEIH 2017: Proceedings of the International Conference on Medical and Biological Engineering 2017*, pages 531–538. Springer, 2017. 1
- [3] Julia F Barrett and Nicholas Keat. Artifacts in ct: recognition and avoidance. *Radiographics*, 24(6):1679–1691, 2004. 1
- [4] F Edward Boas, Dominik Fleischmann, et al. Ct artifacts: causes and reduction techniques. *Imaging Med*, 4(2):229–240, 2012. 1
- [5] Yuanhao Cai, Jiahao Wang, Alan Yuille, Zongwei Zhou, and Angtian Wang. Structure-aware sparse-view x-ray 3d reconstruction. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 11174–11183, 2024. 5, 6
- [6] Qiyong Fan, Bo Lu, Justin C Park, Tianye Niu, Jonathan G Li, Chihray Liu, and Lei Zhu. Image-domain shading correction for cone-beam ct without prior patient information. *Journal of applied clinical medical physics*, 16(6):65–75, 2015. 1
- [7] Shuqin Gao, Min Han, and Xu Cheng. The fast iris image clarity evaluation based on tenengrad and roi selection. In *Ninth International Conference on Graphic and Image Processing (ICGIP 2017)*, pages 1391–1396. SPIE, 2018. 5
- [8] Jakob Sauer Jørgensen. Sparse image reconstruction in computed tomography. 2013. 1
- [9] Anish Lahiri, Gabriel Maliakal, Marc L Klasky, Jeffrey A Fessler, and Saiprasad Ravishankar. Sparse-view cone beam ct reconstruction using data-consistent supervised and adversarial learning from scarce training data. *IEEE Transactions on Computational Imaging*, 9:13–28, 2023. 1
- [10] Tianye Niu, Mingshan Sun, Josh Star-Lack, Hwei Gao, Qiyong Fan, and Lei Zhu. Shading correction for on-board cone-beam ct in radiation therapy using planning mdct images. *Medical physics*, 37(10):5395–5406, 2010. 1
- [11] Hyoung Suk Park, Dosik Hwang, and Jin Keun Seo. Metal artifact reduction for polychromatic x-ray ct based on a beam-hardening corrector. *IEEE transactions on medical imaging*, 35(2):480–487, 2015. 1
- [12] Hyoung Suk Park, Kiwan Jeon, and JK Seo. Deep learning-based artefact reduction in low-dose dental cone beam computed tomography with high-attenuation materials. *Philosophical Transactions of the Royal Society A: Mathematical, Physical and Engineering Sciences*, 383(2305), 2025. 7
- [13] Hyoung Suk Park, Jin Keun Seo, and Kiwan Jeon. Implicit neural representation-based method for metal-induced beam hardening artifact reduction in x-ray ct imaging. *Medical Physics*, 52(4):2201–2211, 2025. 3, 7
- [14] Jared Starman, Josh Star-Lack, Gary Virshup, Edward Shapiro, and Rebecca Fahrig. A nonlinear lag correction algorithm for a-si flat-panel x-ray detectors. *Medical Physics*, 39(10):6035–6047, 2012. 1
- [15] Jia Wu, Jinzhao Lin, Xiaoming Jiang, Wei Zheng, Lisha Zhong, Yu Pang, Hongying Meng, and Zhangyong Li. Dual-domain deep prior guided sparse-view ct reconstruction with multi-scale fusion attention. *Scientific Reports*, 15(1):16894, 2025. 1
- [16] Ruyi Zha, Yanhao Zhang, and Hongdong Li. Naf: neural attenuation fields for sparse-view cbct reconstruction. In *International Conference on Medical Image Computing and Computer-Assisted Intervention*, pages 442–452. Springer, 2022. 5
- [17] Wei Zhao, Don Vernekohl, Jun Zhu, Luyao Wang, and Lei Xing. A model-based scatter artifacts correction for cone beam ct. *Medical physics*, 43(4):1736–1753, 2016. 1