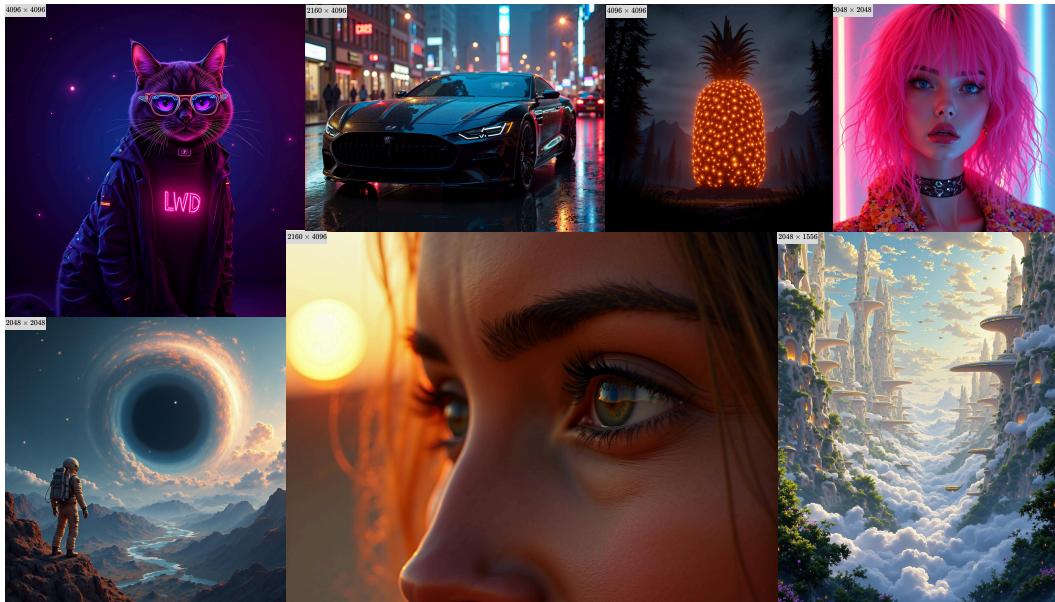


000 001 002 003 004 005 006 007 008 009 010 011 012 013 014 015 016 017 018 019 020 021 022 023 024 025 LATENT WAVELET DIFFUSION FOR ULTRA-HIGH-RESOLUTION IMAGE SYNTHESIS

005 **Anonymous authors**

006 Paper under double-blind review



026 Figure 1: We propose Latent Wavelet Diffusion, achieving 4K image synthesis without architectural
027 changes or additional inference cost to existing latent diffusion models.

028 ABSTRACT

031 High-resolution image synthesis remains a core challenge in generative modeling,
032 particularly in balancing computational efficiency with the preservation of fine-
033 grained visual detail. We present *Latent Wavelet Diffusion (LWD)*, a lightweight
034 training framework that significantly improves detail and texture fidelity in ultra-
035 high-resolution (2K-4K) image synthesis. LWD introduces a novel, frequency-
036 aware masking strategy derived from wavelet energy maps, which dynamically
037 focuses the training process on detail-rich regions of the latent space. This is com-
038 plemented by a scale-consistent VAE objective to ensure high spectral fidelity. The
039 primary advantage of our approach is its efficiency: LWD requires no architectural
040 modifications and adds zero additional cost during inference, making it a practical
041 solution for scaling existing models. Across multiple strong baselines, LWD con-
042 sistently improves perceptual quality and FID scores, demonstrating the power of
043 signal-driven supervision as a principled and efficient path toward high-resolution
044 generative modeling.

045 1 INTRODUCTION

048 Diffusion models have become a dominant paradigm in generative modeling, achieving state-of-the-
049 art results in tasks such as image synthesis (Podell et al., 2024; Sauer et al., 2024; Lopez et al., 2025;
050 Wang et al., 2025), super-resolution (Saharia et al., 2022; Sigillo et al., 2024; Wang et al., 2024),
051 and image editing (Brooks et al., 2023; Meng et al., 2022). Foundational methods like Denoising
052 Diffusion Probabilistic Models (DDPM) Ho et al. (2020) and Denoising Diffusion Implicit Models
053 (DDIM) Song et al. (2021a) have enabled increasingly powerful variants. Latent Diffusion Models
(LDMs) Rombach et al. (2022) improve efficiency by operating in a learned latent space, while recent

054 architectures such as Diffusion Transformers (DiTs) (Peebles & Xie, 2023; Esser et al., 2024) further
 055 enhance scalability and modeling capacity.
 056

057 Despite recent advances, generating ultra-high-resolution (UHR) images at resolutions of 2K to
 058 4K and beyond remains a significant challenge. Simply scaling models trained on lower-resolution
 059 data often fails to generalize, leading to repeated structures, blurred textures, and spatial inconsis-
 060 tencies (Bar-Tal et al., 2023; He et al., 2023). Naive alternatives, such as hierarchical generation
 061 pipelines and post-hoc super-resolution techniques (Du et al., 2024a; He et al., 2023), typically
 062 produce oversmoothed outputs. To overcome these limitations, several strategies have been ex-
 063 plored. One direction involves direct UHR training or fine-tuning (Ren et al., 2024; Chen et al.,
 064 2024), although this typically demands extensive computational resources, access to proprietary high-
 065 resolution datasets (Liu et al., 2024b), and substantial GPU memory for large model backbones (Labs,
 066 2024; Esser et al., 2024). Other efforts focus on architectural modifications to improve long-range
 067 dependency modeling (Liu et al., 2024a), or optimization techniques that enhance the quality of latent
 068 representations (Hahm et al., 2024).

069 Across these varied approaches, a common limitation persists: most methods treat all spatial regions
 070 uniformly during generation, applying the same refinement process to areas with vastly different
 071 visual complexity. This uniformity disregards local frequency variation, failing to distinguish between
 072 smooth regions and areas rich in textures, edges, or semantic structure. The consequences are twofold.
 073 First, computation is wasted on low-detail areas that require minimal refinement. Second, high-detail
 074 regions are not sufficiently supervised, leading to artifacts or loss of fidelity. These issues stem
 075 from both architectural and algorithmic constraints: latent representations often lack the structural
 076 granularity required for UHR synthesis, and current diffusion training objectives do not incorporate
 077 spatial adaptivity into the denoising process. Together, these limitations present a core bottleneck for
 078 perceptually accurate ultra-high-resolution image generation.

079 In this work, we propose *Latent Wavelet Diffusion* (LWD), a general and modular framework that
 080 introduces frequency-sensitive spatial supervision into the latent denoising process of diffusion
 081 models. LWD is motivated by the observation that different regions of an image exhibit varying levels
 082 of structural complexity and perceptual importance. While some areas contain intricate textures or
 083 sharp edges, others are homogeneous or low in detail. Our goal is to exploit this spatial heterogeneity
 084 by allocating greater learning signal to regions with high visual complexity, and reducing supervision
 085 in low-detail areas. Importantly, LWD achieves this adaptivity without modifying the underlying
 086 architecture of the diffusion model, making it broadly applicable across model families.

087 The LWD framework consists of three key components:

- 088 1. A **spectrally-aware VAE fine-tuning objective** that improves the structure and diffusability
 089 of latent representations. By encouraging semantic consistency and frequency regularity, this
 090 objective enhances the suitability of latent spaces for high-resolution generation. It serves as the
 091 foundation for the subsequent components of LWD.
- 092 2. A **wavelet-derived spatial saliency map**, computed via a discrete wavelet transform (DWT)
 093 applied to the latent representation. This map aggregates the local energy of high-frequency
 094 subbands (LH, HL, HH) and is normalized to highlight spatial regions with strong structural detail.
 095 These saliency maps are interpretable, data-driven, and require no additional training, providing a
 096 principled measure of spatial importance directly from the signal.
- 097 3. A **time-dependent masking strategy** that leverages the frequency-based saliency maps to modu-
 098 late the training loss. At each spatial location and timestep, a binary mask determines whether
 099 the denoising loss is applied. Regions with high wavelet energy receive supervision across more
 100 timesteps, while simpler areas are updated less frequently. This mechanism introduces spatial
 101 adaptivity into the learning process and improves the fidelity of fine-scale detail.

102 LWD is compatible with a broad range of latent diffusion models, including both classical diffusion
 103 and flow-matching formulations. Because it operates solely through the training objective, LWD can
 104 be seamlessly integrated into existing pipelines. While ultra-high-resolution generation is inherently
 105 computationally demanding due to the backbone’s scaling properties, LWD incurs zero marginal
 106 cost relative to the baseline. It requires no architectural changes or cascaded upsamplers, making
 107 it a practical solution for improving the generation quality of existing models. We demonstrate its
 108 flexibility and effectiveness by applying it to several state-of-the-art latent diffusion models and
 109 evaluating performance on ultra-high-resolution image synthesis (2K to 4K). Experimental results

108 show that LWD consistently enhances perceptual quality, improves FID scores, and better preserves
 109 fine-grained detail, all without increasing inference complexity.
 110

111 2 RELATED WORK

113 **Diffusion Models and Latent Generation.** Diffusion models have become foundational in generative
 114 modeling, particularly for image synthesis (Shen et al., 2025; Zhan et al., 2024), by progressively
 115 denoising Gaussian noise using a learned score function. Variants based on stochastic differential
 116 equations (Song et al., 2021b), probability flow ordinary differential equations (Lipman et al., 2023),
 117 and reinforcement-trained objectives (Black et al., 2024) have expanded the design space with
 118 improved stability and sampling efficiency.
 119

120 To reduce the cost of high-resolution generation, Latent Diffusion Models (LDMs) (Rombach et al.,
 121 2022) perform synthesis in a compressed space learned via variational autoencoders (VAEs). However,
 122 generation quality is closely tied to the spectral fidelity and structure of these latent representations.
 123 Prior work has sought to improve this through enhanced VAE architectures (Esser et al., 2021),
 124 hierarchical compression (Takida et al., 2024), and frequency-aware regularization (Skorokhodov
 125 et al., 2025; Kouzelis et al., 2025). We build on this direction by integrating frequency-sensitive
 126 supervision both during encoding and throughout the denoising process.
 127

128 **Flow Matching and High-Resolution Diffusion.** Flow matching Lipman et al. (2023); Esser et al.
 129 (2024) offers an alternative to classical diffusion by learning a continuous velocity field that maps
 130 noise to data in latent space, eliminating the need for fixed noise schedules. This formulation
 131 underlies models such as Flux Labs (2024), which, paired with DiT backbones, has demonstrated
 132 strong performance in ultra-high-resolution pipelines (Zhang et al., 2025; Yu et al., 2025). Our
 133 method extends this family by introducing frequency-based spatial masking into the flow-matching
 134 objective. Through wavelet decomposition of the latent space, LWD computes spatial saliency maps
 135 that guide targeted supervision toward detail-rich regions, enhancing fine structure without modifying
 136 model architecture.
 137

138 **Variational Autoencoder Optimization.** The performance of latent diffusion models at high resolu-
 139 tions depends critically on the expressiveness and spectral consistency of the VAE. Improvements in-
 140 clude multi-scale encoders (Vahdat & Kautz, 2020; Takida et al., 2024), spectral loss functions (Björk
 141 et al., 2022), and scale-consistency constraints (Skorokhodov et al., 2025; Kouzelis et al., 2025).
 142 Wavelet-based methods (Esteves et al., 2024; Lin et al., 2024; Agarwal et al., 2025) enrich latent
 143 expressiveness by isolating frequency components, while compression-oriented approaches (Xie
 144 et al., 2025; Tang et al., 2024; HaCohen et al., 2024) aim to reduce token count for improved sampling
 145 efficiency. Our LWD fine-tunes a pretrained VAE with a scale-consistent spectral loss that sup-
 146 presses spurious high-frequency noise. This regularized latent space facilitates downstream wavelet
 147 decomposition and supports our spatially adaptive denoising objective.
 148

149 **Ultra High-Resolution Image Synthesis.** Maintaining global structure and fine detail in ultra-
 150 resolution generation is a persistent challenge. Standard diffusion models tend to produce repetitive
 151 patterns or distortions (He et al., 2023). Existing solutions include cascaded generation (Ho et al.,
 152 2022), progressive upsampling (Gu et al., 2024), domain-specific pipelines (Ren et al., 2024), and
 153 latent-space super-resolution (Jeong et al., 2025). However, training-based methods (Xie et al., 2023;
 154 Zheng et al., 2024; Guo et al., 2024; Chen et al., 2024) are resource-intensive, and training-free
 155 approaches (Bar-Tal et al., 2023; Lee et al., 2023) often yield local artifacts.
 156

157 Methods such as ScaleCrafter He et al. (2023) mitigate repetition through dilated convolutions but
 158 may distort structure, while ResMaster Shi et al. (2025) uses low-resolution references for guided
 159 refinement. HiDiffusion Zhang et al. (2023) introduces architectural changes that risk performance
 160 trade-offs, and progressive strategies like DemoFusion Du et al. (2024a) suffer from slow inference
 161 and irregular patterns. Diffusion-4K Zhang et al. (2025) and URAE Yu et al. (2025) advance
 162 latent modeling and parameter-efficient adaptation at 4K resolution. Our LWD complements these
 163 approaches by introducing signal-driven, spatially adaptive supervision, which improves structural
 164 and perceptual fidelity at no additional cost.
 165

166 **Generative Modeling in the Frequency Domain.** Frequency structure plays an increasingly
 167 important role in generative modeling (Yang et al., 2023). Wavelet-based diffusion methods, such
 168 as WaveDiff Phung et al. (2023), and spectral decomposition approaches have proven useful for
 169 efficient sampling (Qian et al., 2024), super-resolution (Aloisi et al., 2024; Sigillo et al., 2025), and
 170

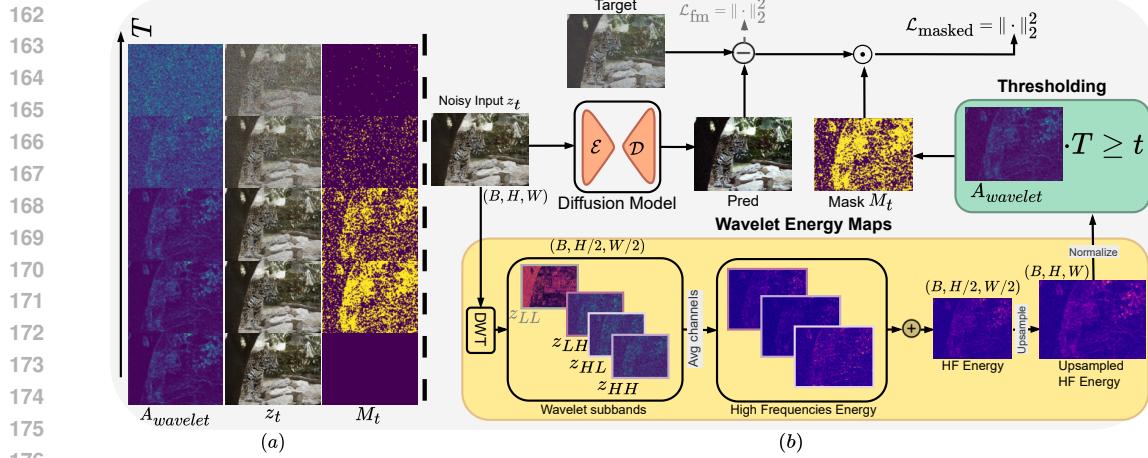


Figure 2: (a) Temporal evolution of latent z_t , wavelet energy maps A_{wavelet} , and attention map M_t across diffusion timesteps. (b) Our wavelet-masked flow matching objective at a timestep t . The model computes a wavelet attention map M_t from latent z_t to modulate the prediction error between target velocity field ($\epsilon - z_0$) and predicted velocity $v_{\Theta}(z_t, t, y)$. This focuses optimization on high-frequency regions with greater perceptual importance. While operations occur in latent space, decoded visualizations are shown for interpretability.

restoration (Huang et al., 2024b; Zhao et al., 2024; Jiang et al., 2023). FouriScale Huang et al. (2024a) demonstrated that frequency filtering enhances coherence, and DiffuseHigh Kim et al. (2025) leveraged low-frequency DWT guidance to improve global structure in UHR synthesis.

Diffusion-4K (Zhang et al., 2025) incorporated wavelet losses in the latent space to balance frequency bands, but applied them uniformly across all spatial locations. In contrast, LWD introduces wavelet-based spatial conditioning through time-dependent masking. Rather than treating frequency as a passive loss signal, our LWD actively uses local frequency energy to modulate supervision across space and time, concentrating on learning where detail matters most. This enables sharper, more coherent synthesis without modifying the underlying model or increasing inference cost.

3 METHODOLOGY

Latent Wavelet Diffusion (LWD) introduces frequency-aware supervision into latent diffusion models by coupling signal-driven saliency analysis with adaptive training. Our key insight is that structural complexity in images is unevenly distributed in space, yet most denoising models refine all positions equally. LWD addresses this by modulating the supervision schedule based on local frequency content, improving detail fidelity without increasing computational cost.

LWD operates in two sequential stages. The first stage fine-tunes a variational autoencoder (VAE) using a scale-consistency objective. This independent step yields spectrally stable latent representations, preparing the space for the subsequent frequency-based modulation.

In the second stage, we fine-tune a latent diffusion model (e.g., Flux) using a modified flow-matching objective that integrates frequency-guided supervision. This stage incorporates three tightly coupled components: (1) extraction of wavelet-based spatial saliency maps from latent codes; (2) construction of a time-dependent mask that adapts the training signal based on local frequency energy; and (3) application of this mask to modulate the training loss dynamically across spatial positions and timesteps. Together, these mechanisms introduce spatial adaptivity into the denoising process, directing learning resources toward detail-rich regions. All components are model-agnostic and can be integrated into standard latent diffusion pipelines without architectural changes.

3.1 VAE FINE-TUNING WITH SCALE-CONSISTENCY LOSS

High-resolution generation places unique demands on the latent space: it must retain both semantic structure and spectral coherence across scales. To ensure this, we fine-tune the variational autoencoder

(VAE) using a multi-resolution reconstruction objective that regularizes frequency content while preserving perceptual fidelity.

Formally, let $z = E(x)$ be the latent encoding of image x , x_{down} a downsampled version of x , and z_{down} the downsampled version of the latent z . Our loss combines four components:

$$\mathcal{L}_{\text{VAE}} = \underbrace{\|D(z) - x\|_2^2}_{\text{Reconstruction}} + \alpha \underbrace{\|D(E(z_{\text{down}})) - x_{\text{down}}\|_2^2}_{\text{Scale Consistency}} + \beta \underbrace{D_{\text{KL}}(q(z|x) \| p(z))}_{\text{Latent Regularization}} + \lambda \underbrace{\mathcal{L}_{\text{LPIPS}}(D(z), x)}_{\text{Perceptual Loss}}, \quad (1)$$

Here, we incorporate a scale-consistency term (Skorokhodov et al., 2025; Kouzelis et al., 2025) that encourages the VAE to maintain structural coherence across resolution scales. While originally proposed for general reconstruction, we identify it as critical for wavelet-guided UHR synthesis. Without this regularization, standard VAEs exhibit spurious high-frequency noise that confounds downstream wavelet masking. This preprocessing naturally suppresses compression artifacts while preserving essential structural information in z , enabling our masking strategy to target meaningful details rather than noise.

Unlike recent approaches that inject frequency conditioning directly into the encoder (Aloisi et al., 2024; Zhang et al., 2025), our formulation decouples signal regularization and generation: we first sculpt the latent space to exhibit desirable frequency properties, and then use this structure to guide the training of the diffusion model. This preserves architectural modularity while enabling effective frequency-aware supervision.

3.2 WAVELET-DERIVED FREQUENCY SALIENCY MAPS

To guide spatial supervision based on structural complexity, we extract saliency maps from latent representations using localized frequency analysis. Given a latent tensor $z \in \mathbb{R}^{C \times H \times W}$, we apply a single-level Discrete Wavelet Transform (DWT), producing four subbands:

$$\text{DWT}(z) \rightarrow \{z_{LL}, z_{LH}, z_{HL}, z_{HH}\}, \quad (2)$$

where z_{LL} captures low-frequency approximations and $\{z_{LH}, z_{HL}, z_{HH}\}$ encode directional high-frequency detail.

We compute the localized high-frequency energy as:

$$E(i, j) = \frac{1}{C} \sum_c \left[(z_{LH}^{c, i, j})^2 + (z_{HL}^{c, i, j})^2 + (z_{HH}^{c, i, j})^2 \right], \quad (3)$$

where (i, j) denotes spatial position and $c \in \{1, \dots, C\}$ indexes feature channels. The resulting energy map $E \in \mathbb{R}^{H/2 \times W/2}$ is bilinearly upsampled and min-max normalized per sample to obtain the final saliency map $A_{\text{wavelet}} \in [0, 1]^{H \times W}$.

This map serves as a proxy for local structural richness, highlighting regions of the latent space associated with high-frequency content (e.g., textures, contours, transitions). Unlike learned attention mechanisms based on semantic similarity (e.g., DINO (Caron et al., 2021)), our approach is deterministic and directly derived from signal properties. While we refer to A_{wavelet} as an "attention map" for interpretability, it is best understood as a frequency-aware saliency measure.

Notably, our VAE fine-tuning objective (Eq. 1) helps suppress high-frequency artifacts and stabilize spectral behavior. This preprocessing step ensures that high-frequency activations captured by the DWT correspond to meaningful structure, rather than encoding noise, thereby improving the utility of A_{wavelet} for guiding spatial supervision.

270 3.3 ADAPTIVE FLOW MATCHING WITH FREQUENCY-GUIDED MASKING
271272 We adopt a continuous-time flow-matching formulation (Lipman et al., 2023; Esser et al., 2024) for
273 training the latent diffusion model. Given a target latent z_0 and noise sample $\epsilon \sim \mathcal{N}(0, I)$, we define
274 interpolated samples as:

275
$$z_t = (1 - t) z_0 + t \epsilon, \quad t \in [0, 1], \quad (4)$$

276 and supervise the predicted velocity field $v_\Theta(z_t, t, y)$, conditioned on text y , via:

277 278
$$\mathcal{L}_{\text{fm}} = \|(\epsilon - z_0) - v_\Theta(z_t, t, y)\|_2^2. \quad (5)$$

279

280 To incorporate frequency-awareness into the training objective, we apply spatially adaptive masking
281 based on the wavelet saliency map A_{wavelet} . Specifically, for each location (i, j) , we define a time-
282 dependent binary mask:

283 284
$$M_t(i, j) = \begin{cases} 1 & \text{if } T \cdot (A_{\text{wavelet}}(i, j) + \ell) \geq t \\ 0 & \text{otherwise} \end{cases}, \quad (6)$$

285

286 where T is the total number of diffusion timesteps and $\ell \in (0, 1)$ sets a lower bound on refinement.
287 This ensures that all regions receive at least ℓT steps of supervision, while high-frequency regions
288 benefit from extended refinement.289 While selective spatial supervision has been investigated using transformer attention (Moser et al.,
290 2025), our wavelet-derived saliency offers a fundamentally different perspective based on signal
291 processing principles rather than learned semantic features, offering computational advantages and a
292 more generalizable solution.

293 The final masked loss becomes:

294 295
$$\mathcal{L}_{\text{masked}} = \|M_t \odot [(\epsilon - z_0) - v_\Theta(z_t, t, y)]\|_2^2, \quad (7)$$

296

297 where \odot denotes element-wise multiplication. This formulation focuses learning capacity on detail-
298 rich regions, improves fidelity in high-frequency content, and does so without increasing inference
299 complexity. Crucially, this mechanism operates purely at the objective level and is compatible with
300 any latent diffusion model using a flow-based or score-based trajectory.301 4 EXPERIMENTS
302303 304 4.1 EXPERIMENTAL SETUP
305306 **Datasets.** We evaluate LWD on two ultra-resolution datasets. Aesthetic-4K Zhang et al. (2025) is a
307 curated 4K benchmark with GPT-4o-generated captions and high visual quality. LAION-High-Res is
308 a filtered subset of LAION-5B Schuhmann et al. (2022), from which we sample 50K 2K-resolution
309 and 20K 4K-resolution image-caption pairs. These two datasets differ in both visual and linguistic
310 distributions, allowing us to assess both generation fidelity and generalization under caption variance.311 **Implementation Details.** We implement LWD using PyTorch and `pytorch-wavelets` (Cotter,
312 2019) for the Haar-based DWT. For the masking strategy, we set the lower bound $\ell = 0.3$, selected
313 via ablation to ensure each spatial location receives at least 30% of supervision.314 **Evaluation Protocol.** To evaluate LWD as a holistic framework, all 'LWD + Model' variants utilize
315 the Scale-Consistent VAE fine-tuning described in Section 3.1, while baseline models are evaluated
316 using their original, off-the-shelf VAE checkpoints. Notably, we observed that LWD significantly
317 accelerates convergence; models required only 10–50% of the training iterations suggested in
318 their original papers to reach convergence. Detailed hyperparameters, training costs, and other
319 configurations are provided in Appendix D.320 **Evaluation Metrics.** We evaluate ultra-resolution text-to-image generation across three key dimensions.
321 For image quality, we use Fréchet Inception Distance (FID) and LPIPS (lower is better),
322 alongside the Gray Level Co-occurrence Matrix (GLCM) Score for texture and JPEG Compression



Figure 4: Visual comparison of 2K image generations. LWD demonstrates improved detail preservation in complex areas while avoiding over-sharpening or texture collapse.

Table 1: Quantitative results on different metrics. The prompts are from the HPD (Wu et al., 2023) dataset. All images are at a resolution of 2048×2048 .

Model	FID \downarrow	LPIPS \downarrow	MAN-IQA \uparrow	QualiCLIP \uparrow	HPSv2.1 \uparrow	PickScore \uparrow
SDEdit Meng et al. (2022)	35.59	0.6456	0.3736	0.4480	30.92	22.86
I-Max Du et al. (2024b)	36.28	0.6750	0.3641	0.4139	30.62	23.02
Diffusion-4K Zhang et al. (2025)	37.10	0.6920	0.3550	0.4815	30.55	22.80
PixArt-Sigma-XL Chen et al. (2024)	36.58	0.6801	0.2949	0.4438	30.66	22.92
Sana-1.6B Xie et al. (2025)	35.75	0.7169	0.3666	0.5796	30.42	22.83
Lumina-Image 2.0 Qin et al. (2025)	54.96	0.6445	0.3663	0.4567	23.08	21.15
FLUX-1.dev Labs (2024)	37.58	<u>0.6371</u>	0.4110	0.5468	28.73	22.68
URAE Yu et al. (2025)	<u>35.25</u>	0.6717	0.4076	<u>0.5423</u>	31.15	22.41
LWD + URAE	32.88	0.6336	0.4099	0.5356	28.78	22.43

Ratio as a proxy for fine-grained detail. For semantic alignment, we report CLIPScore (Hessel et al., 2021) and QualiCLIP Agnolucci et al. (2024). Finally, for perceptual quality, we use MAN-IQA Yang et al. (2022), HPSv2.1 Wu et al. (2023), and PickScore Kirstain et al. (2023). Higher values indicate better performance for all metrics except FID and LPIPS.

4.2 QUANTITATIVE RESULTS

2K Results. Table 1 and the top block of Table 2 demonstrate consistent improvements from integrating LWD across multiple backbone models. On the HPD prompt dataset (Wu et al., 2023), LWD reduces FID by up to 7% and LPIPS up to 6% while also achieving comparable MAN-IQA and QualiCLIP, indicating improved semantic alignment and perceptual quality. Moreover, on the Aesthetic dataset (Zhang et al., 2025), these gains are observed across diverse architectures, reinforcing the generality and model-agnostic nature of our approach.

4K Results. On the Aesthetic-4K (the bottom block of Table 2) and HPD prompt dataset (Table 7), LWD outperforms baselines such as URAE and PixArt-Sigma, particularly in metrics like FID, CLIPScore and Aesthetics. The improvements are especially pronounced in regions with fine structural detail, such as hair, foliage, or architectural elements, highlighting LWD’s ability to scale effectively to ultra-high resolutions. These results suggest that frequency-aware supervision provides

378
 379 Table 2: Quantitative benchmarks of latent diffusion models on Aesthetic-Eval at 2048×2048 and
 380 4096×4096 . GLCM Score measures high-frequency texture richness using gray-level co-occurrence
 381 matrices, while Compression Ratio assesses visual complexity via JPEG file size heuristics, both
 382 introduced in Diffusion-4K Zhang et al. (2025).

	Model	FID \downarrow	CLIPScore \uparrow	Aesthetics \uparrow	GLCM Score \uparrow	Compression Ratio \downarrow
2K	SD3-F16 Esser et al. (2024)	43.82	31.50	5.91	0.75	11.23
	SD3-Diff4k-F16 Zhang et al. (2025)	40.18	34.04	5.96	0.79	11.73
	LWD + SD3-F16	38.74	34.94	6.17	0.74	11.99
	PixArt-Sigma-XL Chen et al. (2024)	39.13	35.02	6.43	0.79	13.66
4K	LWD + PixArt-Sigma-XL	36.14	35.21	6.27	0.87	6.05
	Sana-1.6B Xie et al. (2025)	32.06	35.28	6.15	0.93	24.01
	LWD + Sana-1.6B	34.30	35.58	6.23	0.78	27.34
	SD3-F16 Esser et al. (2024)	-	33.12	5.97	0.73	11.97
8K	SD3-Diff4k-F16 Zhang et al. (2025)	-	33.41	5.97	0.70	11.90
	LWD + SD3-F16	-	34.08	6.03	0.77	12.27
	Sana-1.6B Xie et al. (2025)	-	34.40	6.14	0.39	48.36
	LWD + Sana-1.6B	-	34.59	6.21	0.60	32.62

394
 395 meaningful guidance even in challenging high-frequency regimes where baseline methods often
 396 struggle. LWD also achieves the highest GLCM score when paired with SD3-F16, and substantially
 397 improves Sana’s performance on both GLCM and compression ratio, indicating stronger fine-detail
 398 retention and texture fidelity, without compromising overall quality.

401 4.3 QUALITATIVE RESULTS

402 Figures 4, 5, and 9 compare outputs from LWD and baseline models under identical prompts. LWD
 403 consistently renders sharper textures in high-frequency regions, such as fabric, skin, and hair, while
 404 preserving global structure.

405 Zoomed-in comparisons highlight improved reconstruction of fine details (e.g., hair strands, eyelashes,
 406 small objects) that are often blurred or omitted by baselines. These results suggest that frequency-
 407 aware masking not only enhances local precision but does so in a context-sensitive manner, avoiding
 408 over-sharpening or artifacts. This indicates that LWD effectively balances fine detail refinement with
 409 structural coherence. More full-resolution results can be found in Appendix B.

412 4.4 ABLATION STUDIES

414 4.4.1 EFFECT OF SCALE-CONSISTENCY LOSS ON RECONSTRUCTION QUALITY

416 Table 3 reports quantitative reconstruction metrics for various VAEs, with and without the proposed
 417 Scale-Consistency (SC) loss, evaluated on the Aesthetic-4K validation set. Across different architec-
 418 tures SD3-VAE, Flux-VAE, and Sana-AE, the addition of SC consistently improves performance,
 419 particularly in rFID and perceptual quality (LPIPS). For instance, Flux-VAE-SC outperforms its
 420 baseline with a significant reduction in rFID (0.50 vs. 0.73) and an increase in PSNR and SSIM,
 421 indicating sharper and more faithful reconstructions. Notably, SD3-VAE-F16-SC achieves a sub-
 422 stantial LPIPS improvement (0.18 vs. 0.30), suggesting better perceptual fidelity despite using a
 423 more aggressive compression factor (F16). These results confirm that scale-consistent regularization
 424 enhances latent representations, making them more robust and structurally aligned, critical properties
 425 for downstream diffusion tasks.

426 4.4.2 IMPACT OF VAE REGULARIZATION AND WAVELET MASKING

428 To rigorously evaluate the contribution of each component within the LWD framework, we conducted
 429 a detailed ablation study. Our method comprises two main stages: (1) fine-tuning a VAE with a scale-
 430 consistency (SC) loss and (2) fine-tuning the diffusion model with our proposed wavelet-masked loss
 431 ($\mathcal{L}_{\text{masked}}$). The study presented here isolates the impact of each component on the final text-to-image
 432 generation quality.



Figure 5: 4K images generated by LWD with different architectures.

Table 3: Quantitative reconstruction results of VAEs with and without the Scale-Consistency Loss 3.1 on Aesthetic-4K (Zhang et al., 2025) validation set.

Model	rFID ↓	NMSE ↓	PSNR ↑	SSIM ↑	LPIPS ↓
SD3-VAE (Esser et al., 2024)	1.05	0.01	26.54	0.86	0.08
SD3-VAE-F16 (Zhang et al., 2025)	0.70	0.07	19.82	0.63	0.30
SD3-VAE-F16-SC	0.70	0.04	22.58	0.75	0.18
Flux-VAE (Labs, 2024)	0.73	0.01	27.18	0.89	0.07
Flux-VAE-SC	0.50	0.01	28.14	0.90	0.06
Sana-AE (Xie et al., 2025)	0.74	0.04	22.16	0.70	0.159
Sana-AE-SC	0.55	0.02	23.64	0.73	0.163

The results, shown in Table 4, confirm that both the SC loss and the wavelet-masked loss contribute meaningfully to generation quality, with the full LWD framework yielding the strongest performance. The results in Table 4 show that the GLCM score slightly decreases with the full LWD framework. This reflects a known limitation of classical texture metrics like GLCM, which do not always correlate with perceptual coherence. Our method trades raw statistical complexity for more realistic details, a positive trade-off validated by significant improvements in perceptual metrics like FID and Aesthetics.

486
 487 Table 4: Ablation Study on the Contributions of LWD Components. We evaluate each component’s
 488 impact on final generation quality using the Diffusion4k backbone on the Aesthetic dataset at
 2048×2048 resolution.

489 Configuration	490 FID ↓	491 CLIPScore ↑	492 Aesthetics ↑	493 GLCM Score ↑
490 Baseline (SD3-Diff4k-F16)	491 40.18	492 34.04	493 5.96	494 0.79
490 + VAE Scale-Consistency	491 39.50	492 34.10	493 6.05	494 0.78
490 + Wavelet Masking	491 39.20	492 34.50	493 6.10	494 0.75
490 Full LWD (Ours)	491 38.74	492 34.94	493 6.17	494 0.74

495 4.5 DISCUSSION

496 Across quantitative and qualitative benchmarks, LWD enhances ultra-resolution image synthesis by
 497 integrating signal-derived saliency into the training loss. Compared to both conventional models and
 498 prior wavelet-based methods (e.g., Diffusion-4K), it improves perceptual fidelity, semantic alignment,
 499 and spectral regularity, without increasing inference cost or modifying the underlying architecture.
 500 Its plug-and-play nature makes it broadly compatible with modern latent diffusion pipelines.

501 LWD improves perceptual fidelity while maintaining comparable alignment scores. This reflects an
 502 intentional design choice: LWD prioritizes high-frequency detail recovery to prevent texture collapse
 503 at UHR scales, complementing the base model’s semantic capabilities rather than replacing them.

504 Beyond quality gains, LWD represents a shift toward more interpretable and structure-aware supervi-
 505 sion. Unlike attention mechanisms that rely on semantic priors, LWD leverages wavelet energy as a
 506 transparent, self-supervised signal to prioritize detail-rich regions.

507 This frequency-guided supervision introduces a form of spatial curriculum learning, where complex
 508 regions receive more focused updates. Such adaptive loss weighting opens avenues for dynamic
 509 training strategies, such as frequency-aware learning rates or hybrid spatial-frequency schedules.
 510 These mechanisms may be especially valuable in domains where structural detail is critical but
 511 semantic guidance is weak, such as scientific visualization, material design, or multimodal generation.

512 **Limitations and Future Work.** While LWD improves generation quality without architectural
 513 changes or inference overhead, it inherits limitations common to latent diffusion models. In particular,
 514 VAE compression can cause the loss of fine-grained semantic detail, potentially limiting performance
 515 in tasks requiring precise spatial alignment or photorealistic accuracy.

516 Future work could address this by incorporating higher-fidelity latent spaces or hybrid approaches that
 517 combine latent and pixel-space supervision. Extending LWD to domains such as video generation,
 518 depth-aware synthesis, or multimodal conditioning also represents a promising direction. The
 519 frequency-aware masking mechanism is general and could be adapted to guide temporal attention,
 520 cross-modal alignment, or resolution-specific sampling in broader generative contexts.

521 5 CONCLUSION

522 We introduced *Latent Wavelet Diffusion* (LWD), a general and modular framework that integrates
 523 frequency-based supervision into latent diffusion models for ultra-high-resolution image synthesis.
 524 By computing wavelet energy maps in the latent space and applying spatially and temporally adaptive
 525 masking, LWD selectively emphasizes high-detail regions during training. Without requiring archi-
 526 tectural modifications or incurring additional inference cost, LWD consistently improves perceptual
 527 fidelity and semantic alignment across models such as Flux and SD3. It preserves high-frequency
 528 detail and structural coherence more effectively than prior methods, demonstrating the value of
 529 signal-aware supervision in guiding the generative process. By unifying principles from signal
 530 processing and diffusion modeling, LWD offers a scalable and interpretable approach applicable to a
 531 wide range of generative architectures.

532 **Broader Impact.** LWD promotes efficient and interpretable generation by aligning supervision
 533 with signal-level detail. This may benefit applications requiring controllable high-resolution synthe-
 534 sis, while raising familiar concerns around synthetic media misuse. Incorporating safeguards and
 535 provenance tools remains an important direction.

540

6 REPRODUCIBILITY STATEMENT

541
 542 To ensure the reproducibility of our work, we have included the core Python script detailing our
 543 wavelet-based masking algorithm in the supplementary material. Further implementation details,
 544 including key hyperparameters and computational requirements, are provided in Appendix section ??.
 545 Following the review process, we will release our public GitHub repository, which will contain
 546 the complete implementation, training scripts, evaluation code, and the final pre-trained model
 547 checkpoints.

548

549 7 ETHICAL STATEMENT

550
 551 Our work is built upon publicly available datasets commonly used in the field of generative modeling.
 552 We acknowledge that these large-scale datasets may contain inherent societal biases, which our model
 553 could potentially learn and reproduce. We have used these datasets in accordance with their original
 554 licenses. Below, we discuss the potential societal impacts of our work.

555

556 7.1 POSITIVE IMPACT

557
 558 The primary goal of our research is to advance the state of high-resolution image generation, which
 559 has significant positive applications. These include empowering artists, designers, and content
 560 creators with more powerful creative tools; enhancing visual effects for entertainment and media;
 561 and potentially aiding in scientific visualization and data augmentation. By developing a method that
 562 improves quality with **zero inference overhead**, we aim to make high-fidelity generative AI more
 563 accessible and practical for a wider range of beneficial uses.

564

565 7.2 POTENTIAL RISKS AND MITIGATION

566
 567 We recognize that generative models can be misused for malicious purposes, such as creating
 568 misinformation (“deepfakes”), generating harmful or explicit content, and perpetuating societal
 569 biases. To mitigate these risks, we are committed to the following measures:

- 570 **1. Responsible Release:** We plan to release our code and models under a responsible AI
 571 license (e.g., a variant of the CreativeML Open RAIL-M license) that explicitly prohibits
 572 use for malicious, harmful, or unethical purposes.
- 573 **2. Acknowledging Limitations:** We are transparent about the limitations of our model and its
 574 potential to generate biased or factually incorrect content, as discussed in the main paper.
- 575 **3. Encouraging Safe Deployment:** We strongly encourage all downstream users of our models
 576 to implement their own safety filters, content moderation systems, and ethical guidelines
 577 before deploying any applications.

578

580 8 LLM USAGE STATEMENT

581
 582 A large language model (LLM) was used as a writing assistance tool during the preparation of this
 583 manuscript. The LLM’s role was limited to improving grammar, clarity, and conciseness. All content
 584 was conceived and written by the authors, who take full responsibility for the paper’s scientific
 585 integrity.

586

587 REFERENCES

588
 589 Niket Agarwal, Arslan Ali, Maciej Bala, Yogesh Balaji, Erik Barker, Tiffany Cai, Prithvijit Chat-
 590 topadhyay, Yongxin Chen, Yin Cui, Yifan Ding, et al. Cosmos world foundation model platform
 591 for physical ai. *arXiv preprint arXiv:2501.03575*, 2025.

592
 593 Lorenzo Agnolucci, Leonardo Galteri, and Marco Bertini. Quality-aware image-text alignment for
 594 opinion-unaware image quality assessment. *arXiv preprint arXiv:2403.11176*, 2024.

594 Lorenzo Aloisi, Luigi Sigillo, Aurelio Uncini, and Danilo Comminiello. A wavelet diffusion gan for
 595 image super-resolution. *arXiv preprint arXiv:2410.17966*, 2024.

596

597 Omer Bar-Tal, Lior Yariv, Yaron Lipman, and Tali Dekel. MultiDiffusion: Fusing diffusion paths
 598 for controlled image generation. In Andreas Krause, Emma Brunskill, Kyunghyun Cho, Barbara
 599 Engelhardt, Sivan Sabato, and Jonathan Scarlett (eds.), *Proceedings of the 40th International
 600 Conference on Machine Learning*, volume 202 of *Proceedings of Machine Learning Research*, pp.
 601 1737–1752. PMLR, 23–29 Jul 2023. URL [https://proceedings.mlr.press/v202/
 602 bar-tal23a.html](https://proceedings.mlr.press/v202/bar-tal23a.html).

603 Sara Björk, Jonas Nordhaug Myhre, and Thomas Haugland Johansen. Simpler is better: Spectral
 604 regularization and up-sampling techniques for variational autoencoders. In *ICASSP 2022 - 2022
 605 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, pp. 3778–
 606 3782, 2022. doi: 10.1109/ICASSP43922.2022.9746027.

607 Kevin Black, Michael Janner, Yilun Du, Ilya Kostrikov, and Sergey Levine. Training diffusion models
 608 with reinforcement learning. In *The Twelfth International Conference on Learning Representations*,
 609 2024. URL <https://openreview.net/forum?id=YCWjhGrJFD>.

610

611 Tim Brooks, Aleksander Holynski, and Alexei A Efros. Instructpix2pix: Learning to follow image
 612 editing instructions. In *Proceedings of the IEEE/CVF conference on computer vision and pattern
 613 recognition*, pp. 18392–18402, 2023.

614 Mathilde Caron, Hugo Touvron, Ishan Misra, Hervé Jégou, Julien Mairal, Piotr Bojanowski, and
 615 Armand Joulin. Emerging properties in self-supervised vision transformers. In *Proceedings of the
 616 IEEE/CVF international conference on computer vision*, pp. 9650–9660, 2021.

617

618 Junsong Chen, Chongjian Ge, Enze Xie, Yue Wu, Lewei Yao, Xiaozhe Ren, Zhongdao Wang, Ping
 619 Luo, Huchuan Lu, and Zhenguo Li. Pixart- σ : Weak-to-strong training of diffusion transformer for
 620 4k text-to-image generation. In *European Conference on Computer Vision*, pp. 74–91. Springer,
 621 2024.

622 Fergal Cotter. *Uses of Complex Wavelets in Deep Convolutional Neural Networks*. PhD thesis, Apollo
 623 - University of Cambridge Repository, 2019. URL <https://www.repository.cam.ac.uk/handle/1810/306661>.

624

625 Ruoyi Du, Dongliang Chang, Timothy Hospedales, Yi-Zhe Song, and Zhanyu Ma. Demofusion:
 626 Democratizing high-resolution image generation with no \$\$. In *Proceedings of the IEEE/CVF
 627 Conference on Computer Vision and Pattern Recognition*, pp. 6159–6168, 2024a.

628

629 Ruoyi Du, Dongyang Liu, Le Zhuo, Qin Qi, Hongsheng Li, Zhanyu Ma, and Peng Gao. I-max:
 630 Maximize the resolution potential of pre-trained rectified flow transformers with projected flow,
 631 2024b. URL <https://arxiv.org/abs/2410.07536>.

632

633 Patrick Esser, Robin Rombach, and Bjorn Ommer. Taming transformers for high-resolution image
 634 synthesis. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*,
 635 pp. 12873–12883, 2021.

636

637 Patrick Esser, Sumith Kulal, Andreas Blattmann, Rahim Entezari, Jonas Müller, Harry Saini, Yam
 638 Levi, Dominik Lorenz, Axel Sauer, Frederic Boesel, et al. Scaling rectified flow transformers for
 639 high-resolution image synthesis. In *Forty-first international conference on machine learning*, 2024.

640

641 Carlos Esteves, Mohammed Suhail, and Ameesh Makadia. Spectral image tokenizer. *arXiv preprint
 642 arXiv:2412.09607*, 2024.

643

644 Jiatao Gu, Shuangfei Zhai, Yizhe Zhang, Joshua M. Susskind, and Navdeep Jaitly. Matryoshka
 645 diffusion models. In *The Twelfth International Conference on Learning Representations*, 2024.
 646 URL <https://openreview.net/forum?id=tOzCcDdH9O>.

647

648 Lanqing Guo, Yingqing He, Haoxin Chen, Menghan Xia, Xiaodong Cun, Yufei Wang, Siyu Huang,
 649 Yong Zhang, Xintao Wang, Qifeng Chen, et al. Make a cheap scaling: A self-cascade diffusion
 650 model for higher-resolution adaptation. In *European Conference on Computer Vision*, pp. 39–55.
 651 Springer, 2024.

648 Yoav HaCohen, Nisan Chiprut, Benny Brazowski, Daniel Shalem, Dudu Moshe, Eitan Richardson,
 649 Eran Levin, Guy Shiran, Nir Zabari, Ori Gordon, et al. Ltx-video: Realtime video latent diffusion.
 650 *arXiv preprint arXiv:2501.00103*, 2024.

651

652 Jaehoon Hahm, Junho Lee, Sunghyun Kim, and Joonseok Lee. Isometric representation learning for
 653 disentangled latent space of diffusion models. In *Forty-first International Conference on Machine
 654 Learning*, 2024. URL <https://openreview.net/forum?id=ufCptn28vG>.

655 Yingqing He, Shaoshu Yang, Haoxin Chen, Xiaodong Cun, Menghan Xia, Yong Zhang, Xintao Wang,
 656 Ran He, Qifeng Chen, and Ying Shan. Scalecrafter: Tuning-free higher-resolution visual generation
 657 with diffusion models. In *The Twelfth International Conference on Learning Representations*,
 658 2023.

659

660 Jack Hessel, Ari Holtzman, Maxwell Forbes, Ronan Le Bras, and Yejin Choi. Clipscore: A reference-
 661 free evaluation metric for image captioning. In *EMNLP (1)*, 2021.

662

663 Jonathan Ho, Ajay Jain, and Pieter Abbeel. Denoising diffusion probabilistic models. *Advances in
 664 neural information processing systems*, 33:6840–6851, 2020.

665

666 Jonathan Ho, Chitwan Saharia, William Chan, David J Fleet, Mohammad Norouzi, and Tim Salimans.
 667 Cascaded diffusion models for high fidelity image generation. *Journal of Machine Learning
 668 Research*, 23(47):1–33, 2022.

669

670 Linjiang Huang, Rongyao Fang, Aiping Zhang, Guanglu Song, Si Liu, Yu Liu, and Hongsheng Li.
 671 Fouriscale: A frequency perspective on training-free high-resolution image synthesis. In *European
 672 Conference on Computer Vision*, pp. 196–212. Springer, 2024a.

673

674 Yi Huang, Jiancheng Huang, Jianzhuang Liu, Mingfu Yan, Yu Dong, Jiaxi Lv, Chaoqi Chen, and
 675 Shifeng Chen. Wavedm: Wavelet-based diffusion models for image restoration. *IEEE Transactions
 676 on Multimedia*, 26:7058–7073, 2024b. doi: 10.1109/TMM.2024.3359769.

677

678 Jinho Jeong, Sangmin Han, Jinwoo Kim, and Seon Joo Kim. Latent space super-resolution for
 679 higher-resolution image generation with diffusion models. *arXiv preprint arXiv:2503.18446*, 2025.

680

681 Hai Jiang, Ao Luo, Haoqiang Fan, Songchen Han, and Shuaicheng Liu. Low-light image enhancement
 682 with wavelet-based diffusion models. *ACM Trans. Graph.*, 42(6), December 2023. ISSN 0730-0301.
 683 doi: 10.1145/3618373. URL <https://doi.org/10.1145/3618373>.

684

685 Younghyun Kim, Geunmin Hwang, Junyu Zhang, and Eunbyung Park. Diffusehigh: Training-free
 686 progressive high-resolution image synthesis through structure guidance. In *Proceedings of the
 687 AAAI conference on artificial intelligence*, volume 39, pp. 4338–4346, 2025.

688

689 Yuval Kirstain, Adam Polyak, Uriel Singer, Shahbuland Matiana, Joe Penna, and Omer Levy. Pick-
 690 a-pic: An open dataset of user preferences for text-to-image generation. *Advances in Neural
 691 Information Processing Systems*, 36:36652–36663, 2023.

692

693 Theodoros Kouzelis, Ioannis Kakogeorgiou, Spyros Gidaris, and Nikos Komodakis. Eq-vae:
 694 Equivariance regularized latent space for improved generative image modeling. *arXiv preprint
 695 arXiv:2502.09509*, 2025.

696

697 Black Forest Labs. Flux. <https://github.com/black-forest-labs/flux>, 2024.

698

699 Yuseung Lee, Kunho Kim, Hyunjin Kim, and Minhyuk Sung. Syncdiffusion: Coherent montage via
 700 synchronized joint diffusions. *Advances in Neural Information Processing Systems*, 36:50648–
 701 50660, 2023.

702

703 Bin Lin, Yunyang Ge, Xinhua Cheng, Zongjian Li, Bin Zhu, Shaodong Wang, Xianyi He, Yang Ye,
 704 Shenghai Yuan, Liuhan Chen, et al. Open-sora plan: Open-source large video generation model.
 705 *arXiv preprint arXiv:2412.00131*, 2024.

706

707 Yaron Lipman, Ricky T. Q. Chen, Heli Ben-Hamu, Maximilian Nickel, and Matthew Le. Flow
 708 matching for generative modeling. In *The Eleventh International Conference on Learning Re-
 709 presentations*, 2023. URL <https://openreview.net/forum?id=PqvMRDCJT9t>.

702 Guangyi Liu, Yu Wang, Zeyu Feng, Qiyu Wu, Liping Tang, Yuan Gao, Zhen Li, Shuguang Cui,
 703 Julian McAuley, Zichao Yang, Eric P. Xing, and Zhititing Hu. Unified generation, reconstruction,
 704 and representation: Generalized diffusion with adaptive latent encoding-decoding. In *Forty-first*
 705 *International Conference on Machine Learning*, 2024a. URL <https://openreview.net/forum?id=igRjCCAz2a>.

706

707 Yixin Liu, Kai Zhang, Yuan Li, Zhiling Yan, Chujie Gao, Ruoxi Chen, Zhengqing Yuan, Yue Huang,
 708 Hanchi Sun, Jianfeng Gao, Lifang He, and Lichao Sun. Sora: A review on background, technology,
 709 limitations, and opportunities of large vision models, 2024b.

710

711 Eleonora Lopez, Luigi Sigillo, Federica Colonnese, Massimo Panella, and Danilo Comminiello.
 712 Guess what i think: Streamlined eeg-to-image generation with latent diffusion models. In *ICASSP*
 713 *2025 - 2025 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*,
 714 pp. 1–5, 2025. doi: 10.1109/ICASSP49660.2025.10890059.

715 Stephane Mallat and Wen Liang Hwang. Singularity detection and processing with wavelets. *IEEE*
 716 *transactions on information theory*, 38(2):617–643, 1992.

717

718 Chenlin Meng, Yutong He, Yang Song, Jiaming Song, Jiajun Wu, Jun-Yan Zhu, and Stefano Ermon.
 719 SDEdit: Guided image synthesis and editing with stochastic differential equations. In *International*
 720 *Conference on Learning Representations*, 2022. URL https://openreview.net/forum?id=aBsCjcpU_tE.

721

722 Brian B. Moser, Stanislav Frolov, Federico Raue, Sebastian Palacio, and Andreas Dengel. Dynamic
 723 Attention-Guided Diffusion for Image Super-Resolution . In *2025 IEEE/CVF Winter Conference*
 724 *on Applications of Computer Vision (WACV)*, pp. 451–460, Los Alamitos, CA, USA, March
 725 2025. IEEE Computer Society. doi: 10.1109/WACV61041.2025.00054. URL <https://doi.ieee.org/10.1109/WACV61041.2025.00054>.

726

727 William Peebles and Saining Xie. Scalable diffusion models with transformers. In *Proceedings of*
 728 *the IEEE/CVF international conference on computer vision*, pp. 4195–4205, 2023.

729

730 Hao Phung, Quan Dao, and Anh Tran. Wavelet diffusion models are fast and scalable image generators.
 731 In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pp. 10199–
 732 10208, 2023.

733

734 Dustin Podell, Zion English, Kyle Lacey, Andreas Blattmann, Tim Dockhorn, Jonas Müller, Joe
 735 Penna, and Robin Rombach. SDXL: Improving latent diffusion models for high-resolution image
 736 synthesis. In *The Twelfth International Conference on Learning Representations*, 2024. URL
<https://openreview.net/forum?id=di52zR8xgf>.

737

738 Yurui Qian, Qi Cai, Yingwei Pan, Yehao Li, Ting Yao, Qibin Sun, and Tao Mei. Boosting diffusion
 739 models with moving average sampling in frequency domain. In *Proceedings of the IEEE/CVF*
740 Conference on Computer Vision and Pattern Recognition, pp. 8911–8920, 2024.

741

742 Qi Qin, Le Zhuo, Yi Xin, Ruoyi Du, Zhen Li, Bin Fu, Yiting Lu, Jiakang Yuan, Xinyue Li, Dongyang
 743 Liu, et al. Lumina-image 2.0: A unified and efficient image generative framework. *arXiv preprint*
744 arXiv:2503.21758, 2025.

745

746 Jingjing Ren, Wenbo Li, Haoyu Chen, Renjing Pei, Bin Shao, Yong Guo, Long Peng, Fenglong
 747 Song, and Lei Zhu. Ultrapixel: Advancing ultra high-resolution image synthesis to new peaks.
 748 In *The Thirty-eighth Annual Conference on Neural Information Processing Systems*, 2024. URL
<https://openreview.net/forum?id=voJCPdlw53>.

749

750 Robin Rombach, Andreas Blattmann, Dominik Lorenz, Patrick Esser, and Björn Ommer. High-
 751 resolution image synthesis with latent diffusion models. In *Proceedings of the IEEE/CVF confer-
 752 ence on computer vision and pattern recognition*, pp. 10684–10695, 2022.

753

754 Chitwan Saharia, Jonathan Ho, William Chan, Tim Salimans, David J Fleet, and Mohammad Norouzi.
 755 Image super-resolution via iterative refinement. *IEEE transactions on pattern analysis and machine*
intelligence, 45(4):4713–4726, 2022.

756

757 Axel Sauer, Dominik Lorenz, Andreas Blattmann, and Robin Rombach. Adversarial diffusion
 758 distillation. In *European Conference on Computer Vision*, pp. 87–103. Springer, 2024.

756 Christoph Schuhmann, Romain Beaumont, Richard Vencu, Cade Gordon, Ross Wightman, Mehdi
 757 Cherti, Theo Coombes, Aarush Katta, Clayton Mullis, Mitchell Wortsman, Patrick Schramowski,
 758 Srivatsa Kundurthy, Katherine Crowson, Ludwig Schmidt, Robert Kaczmarczyk, and Jenia Jitsev.
 759 Laion-5b: an open large-scale dataset for training next generation image-text models. In *Proceed-
 760 ings of the 36th International Conference on Neural Information Processing Systems*, NIPS '22,
 761 Red Hook, NY, USA, 2022. Curran Associates Inc. ISBN 9781713871088.

762 Hui Shen, Jingxuan Zhang, Boning Xiong, Rui Hu, Shoufa Chen, Zhongwei Wan, Xin Wang,
 763 Yu Zhang, Zixuan Gong, Guangyin Bao, Chaofan Tao, Yongfeng Huang, Ye Yuan, and Mi Zhang.
 764 Efficient diffusion models: A survey. *Transactions on Machine Learning Research*, 2025. URL
 765 <https://openreview.net/forum?id=wHECkBOwyt>.

766 Shuwei Shi, Wenbo Li, Yuechen Zhang, Jingwen He, Biao Gong, and Yinqiang Zheng. Resmas-
 767 ter: Mastering high-resolution image generation via structural and fine-grained guidance. In
 768 *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 39, pp. 6887–6895, 2025.

769 Luigi Sigillo, Riccardo Fosco Gramaccioni, Alessandro Nicolosi, and Danilo Comminiello. Ship in
 770 sight: Diffusion models for ship-image super resolution. In *2024 International Joint Conference
 771 on Neural Networks (IJCNN)*, pp. 1–8, 2024. doi: 10.1109/IJCNN60899.2024.10650042.

772 Luigi Sigillo, Christian Bianchi, Aurelio Uncini, and Danilo Comminiello. Quaternion wavelet-
 773 conditioned diffusion models for image super-resolution. In *2025 International Joint Conference on
 774 Neural Networks (IJCNN)*, pp. 1–8, 2025. URL <https://arxiv.org/abs/2505.00334>.

775 Ivan Skorokhodov, Sharath Girish, Benran Hu, Willi Menapace, Yanyu Li, Rameen Abdal, Sergey
 776 Tulyakov, and Aliaksandr Siarohin. Improving the diffusability of autoencoders, 2025. URL
 777 <https://arxiv.org/abs/2502.14831>.

778 Jiaming Song, Chenlin Meng, and Stefano Ermon. Denoising diffusion implicit models. In *Inter-
 779 national Conference on Learning Representations*, 2021a. URL <https://openreview.net/forum?id=St1giarCHLP>.

780 Yang Song, Jascha Sohl-Dickstein, Diederik P Kingma, Abhishek Kumar, Stefano Ermon, and Ben
 781 Poole. Score-based generative modeling through stochastic differential equations. In *International
 782 Conference on Learning Representations*, 2021b. URL <https://openreview.net/forum?id=PxTIG12RRHS>.

783 Yuhta Takida, Yukara Ikemiya, Takashi Shibuya, Kazuki Shimada, Woosung Choi, Chieh-Hsin Lai,
 784 Naoki Murata, Toshimitsu Uesaka, Kengo Uchida, Wei-Hsiang Liao, and Yuki Mitsufuji. HQ-VAE:
 785 Hierarchical discrete representation learning with variational bayes. *Transactions on Machine
 786 Learning Research*, 2024. ISSN 2835-8856. URL <https://openreview.net/forum?id=xqAVkqrLjx>.

787 Anni Tang, Tianyu He, Junliang Guo, Xinle Cheng, Li Song, and Jiang Bian. Vidtok: A versatile and
 788 open-source video tokenizer. *arXiv preprint arXiv:2412.13061*, 2024.

789 Arash Vahdat and Jan Kautz. Nvae: A deep hierarchical variational autoencoder. *Advances in neural
 790 information processing systems*, 33:19667–19679, 2020.

791 Yufei Wang, Wenhan Yang, Xinyuan Chen, Yaohui Wang, Lanqing Guo, Lap-Pui Chau, Ziwei Liu,
 792 Yu Qiao, Alex C Kot, and Bihan Wen. Sinsr: diffusion-based image super-resolution in a single
 793 step. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*,
 794 pp. 25796–25805, 2024.

795 Z. Wang, E.P. Simoncelli, and A.C. Bovik. Multiscale structural similarity for image quality
 796 assessment. In *The Thirly-Seventh Asilomar Conference on Signals, Systems & Computers,
 797 2003*, volume 2, pp. 1398–1402 Vol.2, 2003. doi: 10.1109/ACSSC.2003.1292216.

798 Zhendong Wang, Jianmin Bao, Shuyang Gu, Dong Chen, Wengang Zhou, and Houqiang Li. De-
 799 signdiffusion: High-quality text-to-design image generation with diffusion models, 2025. URL
 800 <https://arxiv.org/abs/2503.01645>.

810 Xiaoshi Wu, Yiming Hao, Keqiang Sun, Yixiong Chen, Feng Zhu, Rui Zhao, and Hongsheng Li.
 811 Human preference score v2: A solid benchmark for evaluating human preferences of text-to-image
 812 synthesis, 2023. URL <https://arxiv.org/abs/2306.09341>.

813

814 Enze Xie, Lewei Yao, Han Shi, Zhili Liu, Daquan Zhou, Zhaoqiang Liu, Jiawei Li, and Zhenguo
 815 Li. Difffit: Unlocking transferability of large diffusion models via simple parameter-efficient
 816 fine-tuning. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pp.
 817 4230–4239, 2023.

818 Enze Xie, Junsong Chen, Junyu Chen, Han Cai, Haotian Tang, Yujun Lin, Zhekai Zhang, Muyang
 819 Li, Ligeng Zhu, Yao Lu, and Song Han. SANA: Efficient high-resolution text-to-image synthesis
 820 with linear diffusion transformers. In *The Thirteenth International Conference on Learning
 821 Representations*, 2025. URL <https://openreview.net/forum?id=N8Oj1XhtYZ>.

822 Sidi Yang, Tianhe Wu, Shuwei Shi, Shanshan Lao, Yuan Gong, Mingdeng Cao, Jiahao Wang, and
 823 Yujiu Yang. Maniqa: Multi-dimension attention network for no-reference image quality assessment.
 824 In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pp. 1191–
 825 1200, 2022.

826 Xingyi Yang, Daquan Zhou, Jiashi Feng, and Xinchao Wang. Diffusion probabilistic model made
 827 slim. In *2023 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pp.
 828 22552–22562, 2023. doi: 10.1109/CVPR52729.2023.02160.

829

830 Ruonan Yu, Songhua Liu, Zhenxiong Tan, and Xinchao Wang. Ultra-resolution adaptation with ease.
 831 *International Conference on Machine Learning*, 2025.

832

833 Zheyuan Zhan, Defang Chen, Jian-Ping Mei, Zhenghe Zhao, Jiawei Chen, Chun Chen, Siwei Lyu, and
 834 Can Wang. Conditional image synthesis with diffusion models: A survey. *CoRR*, abs/2409.19365,
 835 2024. URL <https://doi.org/10.48550/arXiv.2409.19365>.

836 Jinjin Zhang, Qiuyu Huang, Junjie Liu, Xiefan Guo, and Di Huang. Diffusion-4k: Ultra-high-
 837 resolution image synthesis with latent diffusion models. In *IEEE/CVF Conference on Computer
 838 Vision and Pattern Recognition (CVPR)*, 2025.

839 Lin Zhang, Lei Zhang, Xuanqin Mou, and David Zhang. Fsim: A feature similarity index for
 840 image quality assessment. *IEEE Transactions on Image Processing*, 20(8):2378–2386, 2011. doi:
 841 10.1109/TIP.2011.2109730.

842

843 Lvmin Zhang, Anyi Rao, and Maneesh Agrawala. Adding conditional control to text-to-image
 844 diffusion models. In *Proceedings of the IEEE/CVF international conference on computer vision*,
 845 pp. 3836–3847, 2023.

846 Chen Zhao, Weiling Cai, Chenyu Dong, and Chengwei Hu. Wavelet-based fourier information
 847 interaction with frequency diffusion adjustment for underwater image restoration. In *Proceedings
 848 of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 8281–8291, 2024.

849

850 Qingping Zheng, Yuanfan Guo, Jiankang Deng, Jianhua Han, Ying Li, Songcen Xu, and Hang Xu.
 851 Any-size-diffusion: Toward efficient text-driven synthesis for any-size hd images. In *Proceedings
 852 of the AAAI Conference on Artificial Intelligence*, volume 38, pp. 7571–7578, 2024.

853

854

855

856

857

858

859

860

861

862

863

864 **A THEORETICAL AND IMPLEMENTATION DETAILS OF WAVELET-GUIDED
865 MASKING**
866

867 This section provides a detailed theoretical background on the Discrete Wavelet Transform (DWT)
868 and its application in our framework, discusses specific implementation choices and hyperparameters,
869 and analyzes the properties and trade-offs of our proposed method.
870

871 **A.1 THEORETICAL FOUNDATION: DWT AND RELEVANCE MAPS**
872

873 **Discrete Wavelet Transform (DWT).** The Discrete Wavelet Transform (DWT) decomposes a
874 2D signal, such as an image or a latent tensor, into orthogonal frequency subbands by applying
875 separable cascades of low-pass and high-pass filters along both spatial dimensions. Given an input
876 $\mathbf{X} \in \mathbb{R}^{H \times W \times D}$, a single-level DWT produces four spatial subbands:
877

878

879 - **LL** (Low-Low): Approximation coefficients capturing global structure and coarse semantics.
880 - **LH** (Low-High): Horizontal detail coefficients, sensitive to vertical edges.
881 - **HL** (High-Low): Vertical detail coefficients, highlighting horizontal edges.
882 - **HH** (High-High): Diagonal detail coefficients, encoding fine textures and high-frequency
883 transitions.

884 Formally, the DWT of \mathbf{X} is given by:
885

886
$$\text{DWT}(\mathbf{X}) = \begin{cases} \text{LL} = \mathbf{X} * h_{\text{low}} \downarrow 2 * h_{\text{low}} \downarrow 2 \\ \text{LH} = \mathbf{X} * h_{\text{low}} \downarrow 2 * h_{\text{high}} \downarrow 2 \\ \text{HL} = \mathbf{X} * h_{\text{high}} \downarrow 2 * h_{\text{low}} \downarrow 2 \\ \text{HH} = \mathbf{X} * h_{\text{high}} \downarrow 2 * h_{\text{high}} \downarrow 2 \end{cases}$$
887

888 where h_{low} , h_{high} are orthogonal wavelet filters (e.g., Haar or Daubechies), $*$ denotes convolution,
889 and $\downarrow 2$ indicates downsampling by a factor of 2.
890

891 **Relevance Map Construction.** To identify spatial regions that require enhanced refinement during
892 generation, we compute a wavelet-based relevance map from the high-frequency subbands.
893 Specifically, we calculate the aggregated energy of the directional detail coefficients:
894

895
$$\mathbf{M}_{\text{relevance}} = \text{LH}^2 + \text{HL}^2 + \text{HH}^2$$
896

897 This yields a saliency map that highlights localized frequency-rich regions, such as edges, textures,
898 and fine details. The relevance map is then resized to match the original latent resolution via bilinear
899 interpolation and normalized to the range $[0, 1]$:
900

901
$$\mathbf{M}_{\text{norm}} = \frac{\mathbf{M}_{\text{relevance}} - \min(\mathbf{M}_{\text{relevance}})}{\max(\mathbf{M}_{\text{relevance}}) - \min(\mathbf{M}_{\text{relevance}}) + \epsilon}$$
902

903 where ϵ is a small constant to avoid division by zero. The LL subband, which encodes coarse
904 spatial content, offers limited information about local complexity and is thus excluded. In contrast,
905 the aggregated energy of the LH, HL, and HH bands approximates the local gradient magnitude,
906 similar in spirit to edge detectors like Sobel and Laplacian filters, and aligns with the intuition that
907 visually salient regions often correspond to areas with rich high-frequency content. This approach
908 is theoretically supported by prior work in multiscale signal analysis, which demonstrates that the
909 local maxima of wavelet detail coefficients correspond to structural singularities and perceptual
910 boundaries (Mallat & Hwang, 1992).
911

912 **A.2 IMPLEMENTATION AND HYPERPARAMETERS**
913

914 **Choice of Wavelet Basis.** We deliberately selected the Haar wavelet for its ideal trade-off of
915 properties for our framework:
916

917

- **Computational Efficiency:** It is the most computationally efficient wavelet, minimizing
918 training overhead.

- **Preservation of Discontinuities:** Its discontinuous nature makes it exceptionally effective at localizing and preserving sharp edges and contours.
- **Sparsity and Non-Redundancy:** Its orthogonality and compact support induce sparse, non-redundant representations, making our energy maps precise and ideal for our masking strategy.

Comparative Analysis with Other Transforms. To empirically validate the selection of Haar wavelets over other frequency analysis methods, we conducted a rigorous ablation study comparing our approach against Daubechies wavelets (db2) and FFT-based High-Pass filtering. The results are summarized in Table 5.

Two key principles dictate the superior performance of Haar in this context:

1. **Spatial Localization:** LWD requires a spatially precise mask $M_t(i, j)$ to target specific latent regions. While global transforms like FFT/DCT provide excellent frequency resolution, they sacrifice spatial localization; a high-frequency coefficient corresponds to periodic patterns across the entire image, not specific positions. Recovering spatial energy maps via inverse transformation introduces Gibbs ringing near sharp transitions. These artifacts cause signal leakage into neighboring latent positions, blurring the mask and degrading texture precision (GLCM 0.71 vs 0.74).
2. **Compact Support:** Among wavelets, Haar has the most compact support (2 coefficients), minimizing cross-position interference. This is critical for generating sharp binary training masks. Smoother wavelets (e.g., Daubechies) introduce wider receptive fields, creating “gray areas” at mask boundaries that dilute supervision without semantic benefit.

As shown in Table 5, while FFT is computationally faster due to hardware optimizations, the spatial artifacts degrade generation quality (FID 39.45). Haar achieves the optimal balance, outperforming Daubechies in both efficiency and final texture fidelity.

Table 5: Ablation of Frequency Decomposition Methods. Comparison on Diffusion4k backbone (Aesthetic dataset, 2048 × 2048).

Method	Cost (ms) ↓	FID ↓	Aesthetics ↑	GLCM ↑
LWD (Haar)	1.136	38.74	6.17	0.74
LWD (Daubechies)	1.274	38.92	6.14	0.73
LWD (FFT High-Pass)	0.875	39.45	6.08	0.71

Wavelet Masking Lower Bound. The primary hyperparameter for our wavelet masking strategy is the lower bound l . The value $l = 0.3$ was chosen based on an ablation study (Table 6). This study revealed a clear trade-off: a very low value (e.g., $l < 0.1$) can cause smooth regions to be under-trained, while a very high value (e.g., $l > 0.7$) diminishes the benefit of targeted refinement, causing performance to regress towards the baseline. The value $l = 0.3$ was found to be a robust sweet spot.

Table 6: Ablation on the Masking Lower Bound (l).

Masking Lower Bound (l)	FID ↓	GLCM Score ↑	CLIPScore ↑
0.0	34.15	0.68	0.5411
0.1	33.21	0.72	0.5420
0.3	32.88	0.74	0.5423
0.5	33.46	0.71	0.5419
0.7	34.02	0.69	0.5407

Intuition for the Masking Strategy. The time-dependent masking schedule is designed to allocate more training attention to structurally rich regions of the image, which are identified via higher

972 wavelet energy. The schedule ensures that these areas are refined over more training steps, while
 973 still providing a minimum level of supervision to all regions. This enhances high-frequency details
 974 without overfitting to them.

976 A.3 ANALYSIS OF METHOD PROPERTIES AND TRADE-OFFS

978 **Preservation of Global Structure.** Our wavelet-masked loss is designed to preserve global coherence.
 979 The masking computation targets only the high-frequency subbands (LH, HL, HH), which
 980 encode localized detail. The LL subband, which captures low-frequency, global structure, is not
 981 involved. This ensures that while local refinement is emphasized, the global scene layout and object
 982 structure remain intact.

984 **Robustness and Potential Artifacts.** Our wavelet-based masking strategy is agnostic to the source
 985 of high-frequency information and has proven robust across diverse scenes without introducing
 986 noticeable artifacts. The VAE fine-tuning stage is key to this stability, as it regularizes the latent space
 987 to ensure that the high-frequency energy used for masking corresponds to meaningful content rather
 988 than spurious signals. While we have not observed failure cases in our benchmarks, investigating
 989 domain-specific behavior is an important direction for future work.

991 **On the Synergy of Frequency Suppression and Utilization.** A natural question arises from the
 992 apparent tension between our use of a multi-scale VAE loss to suppress spurious high-frequency
 993 components, and our later use of high-frequency energy to guide the adaptive masking. These two
 994 strategies serve complementary and sequential roles. The VAE loss does not eliminate all high-
 995 frequency content; rather, it penalizes frequency components inconsistent across scales, which often
 996 correspond to noise or artifacts. This regularization aligns the latent spectral distribution more closely
 997 with that of clean, natural images.

998 Crucially, it is precisely this filtered latent space that makes our frequency-based attention meaningful.
 999 Once the latent tensor is regularized, the remaining high-frequency energy (extracted via DWT)
 1000 is more strongly correlated with visually salient features like edges and textures. In other words,
 1001 by denoising the latent representation, the VAE enhances the signal-to-noise ratio of our wavelet
 1002 attention mechanism. This sequential design, first purifying the latent space, then exploiting its
 1003 structured frequency characteristics, ensures our method combines signal-domain regularization and
 1004 content-adaptive computation in a synergistic manner.

1006 B ADDITIONAL RESULTS FOR 4K

1008 To assess the efficacy of our proposed Latent Wavelet Diffusion (LWD), we conduct a detailed
 1009 evaluation focusing on the challenging 4K resolution (4096×4096).

1012 B.1 QUANTITATIVE RESULTS

1014 Table 7 presents the quantitative comparison on 4K image generation. We evaluate the generated
 1015 images using several key metrics: MAN-IQA (Yang et al., 2022) and QualiCLIP (Agnolucci et al.,
 1016 2024), which assess perceptual image quality and alignment with textual prompts, respectively.
 1017 Additionally, we compute the GLCM (Gray-Level Co-occurrence Matrix) score as a measure of
 1018 texture complexity and detail richness in the generated high-resolution outputs. Finally, we report the
 1019 Compression Ratio, indicating the compressibility of the generated images, which can be indicative
 1020 of redundancy or lack of fine details.

1021 Our LWD-enhanced URAE demonstrates competitive performance across all evaluated metrics.
 1022 Notably, it achieves the highest MAN-IQA score and GLCM score, suggesting superior perceptual
 1023 quality and richer textural details compared to the baselines. Furthermore, our LWD + URAE achieves
 1024 a favorable Compression Ratio (28.77), better than URAE and PixArt-Sigma-XL, suggesting a good
 1025 balance between detail and redundancy. These quantitative results underscore the effectiveness of our
 LWD approach in enhancing the visual quality and detail of images generated at 4K resolution.

1026

1027

1028

1029

1030

1031

1032

1033

1034

1035

1036

1037

1038

1039

1040

1041

1042

1043

1044

1045

1046

1047

1048

1049

1050

A man with an afro hairstyle wears futuristic reflective sunglasses and a coat with fur lining, standing in front of a vibrant pink and blue neon sign.



1051

1052

1053

1054

1055

1056

1057

1058

1059

1060

1061

1062

1063

1064

1065

1066

1067

1068



1069

1070

1071

Figure 6: Images generated at 4K resolution with LWD+SANA.

1072

1073

1074

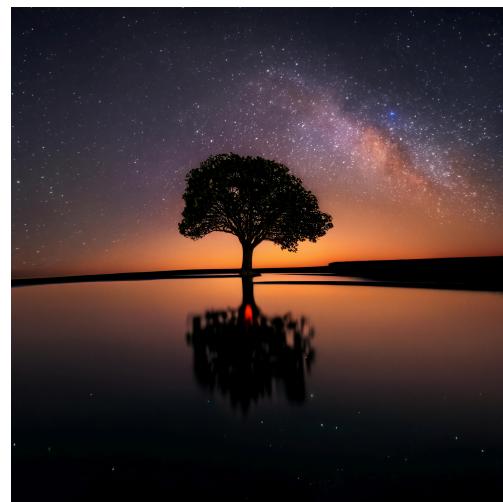
1075

1076

1077

1078

1079



A lone tree stands in calm water reflecting the starry night sky, with the Milky Way stretching above and warm orange hues from a distant horizon.



A person wearing a Spider-Man suit in the game Half-Life Alyx.™

1080

1081

1082

1083

1084

1085

1086

1087

1088

1089

1090

1091

1092

1093

1094

1095

1096

1097

1098

1099

1100

1101

1102

1103



1104

A surreal landscape depicting an ethereal fusion of natural beauty and fantastical architecture, reminiscent of Salvador Dali's dreamlike paintings. From above the clouds, one gazes upon a colossal tower emerging from the earth, its intricate gears visible as it merges seamlessly with a tranquil mountain lake. The scene is bathed in an otherworldly glow, casting lavender and gold hues across the sky, while delicate cherry blossoms flutter gently in the foreground, adding a sense of serenity to this breathtaking vision where time and nature intertwine.

1107

1108

1109

1110

1111

1112

1113

1114

1115

1116

1117

1118

1119

1120

1121

1122

1123

1124



1125

Steampunk airship floating above a misty Victorian cityscape, intricate brass and copper mechanical details, golden hour lighting, billowing clouds, detailed architectural elements, rich warm color palette, cinematic composition.

1126

Figure 7: Images generated at 4K resolution with LWD+URAE.

1127

1128

1129

1130

1131

1132

1133

1134
1135
1136
1137
1138
1139
1140
1141
1142
1143
1144
1145
1146
1147
1148
1149
1150
1151
1152
1153
1154
1155
1156
1157
1158
1159
1160
1161
1162

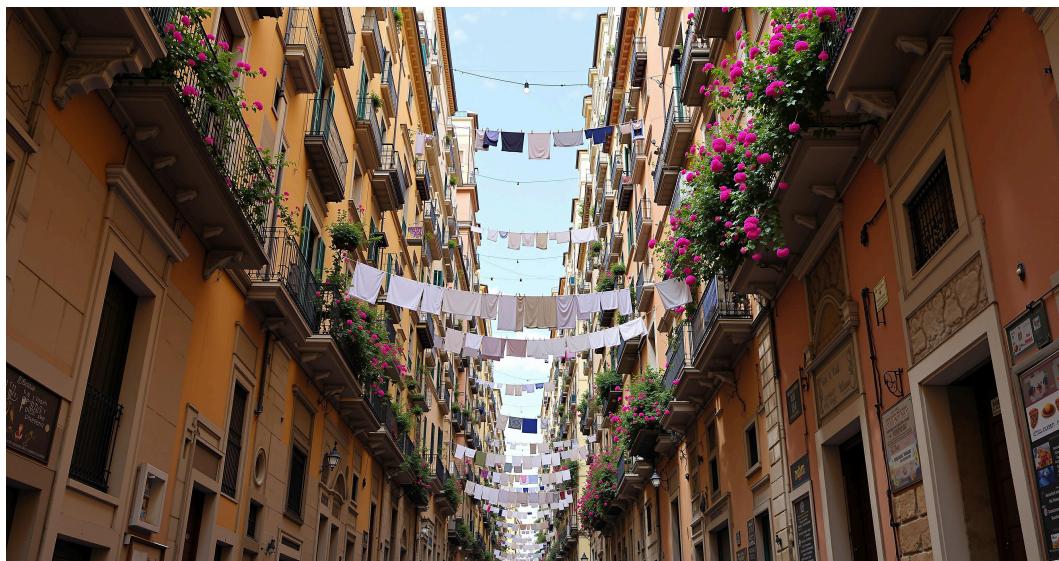


1163 Girl with pink hair, vaporwave style, retro aesthetic,
1164 cyberpunk, vibrant, neon colors, vintage 80s and 90s style,
1165 highly detailed.



1166 A sleek black luxury sedan parked on a rain-soaked city street at night, reflecting
1167 neon lights from nearby buildings. The wet pavement glistens, and the car's
1168 smooth curves are highlighted by the ambient glow of the urban environment.

1169
1170
1171
1172
1173
1174
1175
1176
1177
1178
1179
1180
1181
1182



1183 A narrow and picturesque alley in the historic center of Naples, with laundry hanging out to dry between flower-filled balconies and the inviting
1184 aroma of freshly baked pizza in the air.

1185
1186
1187

Figure 8: Images generated at 4K resolution with LWD+URAE.

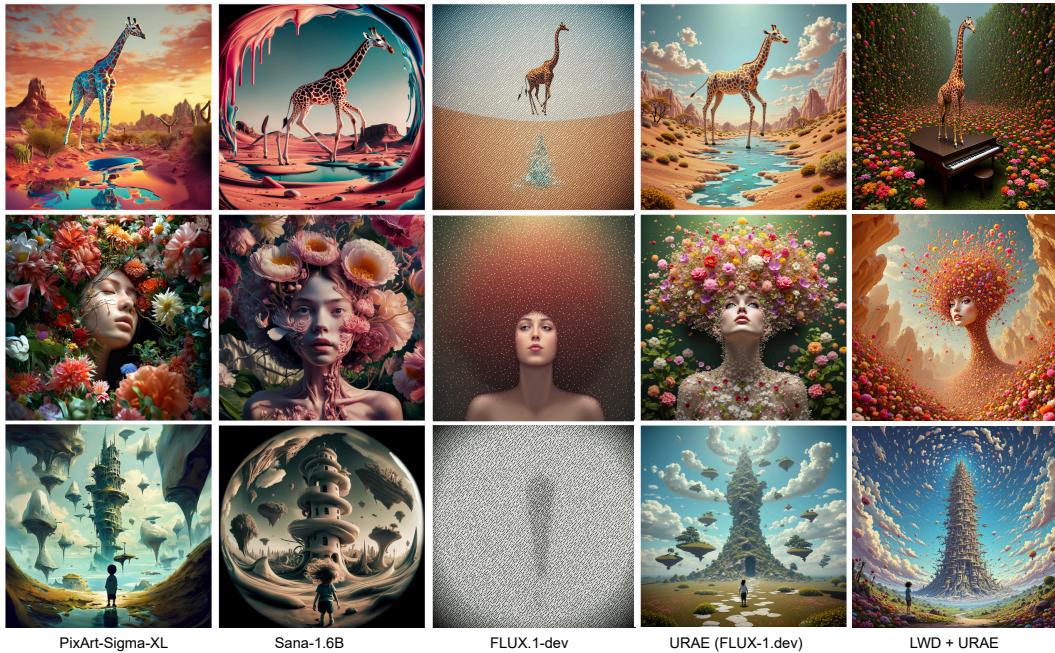


Figure 9: Visual comparison of 4K image generations from LWD and competing baselines.

B.2 QUALITATIVE ANALYSIS

To complement the quantitative evaluation, Figures 10, 11, 13, and 12 provide a qualitative comparison of LWD against selected baselines, presenting side-by-side comparisons of the generated images.

The visual comparisons highlight the benefits of our LWD enhancement. Our method demonstrates the generation of images with finer and more intricate details, particularly noticeable in complex textures and object boundaries. These qualitative observations align with our quantitative findings, reinforcing the effectiveness of the proposed LWD for high-resolution image generation.

Table 7: Evaluation on ultra resolution (4096×4096) image generation task with (Wu et al., 2023) prompts.

Method	MAN-IQA (\uparrow)	QualiCLIP (\uparrow)	GLCM Score \uparrow	Compression Ratio \downarrow
PixArt-Sigma-XL (Chen et al., 2024)	0.2935	0.2308	0.48	45.15
Sana-1.6B (Xie et al., 2025)	0.3288	0.4979	<u>0.71</u>	25.89
FLUX-1.dev (Labs, 2024)	0.3673	0.2564	-	-
URAE (Yu et al., 2025)	<u>0.3850</u>	0.3758	0.41	38.86
LWD + URAE	0.4011	<u>0.3855</u>	0.74	<u>28.77</u>

Selection Protocol All qualitative examples shown in the paper were generated following a strict, reproducible protocol to prevent cherry-picking. Prompts were sourced directly from the original papers of the baseline models (e.g., URAE, Diffusion-4K) or the HPD benchmark dataset. For each prompt, the displayed image is the first output generated using a fixed random seed, applied identically to both the baseline and our LWD-enhanced model. We have included our code in the supplementary materials to ensure full transparency. While subjective preferences for certain images may vary, our method consistently improves objective indicators of texture and detail.

C FREQUENCY-AWARE EVALUATION

To rigorously assess the frequency characteristics of generated images, we propose a suite of frequency-sensitive metrics that extend beyond standard perceptual scores. These metrics are designed to quantify the degree to which generated images preserve the natural frequency distribution

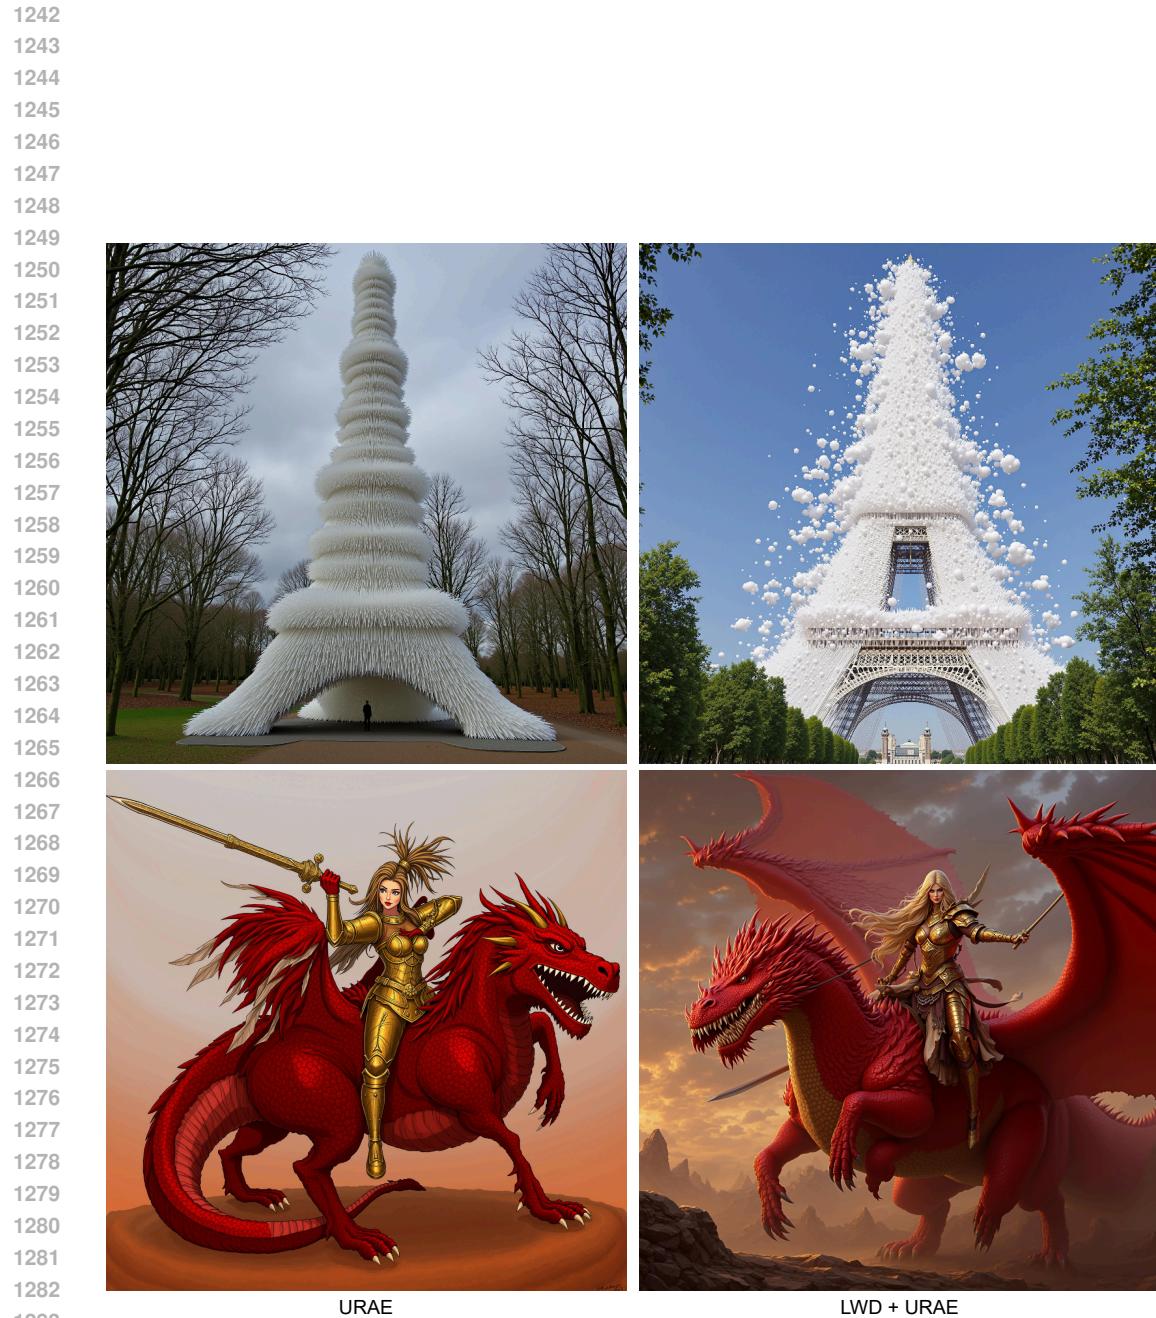


Figure 10: 4K generation of URAE vs LWD + URAE.

Upper caption: "*Eiffel Tower was Made up of more than 2 million translucent straws to look like a cloud, with the bell tower at the top of the building, Michel installed huge foam-making machines in the forest to blow huge amounts of unpredictable wet clouds in the building's classic architecture.*".

Lower caption: "*Barbarian woman riding a red dragon, holding a broadsword, in gold armour.*"

1296 observed in real images, with particular attention to the presence and quality of high-frequency
 1297 details.

1298 **High/Low Frequency Ratio (HLFR):** We decompose each image using a 2D discrete wavelet
 1299 transform (DWT) and compute the total energy in the detail coefficients (high-frequency subbands:
 1300 LH, HL, HH) and in the approximation coefficients (low-frequency LL subband). The HLFR is
 1301 defined as the ratio of high-frequency to low-frequency energy:

$$1303 \quad \text{HLFR} = \frac{E_{LH} + E_{HL} + E_{HH}}{E_{LL}}.$$

$$1304$$

1305 This ratio reflects the relative emphasis on fine-scale structures. A value similar to the reference
 1306 (real images) indicates a natural distribution of frequency content. Large deviations can signal
 1307 oversmoothing or hallucinated detail.

1308 **Ratio Difference from Real (RDR):** To quantify deviation from the natural HLFR, we compute the
 1309 absolute difference between the HLFR of the generated image and the real reference:

$$1311 \quad \text{RDR} = |\text{HLFR}_{\text{gen}} - \text{HLFR}_{\text{real}}|.$$

$$1312$$

1313 Lower values are better, indicating better alignment with the natural frequency distribution.

1314 **Wavelet Quality Score (WQS):** This metric evaluates the structural similarity between generated and
 1315 real images in the wavelet domain, where frequency components are explicitly separated by scale and
 1316 orientation. Given a multi-level discrete wavelet transform (DWT) of both the generated image I_g
 1317 and the reference image I_r , we compute the SSIM and MSE for each corresponding subband s across
 1318 all decomposition levels l . The final WQS aggregates the per-subband scores using frequency-aware
 1319 weights:

$$1320 \quad \text{WQS} = \sum_{l=1}^L \sum_{s \in \{LL, LH, HL, HH\}} w_{l,s} \cdot \text{SSIM}(I_r^{l,s}, I_g^{l,s}) - \lambda \cdot \text{MSE}(I_r^{l,s}, I_g^{l,s}),$$

$$1321$$

$$1322$$

$$1323$$

1324 where $w_{l,s}$ are weights that can prioritize perceptually important subbands (e.g., low-frequency LL or
 1325 high-frequency HH), and λ is a scaling factor that penalizes distortion. The score is normalized to
 1326 $[0, 1]$, where 1 indicates perfect structural alignment. Higher WQS values reflect better reconstruction
 1327 fidelity across frequency scales, meaning the model preserves both coarse structure and fine texture.

1328 **High-Frequency Energy (HFE):** This metric quantifies the total energy of the image's high-frequency
 1329 components after wavelet decomposition. For a given decomposition level, we define:

$$1330 \quad \text{HFE} = \sum_{l=1}^L (\|I^{l,LH}\|^2 + \|I^{l,HL}\|^2 + \|I^{l,HH}\|^2).$$

$$1331$$

$$1332$$

$$1333$$

$$1334$$

1335 This value provides an absolute measure of fine-scale activity in the image. While real images have
 1336 characteristic HFE ranges, excessive HFE may indicate artifacts or noise, and too little HFE suggests
 1337 oversmoothing. Alignment with the real HFE is typically ideal.

1338 **High-Frequency Emphasis Index (HFEI):** This metric evaluates how much the model over- or
 1339 under-emphasizes high-frequency content relative to the real distribution. We define it as:

$$1340 \quad \text{HFEI} = \left(\frac{\text{HFE}_{\text{gen}}}{\text{TotalEnergy}_{\text{gen}}} \right) - \left(\frac{\text{HFE}_{\text{real}}}{\text{TotalEnergy}_{\text{real}}} \right),$$

$$1341$$

$$1342$$

$$1343$$

1344 where total energy is computed over all wavelet subbands. $\text{HFEI} > 0$ indicates the generated image
 1345 places more emphasis on high frequencies than real images (potentially hallucinated detail), while
 1346 $\text{HFEI} < 0$ indicates a loss of fine detail. An HFEI near zero is ideal.

1347 **Perceptual Metrics:** For completeness, we also report FSIM (Zhang et al., 2011) and MS-SSIM
 1348 (Wang et al., 2003), which capture visual similarity and structural coherence, respectively. Both
 1349 metrics are bounded between 0 and 1, with higher values indicating better perceptual quality.

1350
1351
1352
1353
1354
1355
1356
1357
1358
1359
1360
1361
1362
1363
1364
1365
1366
1367
1368
1369
1370
1371
1372
1373



1374
1375
1376
1377
1378
1379
1380
1381
1382
1383
1384
1385
1386
1387
1388
1389
1390
1391



PixArt-Sigma-XL



LWD + PixArt-Sigma-XL

1392
1393
1394
1395
1396
1397
1398
1399
1400
1401
1402
1403

Figure 11: 2K generation of PixArt-Sigma-XL vs LWD + PixArt-Sigma-XL.

Upper caption: "An elderly man with a prominent, bushy beard and deep-set eyes wears a vibrant orange turban, his weathered face marked by lines of age and experience.". Lower caption: "A lone figure on a horse stands in a misty forest, gazing up at a tall, multi-tiered temple surrounded by towering trees and soft, diffused light. Steam rises from the rocks near a stream, creating an atmospheric scene of tranquility and mystery."

1458

1459

1460

1461

1462

1463

1464

1465

1466

1467

1468

1469

1470

1471

1472

1473

1474

1475

1476

1477

1478

1479

1480

1481

1482

1483

1484

1485

1486

1487

1488

1489

1490

1491

1492

1493

1494

1495

1496

1497



SD3-Diff4k-F16



LWD + SD3-F16

1500 Figure 13: 2K generation of SD3-Diff4k-F16 vs LWD + SD3-F16.

1501 Upper caption: "A serene landscape features a winding river, flanked by trees with autumn foliage,
1502 leading to a rustic wooden cabin with a corrugated roof, set against a softly blurred background.".
1503 Lower caption: "A grand interior featuring intricate stained glass windows, an elaborate rose
1504 window, ornate frescoes depicting biblical scenes, and elegant chandeliers illuminating the richly
1505 decorated walls and arches."

1506

1507

1508

1509

1510

1511

1512
 1513 Table 8: Comparison of frequency-sensitive metrics across different methods on Aesthetic-4K (Zhang
 1514 et al., 2025) validation set.

Metric	HLFR	RDFR ↓	WQS ↑	HFE	HFEI ↓	FSIM ↑	MS-SSIM ↑
Real	0.0560	0.0000	1.0000	0.0140	0.0000	1.0000	1.0000
Sana-1.6B (Xie et al., 2025)	0.0784	0.0558	0.4673	0.0196	0.6108	0.6128	0.1324
LWD + Sana-1.6B	0.0610	0.0537	0.4701	0.0144	0.5227	0.6217	0.1324
SD3-Diff4k-F16 (Zhang et al., 2025)	0.0691	0.0470	0.4624	0.0158	0.8064	0.6155	0.1296
LWD + SD3-F16	0.0555	0.0437	0.4735	0.0144	0.4826	0.6245	0.1521
PixArt-Sigma-XL	0.0550	0.0409	0.4763	0.0119	0.6296	0.1354	0.4255
LWD + PixArt-Sigma-XL	0.0564	0.0500	0.4730	0.0150	0.6239	0.1478	0.5094

1523
 1524 Table 9: Training configurations and efficiency gains for LWD across different backbones.

Backbone	Res.	Batch Size	Iterations	Training Time
LWD + URAE (Flux)	2048	1	2k	~4 hours
LWD + URAE (Flux)	4096	1	2k	~24 hours
LWD + Diff4K (SD3)	2048	8	10k	~48 hours
LWD + SANA	2048/4096	2/1	33k	~24 hours
LWD + PixArt- Σ	2048	2	1.5k	~24 hours

1533 **Hyperparameters** All experiments were conducted on a system with 4 NVIDIA A100 GPUs. Our
 1534 VAE fine-tuning objective (Equation 1) balances four terms. The weights were adopted from prior
 1535 work (Wu et al., 2023), which provides extensive validation for these values. Following established
 1536 practice, we set the weights to $\alpha = 0.25$, $\beta = 0.001$, and $\lambda = 0.05$.

1537
 1538 **Training Overhead** LWD introduces a marginal overhead during training. The minor increase in
 1539 peak GPU memory usage (Table 10) is due to the storage of intermediate tensors for the wavelet
 1540 transform and energy masks. These tensors are small (the size of the latent map) and their memory
 1541 footprint is insignificant compared to the large diffusion model backbone.

1542
 1543 Table 10: Computational Overhead Analysis during training on a single A100 GPU (64GB).

Method	Mem. Usage (%)	Mem. Usage (GB)	Time per 20 Steps (s)
Sana	90.5	57.9	~47
Sana + LWD	93.9	60.1	~47

1548 Another key advantage of LWD is its efficiency. It is a training-only strategy that requires zero
 1549 architectural modifications. Consequently, an LWD-enhanced model has the exact same number of
 1550 parameters and identical inference time as its baseline counterpart.