

LENS: Low-Frequency Eigen Noise Shaping for Efficient Diffusion Sampling


Haewon Jeon, Si-Hyeon Lee

School of EE, KAIST, South Korea

40th Conference on Neural Information Processing Systems (NeurIPS 2026)

01 Introduction

Distilled Diffusion Model



Faster sampling → Potential quality degradation

- ✓ Prior approaches improve generation by optimizing the initial noise

01 Introduction

Challenge


- Existing noise optimization is computationally expensive

[1] Iterative Optimization



- Repeated model evaluations
- High test-time optimization cost

[2] Learned Noise Modulation



- Uses diffusion backbone
- Full-dimensional modulation

01 Introduction

Why Focus on Low-Frequency Noise?

Low-Frequency Components

Global Structure and Object Layout

High-Frequency Components


Fine details,
with limited impact on global structure

Key Insight

Efficient noise modulation can be achieved by selectively modifying low-frequency components


02 Our Method

LENS (Low-Frequency Eigen Noise Shaping)




02 Our Method

LENS (Low-Frequency Eigen Noise Shaping)




02 Our Method

LENS (Low-Frequency Eigen Noise Shaping)




02 Our Method

LENS (Low-Frequency Eigen Noise Shaping)




02 Our Method

LENS (Low-Frequency Eigen Noise Shaping)



02 Our Method

LENS (Low-Frequency Eigen Noise Shaping)



02 Our Method

Principled Low-Frequency Modulation

[Theoretical Justification]

Assumption

$$|r(g(w_L, w_H)) - \bar{r}(w_L)| \leq \epsilon \text{ where } \bar{r}(w_L) := \mathbb{E}_{w_H \sim q_H}[r(g(w_L, w_H))]$$

KL Divergence
Bound

$$D_{\text{KL}}(q^* || \tilde{q}^*) \leq 2\epsilon$$

q^* : full tilted distribution, $\tilde{q}^*$: low-frequency-only tilted distribution

[Training Objective]

$$\mathcal{L}(\phi) = \mathbb{E}_{w \sim q}[\frac{1}{2} \|h_\phi(w_L)\|^2 - r(g(w_L + h_\phi(w_L), w_H))]$$

→ LENS is trained to maximize generation rewards while regularizing deviations from the original noise distribution.

03 Experiment

Generation Quality (GenEval2)

Model	Time(s) ↓	Object ↑	Attribute ↑	Count ↑	Position ↑	Verb ↑	Mean ↑
SD-Turbo	0.107	72.27	42.65	21.15	28.1	5.60	33.95
+ ReNo	11.371	81.02	50.42	31.91	32.58	6.90	40.57
+ HyperNoise	0.194	74.94	43.00	23.95	27.67	5.23	34.96
+ LENS (Ours)	0.112	75.07	43.08	26.08	28.46	5.78	35.69
Hyper-SDXL	0.414	71.94	51.97	26.88	21.14	3.28	35.04
+ ReNo	45.13	80.06	56.79	32.61	27.53	3.59	40.12
+ HyperNoise	0.694	71.97	52.02	27.06	21.58	4.76	35.48
+ LENS (Ours)	0.424	72.96	53.47	27.84	21.45	4.39	36.02
SANA-Sprint	0.336	80.23	54.73	33.01	33.31	15.53	43.36
+ ReNo	14.964	83.16	56.57	34.73	38.36	13.44	45.25
+ HyperNoise	0.532	83.59	54.18	34.84	36.32	20.66	45.92
+ LENS (Ours)	0.352	82.98	55.27	36.12	36.38	18.21	45.79

03 Experiment

Computational Efficiency

Model	Method	FLOPs	Parameters	Inference-time
SD-Turbo	HyperNoise	997.1 G	136.5 M	70.78 ms
	LENS (Ours)	2.4 G	4.8 M	5.68 ms
Hyper-SDXL	HyperNoise	7.8 T	395.3 M	207.91 ms
	LENS (Ours)	12.3 G	5.3 M	9.83 ms
SANA-Sprint	HyperNoise	1.7 T	314.7 M	173.62 ms
	LENS (Ours)	2.2 G	10.0 M	10.84 ms

03 Experiment

Qualitative Results



04 Key Contribution

- ✓ **Low-Dimensional Noise Modulation**
LENS selectively modulates low-frequency PCA coefficients of the diffusion initial noise.
- ✓ **Principled and Lightweight Framework**
LENS uses a lightweight, standalone network supported by theoretical justification.
- ✓ **Competitive Quality and Efficient Generation**
LENS achieves competitive quality with 400–700× fewer FLOPs and 10–20× faster noise modulation than HyperNoise.

Thank you



Project