UC San Diego

Frequency-domain Transforms as Priors for Computer Vision

Video Processing Lab's Seminar

Hung Nguyen

Advisor: Prof. Truong Nguyen

November 2, 2026

Frequency-domain transforms

e.g., DFT, DCT, DWT

Frequency-aware loss functions

- **WGSR (CVPR 2024):** wavelet-domain loss for image super-resolution
- **FreGS (CVPR 2024):** Fourier-domain loss for 3D Gaussian Splatting
- **DWTGS (Ours, VCIP 2025):** wavelet-domain loss for sparse-view 3D Gaussian Splatting

Frequency-aware image generation

- **WaveDiff (CVPR 2023):** wavelet-domain diffusion
- **DCTDiff (ICML 2025):** cosine-domain diffusion
- **WaveletGaussian (Ours, ICASSP 2026):** wavelet-domain diffusion for sparse-view 3D Gaussian Splatting

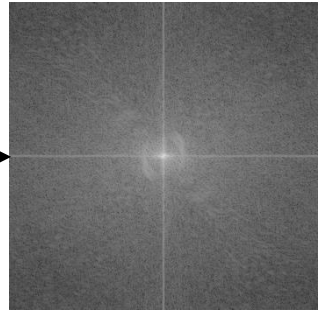
Frequency-aware modulation

- **Opti3DGS (CVMP 2025)**: Fourier-based coarse-to-fine frequency modulation for 3D Gaussian Splatting pruning
- **AutoOpti3DGS and its extension (ICCVW 2025, EUSIPCO 2026):** wavelet-based extensions



DFT

Complex exponential bases

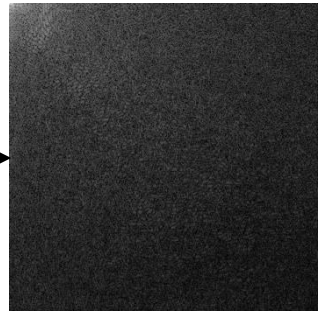


- Complex coefficients



DCT

Real cosine bases



- Real coefficients
- Energy compaction

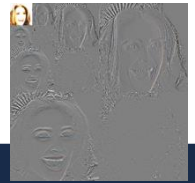


DWT

Wavelet bases

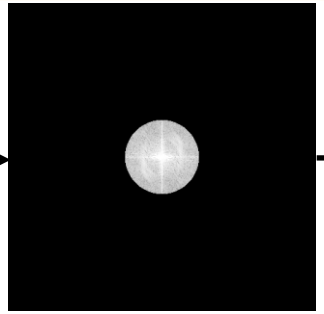


- Spatial localization
- Multi-resolution analysis





DFT

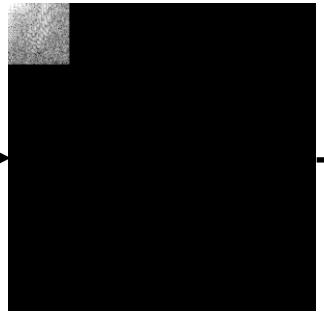


IDFT





DCT



IDCT



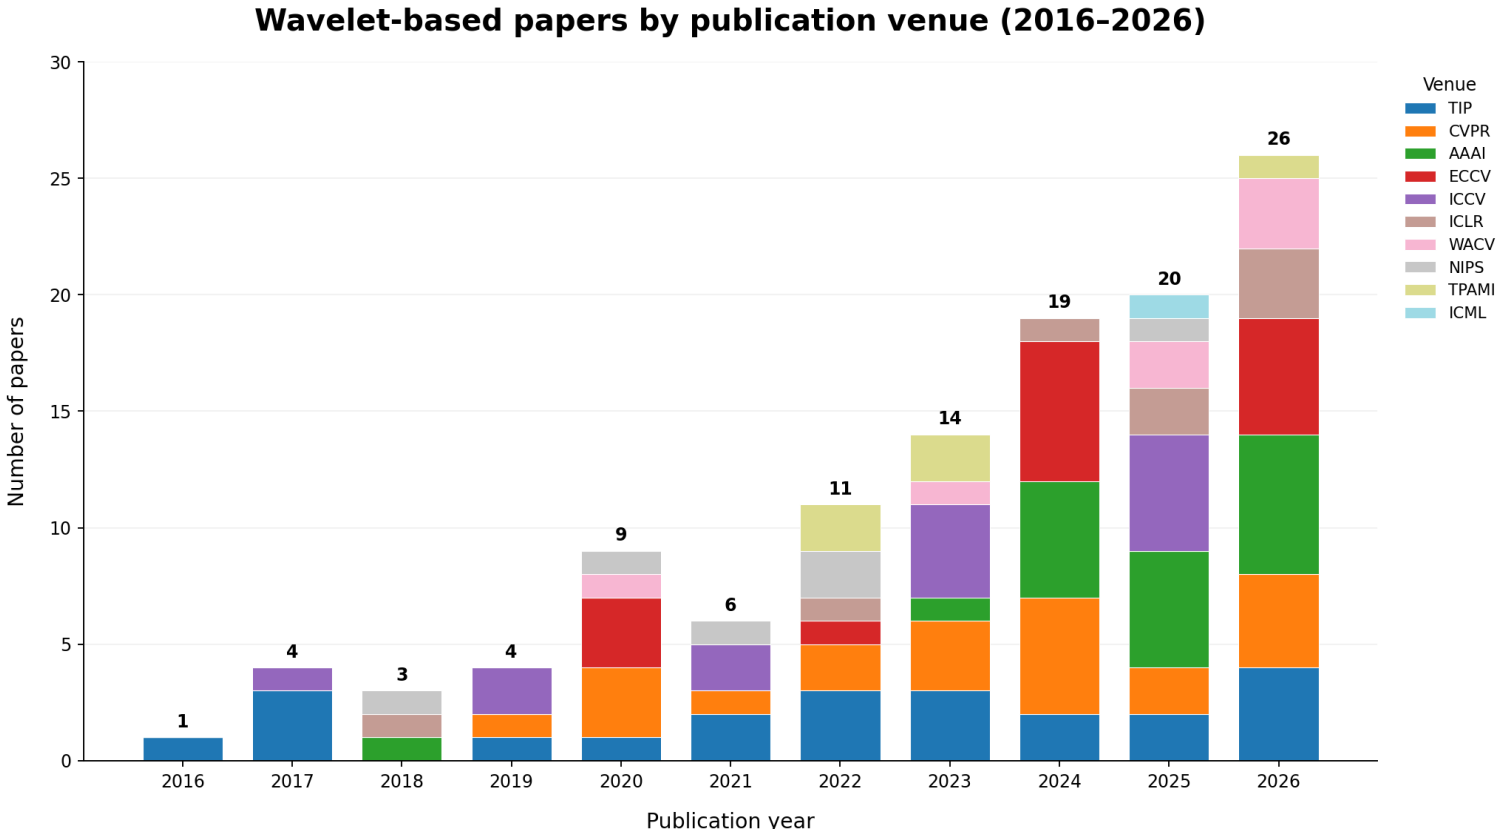


DWT

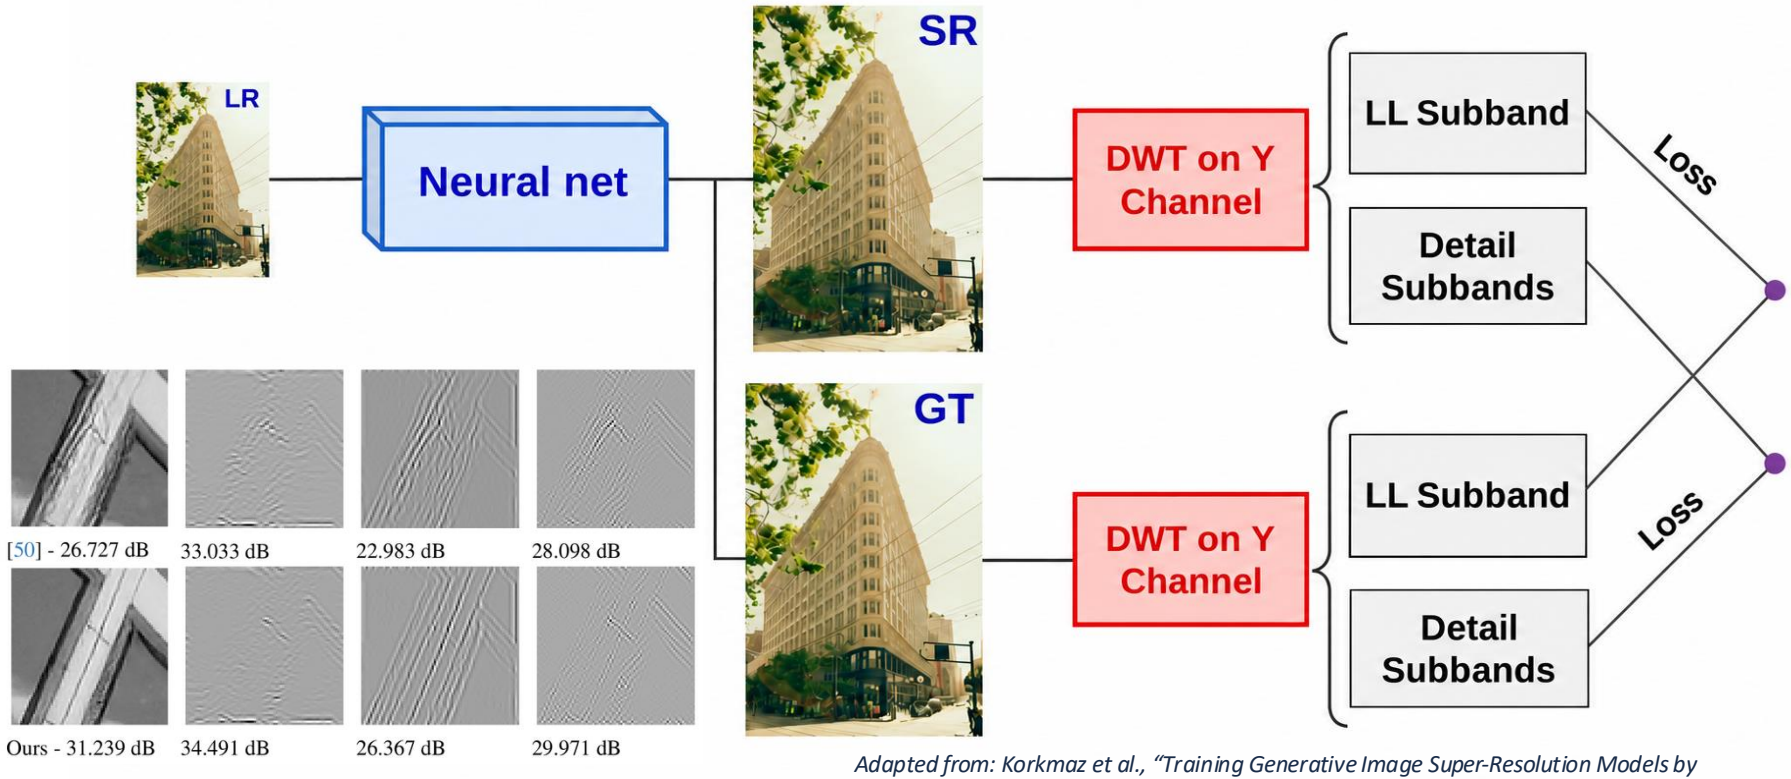


IDWT




Wavelet-based papers by publication venue (2016–2026)
Number of papers
Publication year
Venue
TIP
CVPR
AAAI
ECCV
ICCV
ICLR
WACV
NIPS
TPAMI
ICML
1
4
3
4
9
6
11
14
19
20
26
0
5
10
15
20
25
30
2016
2017
2018
2019
2020
2021
2022
2023
2024
2025
2026

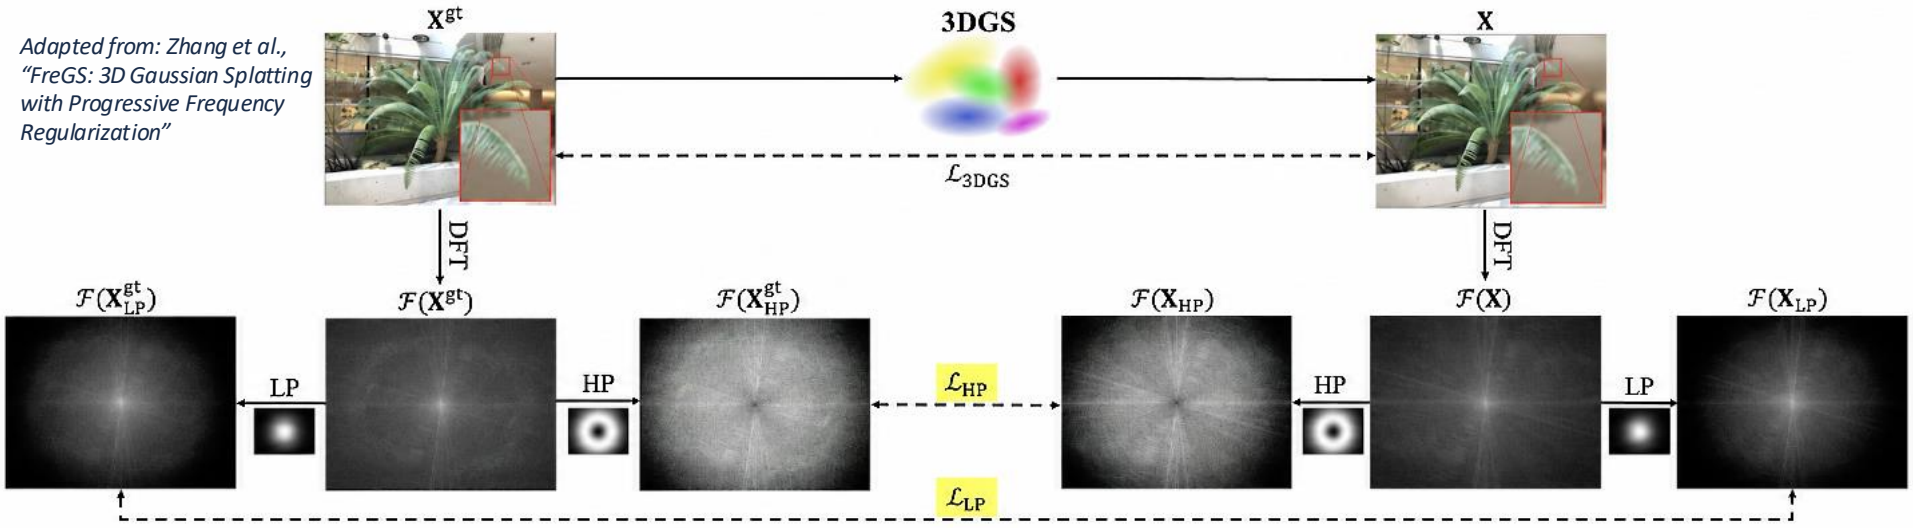
Frequency-aware loss (WGSR, CVPR 2024)

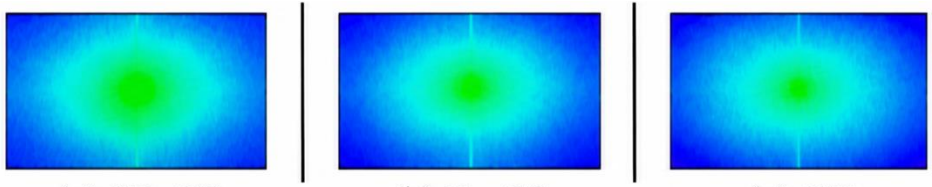


Adapted from: Korkmaz et al., "Training Generative Image Super-Resolution Models by Wavelet-Domain Losses Enables Better Control of Artifacts"

Frequency-aware loss (FreGS, CVPR 2024)

Adapted from: Zhang et al., "FreGS: 3D Gaussian Splatting with Progressive Frequency Regularization"



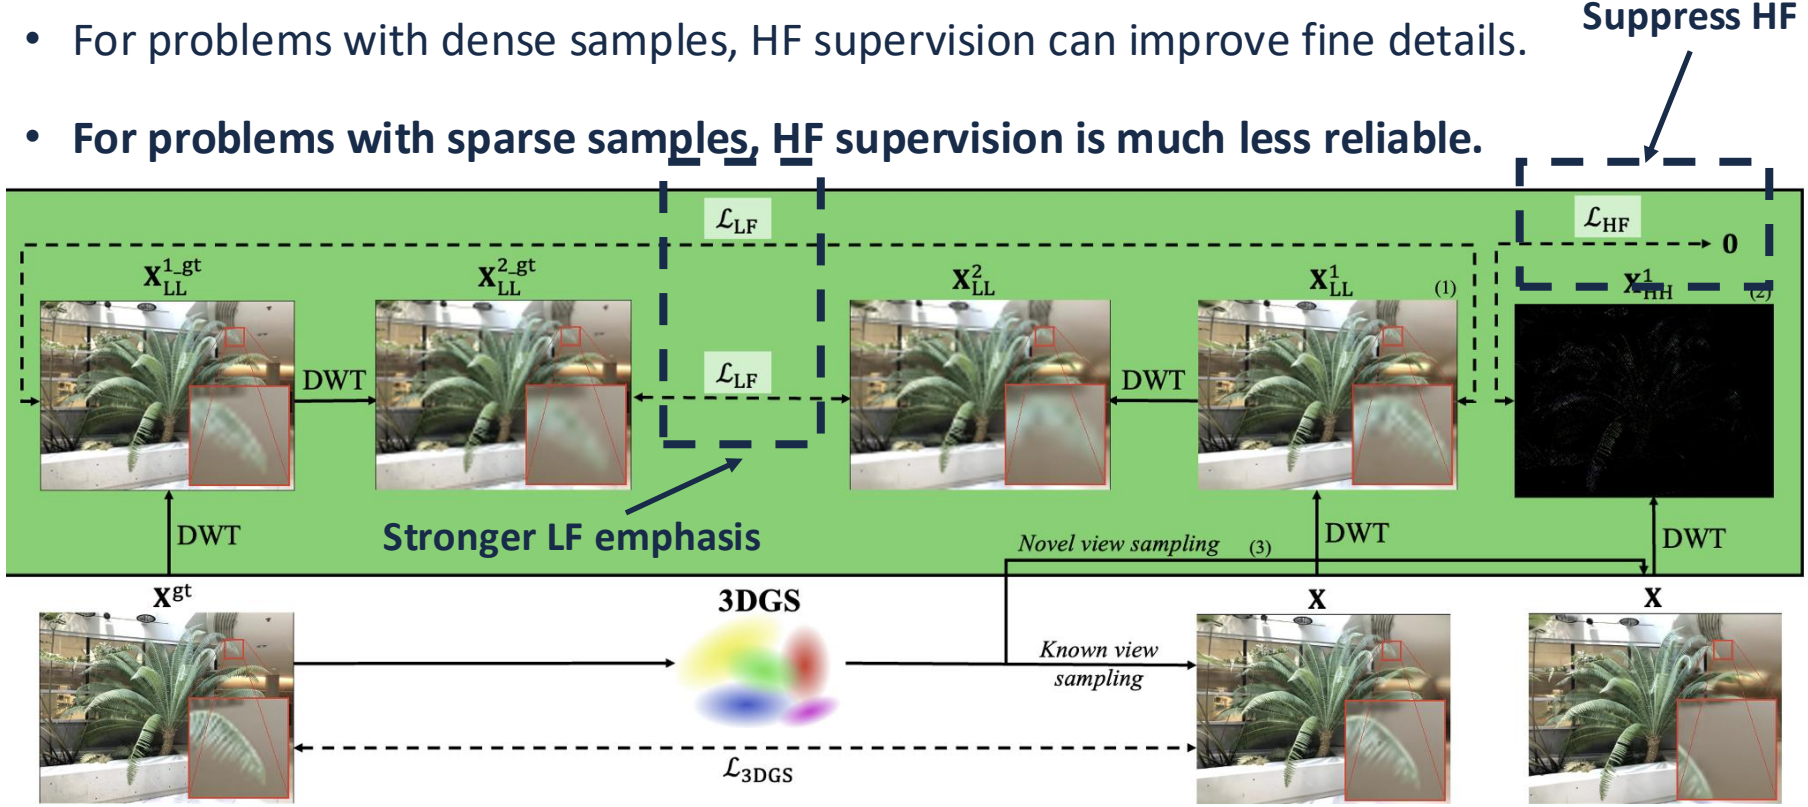


(a) 3D-GS (b) FreGS (c) GT

Datasets	Mip-NeRF360		
Methods	SSIM↑	PSNR↑	LPIPS↓
3D-GS	0.815	27.21	0.214
FreGS(Ours)	**0.826**	**27.85**	**0.209**

Frequency-aware loss (DWTGS, VCIP 2025)

- For problems with dense samples, HF supervision can improve fine details.
- **For problems with sparse samples, HF supervision is much less reliable.**



Adapted from: Nguyen et al., "DWTGS: Rethinking Frequency Regularization for Sparse-view 3D Gaussian Splatting"

Frequency-aware loss (DWTGS, VCIP 2025)



3DGS (PSNR: 22.68)



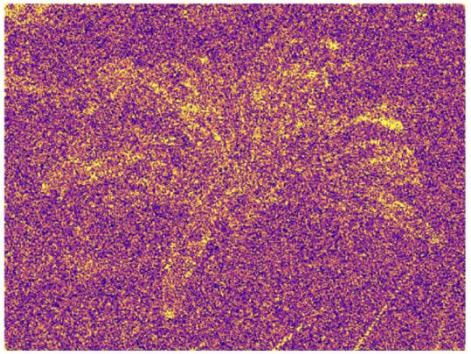
FreGS (PSNR: **22.65**)



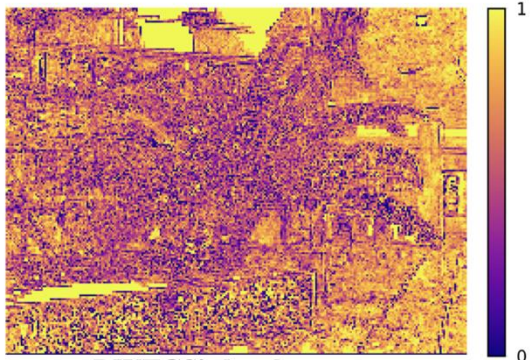
DWTGS (PSNR: **23.04**)



Ground-truth



FreGS's loss heatmap



DWTGS's loss heatmap

Adapted from: Nguyen et al., "DWTGS: Rethinking Frequency Regularization for Sparse-view 3D Gaussian Splatting"

Open Research Questions

- Related works often introduce frequency-aware loss as a subloss:

$$\mathcal{L} = \mathcal{L}_{task} + \lambda_{freq}\mathcal{L}_{freq}$$

- In the case of wavelet losses:

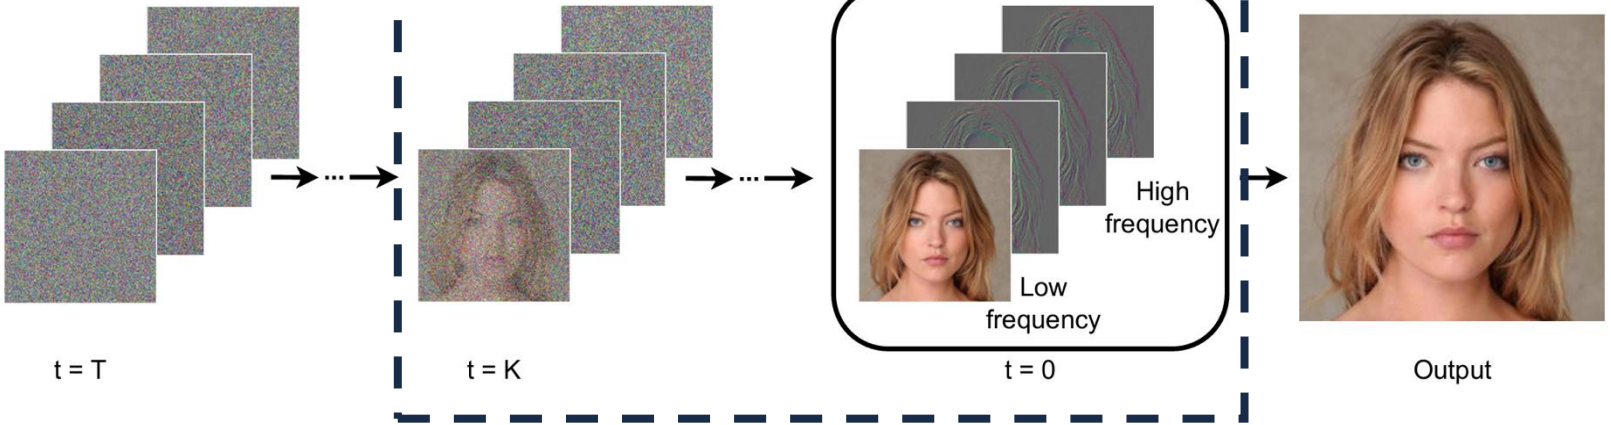
$$\mathcal{L} = \mathcal{L}_{task} + \lambda_{LL}\mathcal{L}_{LL} + \lambda_{LH}\mathcal{L}_{LH} + \lambda_{HL}\mathcal{L}_{HL} + \lambda_{HH}\mathcal{L}_{HH}$$

- **All subbands have different pixel distributions.**
- Most works require careful tuning for the subband weights.
- **Future works should address adaptive subband weight tuning.**
- Since HF subbands are sparse, pixel-wise losses might not be ideal.

Frequency-aware generation (WaveDiff, CVPR 2023)

Wavelet-domain diffusion

Adapted from: Phung et al., "Wavelet Diffusion Models are fast and scalable Image Generators"

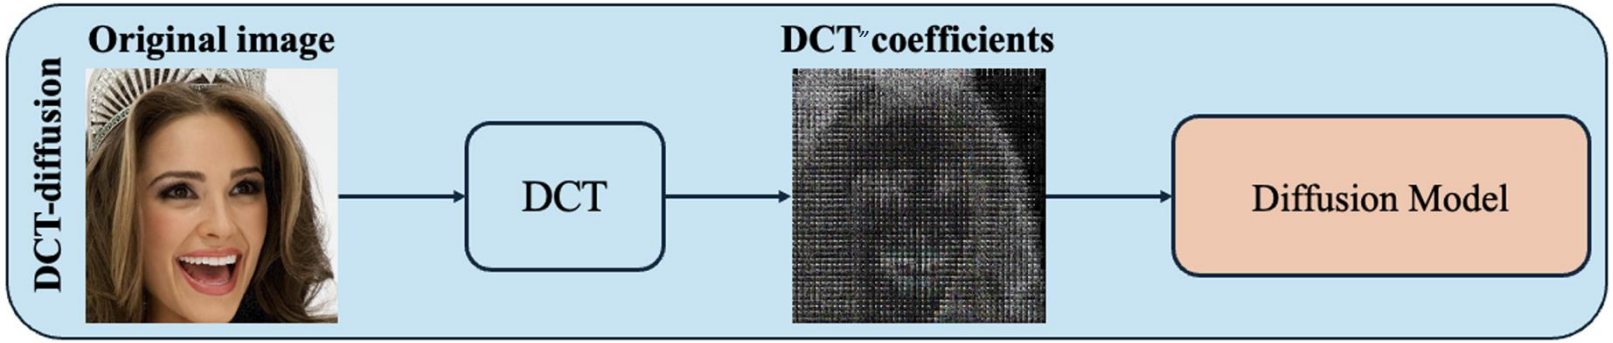




Model	FID↓	Recall↑	Time (s)↓
Ours + W-Generator	**6.40**	**0.35**	**0.59**
DDGAN [49]	8.43	0.33	1.49

Frequency-aware generation (DCTDiff, ICML 2025)

Adapted from: Ning et al., "DCTdiff: Intriguing Properties of Image Generative Modeling in the DCT Space"

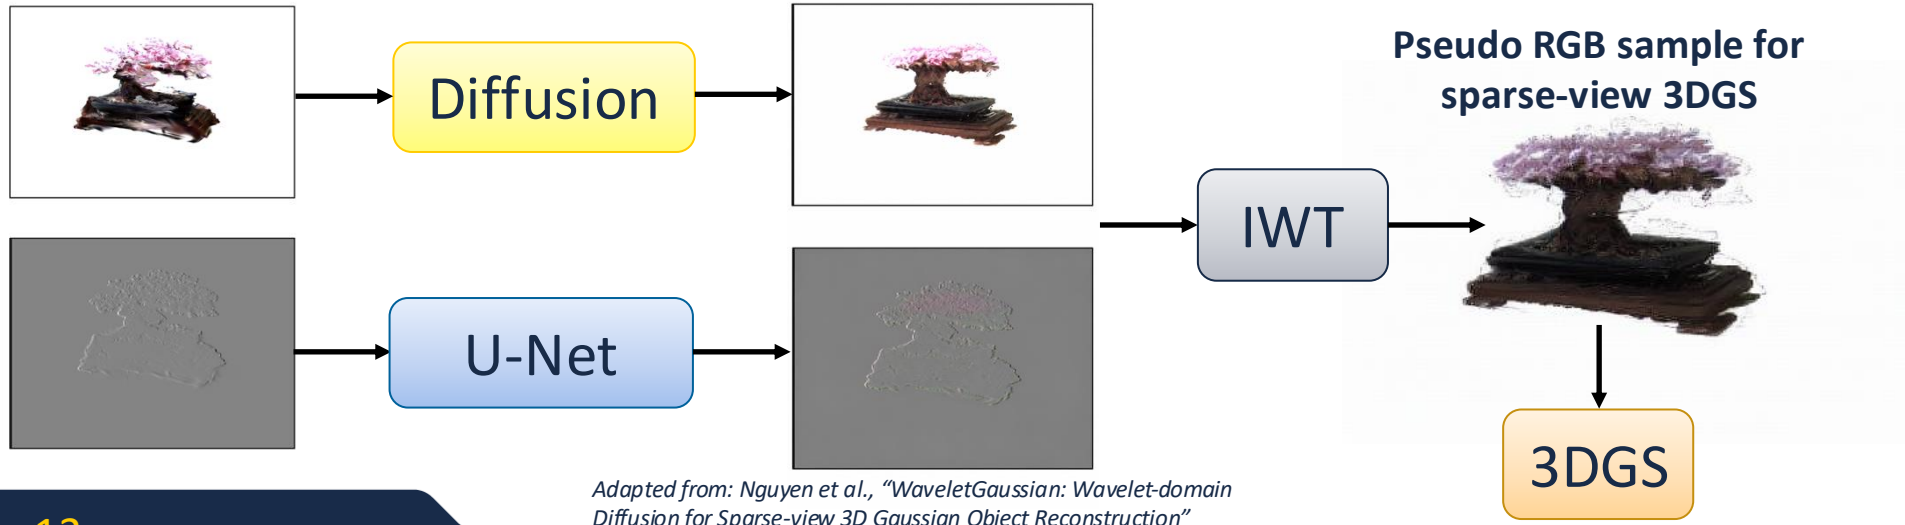




Model	Dataset	
	FFHQ 512	AFHQ 512
UViT (latent)	10.89	10.86
DCTdiff	**7.07**	**8.76**

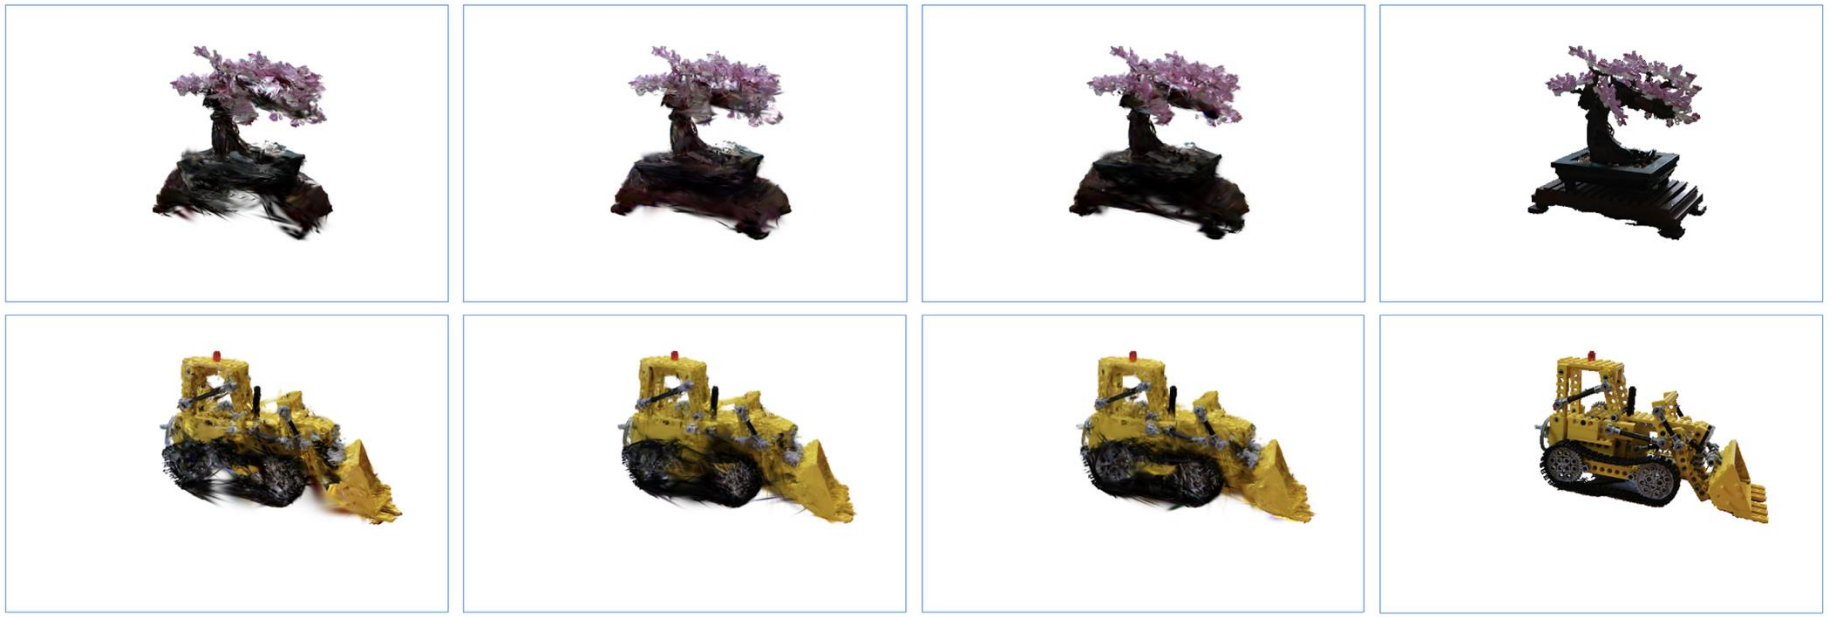
Frequency-aware generation (WaveletGaussian, ICASSP 2026)

- Wavelet-diffusion models can be accelerated by diffusing the LL only.
- **LL-diffusion works often use the frequency separation principle.**
- The HF subbands are refined through a lightweight network, e.g., U-Net.



Adapted from: Nguyen et al., "WaveletGaussian: Wavelet-domain Diffusion for Sparse-view 3D Gaussian Object Reconstruction"

Frequency-aware generation (WaveletGaussian, ICASSP 2026)



3DGS
(PSNR: 20.31)

GaussianObject
(PSNR: 24.81, 51 mins)

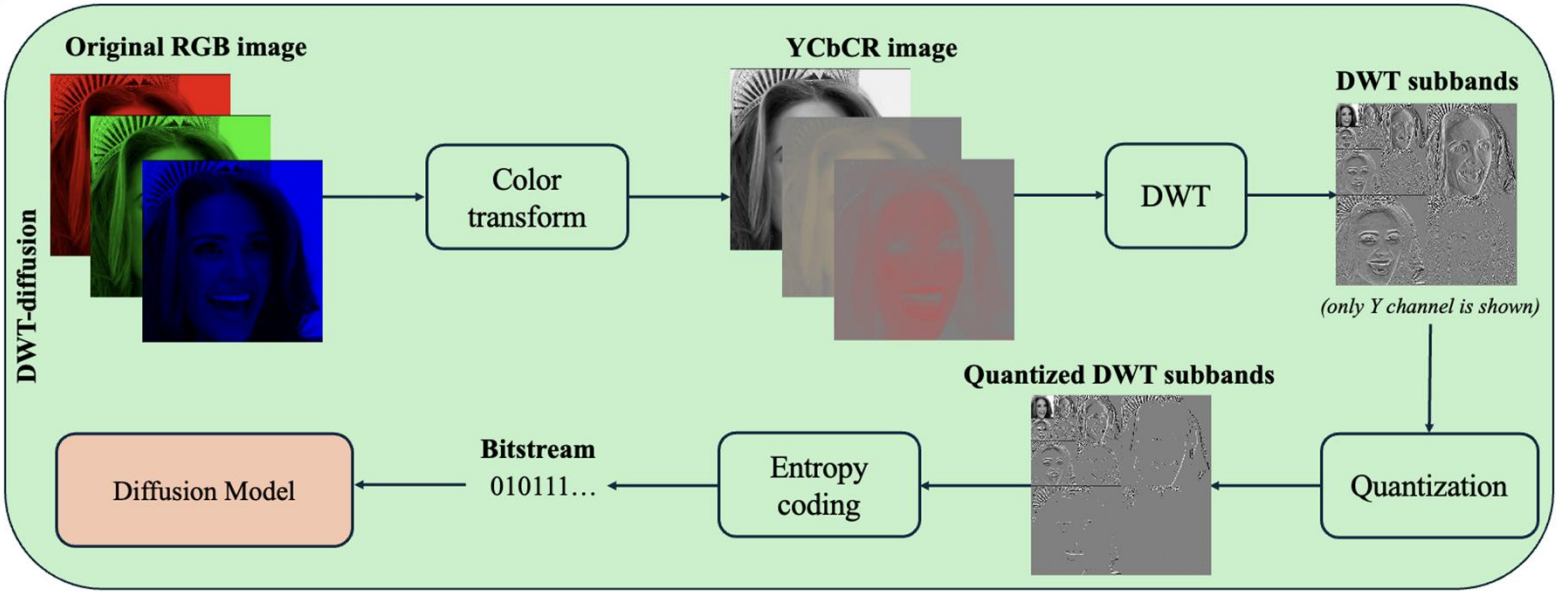
WaveletGaussian
(PSNR: **25.31**, **33** mins)

Ground-truth

Adapted from: Nguyen et al., "WaveletGaussian: Wavelet-domain Diffusion for Sparse-view 3D Gaussian Object Reconstruction"

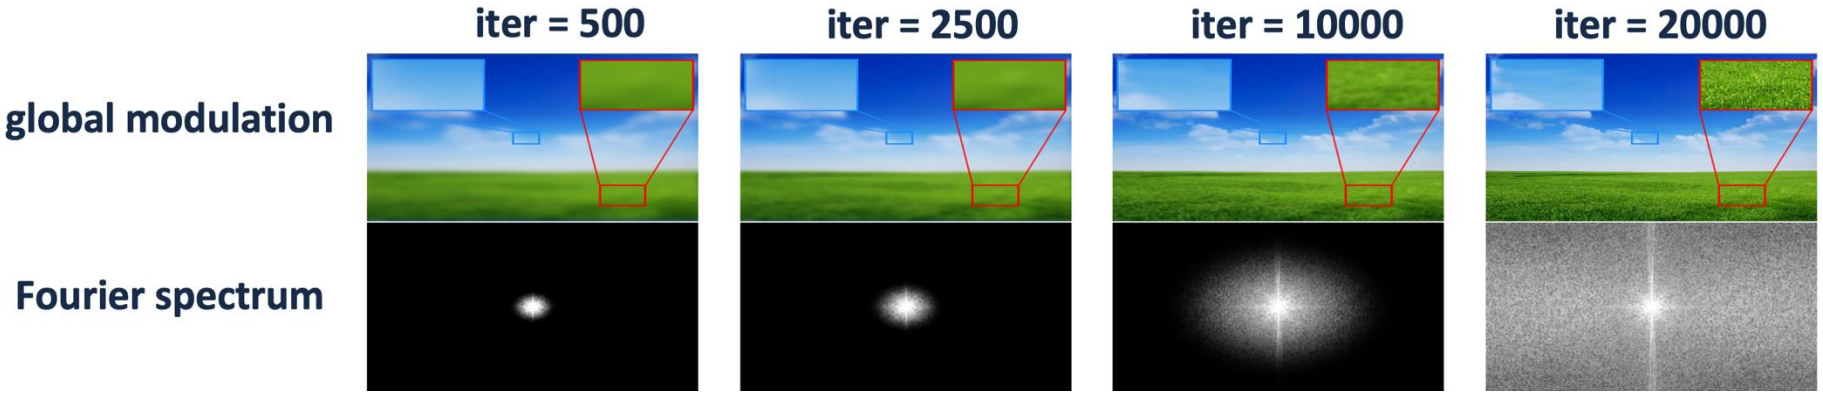
Open Research Questions

- Current methods ignore the sparsity of the DWT subbands/DCT coefficients.
- **Future works need to compact the diffused targets further.**



Frequency-aware modulation (Opti3DGS, CVMP 2025)

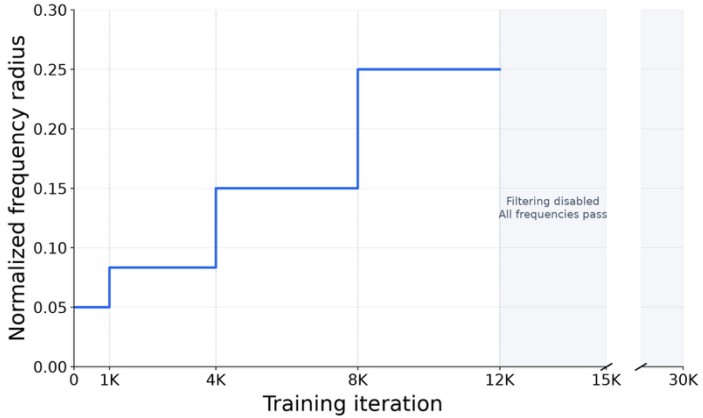
- Frequency modulation is a method for 3DGS pruning.
- **Gradually introduces HF details to the training images.**
- Prevents over-densification by only creating HF Gaussians when needed.



Adapted from: Farooq et al., "Optimized 3D Gaussian Splatting using Coarse-to-Fine Image Frequency Modulation"

Frequency-aware modulation (Opti3DGS, CVMP 2025)

Adapted from: Farooq et al., "Optimized 3D Gaussian Splatting using Coarse-to-Fine Image Frequency Modulation"



Manually define filter sizes for different training stages

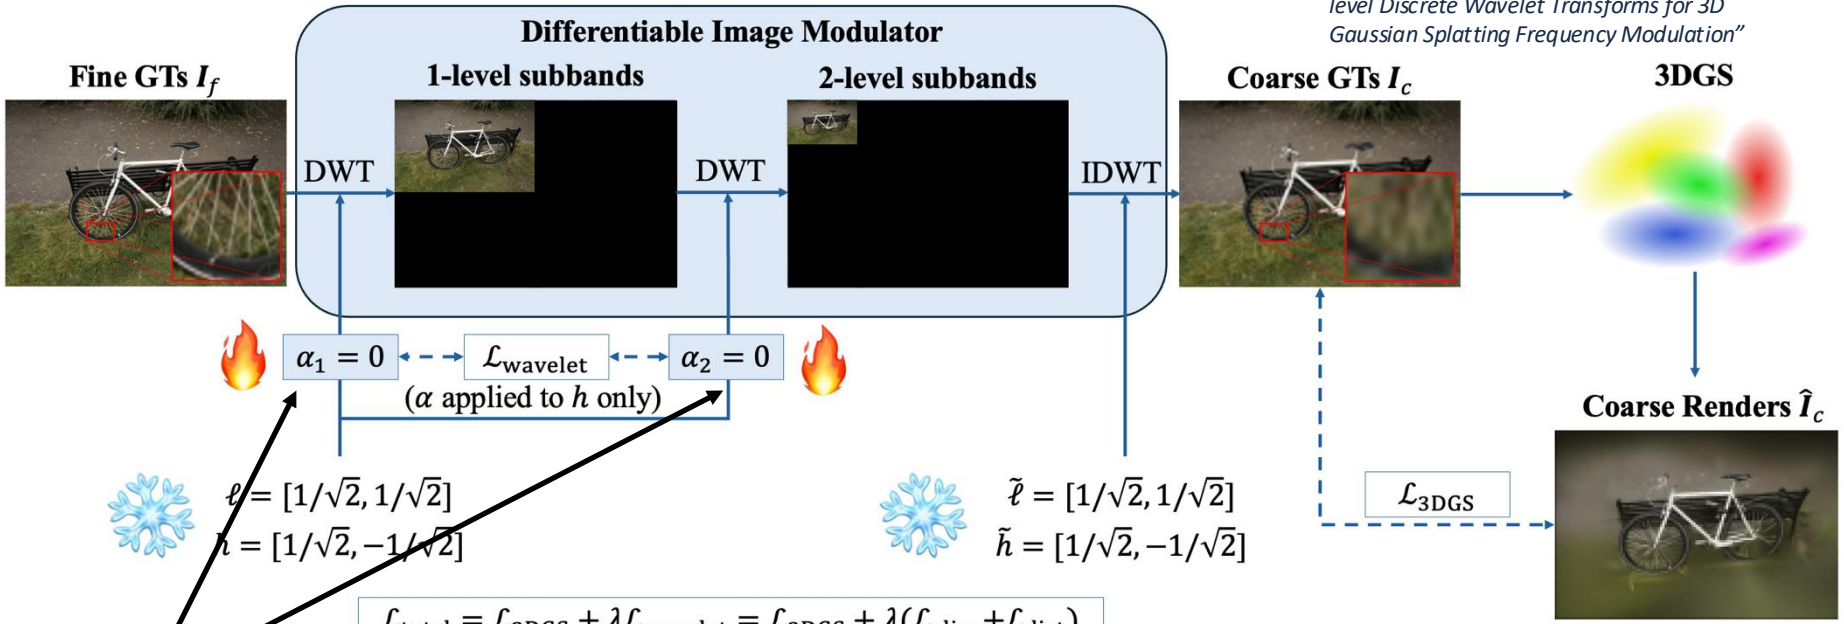
Reduces Gaussian counts while maintains quality

Method	Deep Blending[Hedman et al. 2021]						
	SSIM ↑	PSNR ↑	LPIPS ↓	FPS ↑	Time $(m:s)$ ↓	#G (10^6) ↓	Peak #G ↓
3DGS[Kerbl et al. 2023]	0.900	29.54	**0.247**	293	11:34	2.795	-
3DGS + Ours*	**0.902**	**29.56**	0.255	**611**	**10:26**	**0.953** (-66%)	-

Frequency-aware modulation (AutoOpti3DGS, ICCVW 2025 & EUSIPCO 2026)

UC San Diego

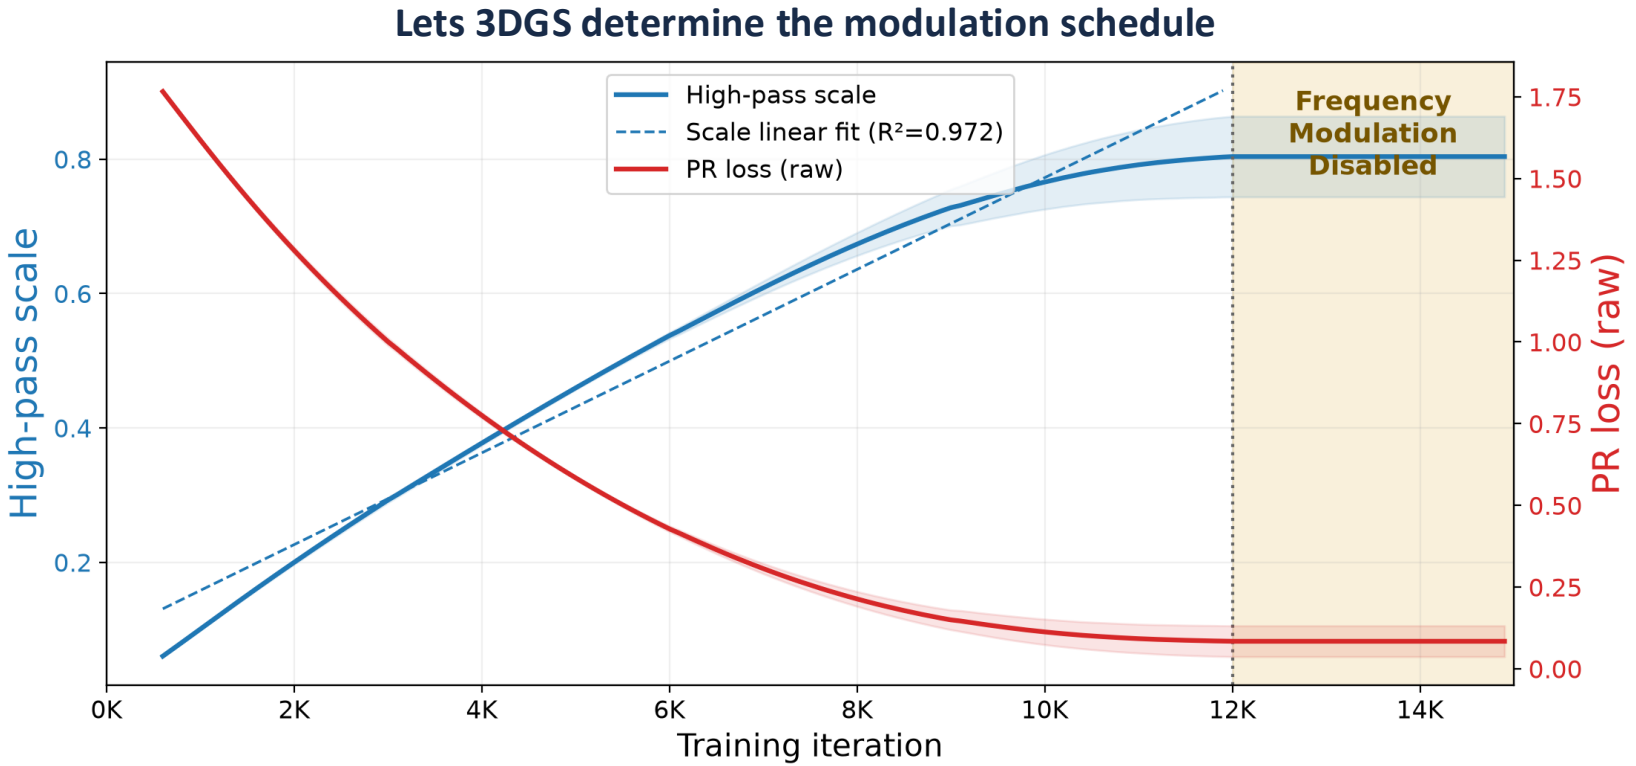
Adapted from: Nguyen et al., "Learnable Multi-level Discrete Wavelet Transforms for 3D Gaussian Splatting Frequency Modulation"



ZERO-INITIALIZE THE HP FILTER'S GAIN AND INVERSE DWT → COARSE RECONSTRUCTION.

A PERFECT RECONSTRUCTION-BASED SUBLOSS ENCOURAGES THE GAIN TO CONVERGE → 1, RESTORING DETAILS.

Frequency-aware modulation
(AutoOpti3DGS, ICCVW 2025 & EUSIPCO 2026)



Frequency-aware modulation
(AutoOpti3DGS, ICCVW 2025 & EUSIPCO 2026)



PSNR: 18.90, **Gaussian Counts:** 352K
3DGS



PSNR: 18.78, **Gaussian Counts:** 301K
Opti3DGS



PSNR: 19.13, **Gaussian Counts:** 257K
AutoOpti3DGS (Ours)

Adapted from: Nguyen et al., "From Coarse to Fine: Learnable Discrete Wavelet Transforms for Efficient 3D Gaussian Splatting"

Open Research Questions

- Extends the frequency modulation principle to sparse-view 3DGS and 4DGS.
- Integrates the principle into SOTA 3DGS compression methods.
- Explores better wavelet formulations to achieve better rate-distortion.

Conclusion

- Conclusion
  - Frequency-domain priors are increasingly injected into computer vision architectures.
  - Three covered applications: loss function, image generation, modulation.
- Next steps
  - Designs adaptive filters for data-dependent frequency decomposition.
  - Applies the priors to non-CV problems, e.g., wireless communications, robot learning.
- Questions?