

Latent Diffusion Models with Masked AutoEncoders

Junho Lee^{1*} Jeongwoo Shin^{1*} Hyungwook Choi¹ Joonseok Lee^{1†}

¹Seoul National University, Seoul, Korea

{joon2003, sswss, chooi221, joonseok}@snu.ac.kr



Figure 1. **Class-conditional generation samples of our proposed LDMAE.** Trained on ImageNet-1K, resolution of 256×256 .

Abstract

In spite of the remarkable potential of Latent Diffusion Models (LDMs) in image generation, the desired properties and optimal design of the autoencoders have been underexplored. In this work, we analyze the role of autoencoders in LDMs and identify three key properties: latent smoothness, perceptual compression quality, and reconstruction quality. We demonstrate that existing autoencoders fail to simultaneously satisfy all three properties, and propose Variational Masked AutoEncoders (VMAEs), taking advantage of the hierarchical features maintained by Masked AutoEncoders. We integrate VMAEs into the LDM framework, introducing Latent Diffusion Models with Masked AutoEncoders (LDMAEs). Through comprehensive experiments, we demonstrate significantly enhanced image generation quality and computational efficiency.

1. Introduction

Diffusion models (DMs) [8, 17, 19, 29, 31, 41] have recently set new standards in image generation, but their iterative denoising process incurs high computational cost. While deterministic sampling [26, 39] reduces the number of sampling steps, the core inefficiency of operating directly in pixel space remains.

To address these challenges, Latent Diffusion Models (LDMs)[32] shift the denoising process from pixels to a compressed space, focusing on semantic representations with lower computational cost. For this, Variational Autoencoders (VAEs)[20] are adopted to enforce probabilistic alignment with a prior, enabling smooth interpolation, generative sampling, and robust generalization.

However, this latent space alignment loss often negatively impacts the reconstruction performance, leading to a trade-off between latent space alignment and reconstruction quality. To mitigate this, LDMs, such as Stable Diffusion [11, 32], incorporate adversarial and perceptual losses to further improve the VAEs, which is commonly referred to as the StableDiffusion-VAEs (SD-VAEs). Subsequent stud-

*Equal contribution.

†Corresponding author.




Figure 2. Comparison of autoencoding methods across key metrics: latent smoothness, perceptual compression quality, reconstruction quality (at pixel and perceptual levels), and overall generation quality. Each score is linearly rescaled so that the outermost and innermost grids indicate the highest and lowest scores, respectively.

ies further improved it by enhancing reconstruction quality via increased latent dimensionality [5, 10, 22, 45] or by improving efficiency with more lightweight architectures [34].

Despite the notable success of the additional loss terms in autoencoders, their specific roles in shaping the characteristics of autoencoders, as well as their impact on the entire LDM framework, remain underexplored. A fundamental question arises: what properties should an autoencoder possess to effectively support LDMs? Since the choice of autoencoder directly influences the quality and efficiency of LDMs, this is a crucial question for designing an autoencoder. To address this question, we propose three key properties that an autoencoder should satisfy to effectively compress data, preserve essential semantics, and eventually facilitate high-quality image generation in the LDM framework: 1) latent smoothness, 2) perceptual compression quality, and 3) reconstruction quality.

First, *latent smoothness* ensures that minor perturbations in the latent representation do not drastically change the generated output. A sufficiently smooth latent space allows the model to generate high-quality images robustly, even when the diffusion model predicts latent codes with some error. We reveal through an empirical analysis on the stability of autoencoder embeddings under small perturbations (Sec. 5.1) that autoencoders with deterministic encoding, such as vanilla AutoEncoders (AEs), lack this property, whereas probabilistic encoding, as used in VAEs, much better satisfies this requirement.

Second, *perceptual compression quality* measures how effectively an autoencoder encodes the images into a simplified latent space, maximizing the reduction of perceptual details while preserving semantics as much as possible. However, the boundary between ‘semantics’ and ‘perceptual details’ is inherently ambiguous, as they exist on a continuous spectrum. This spectrum ranges from pixel-level (minimal compression) to object-level (maximum compression). In

Sec. 5.1, we observe that SD-VAEs aggressively compress input images, resulting in compression close to the object-level. This simplifies the latent space, making it easier for diffusion models to predict latent codes, but sacrifices fine-grained details. To mitigate this, we propose *hierarchical compression*, which organizes the latent space into a continuous hierarchy: first grouping images into object-level clusters, then subdividing each object into part-level clusters (e.g., facial features), and finally refining them into fine-detail clusters. This structure preserves global simplicity while retaining local details, effectively balancing compression and reconstruction quality.

Lastly, *reconstruction quality* ensures the decoder to accurately reconstruct the original image from its latent representation. We claim that this property needs to be assessed both at the perceptual level (e.g., measured by rFID or LPIPS) and at the pixel level (e.g., evaluated by PSNR and SSIM). In literature, perceptual metrics such as FID are widely adopted, while pixel-level reconstruction quality has been overlooked, potentially due to the lack of reference images. Nevertheless, ensuring high pixel-level reconstruction quality is crucial, particularly for tasks that require precise spatial or numerical consistency, such as image editing, inpainting, or super-resolution. Our empirical analysis reveals that SD-VAEs achieve high reconstruction quality at perceptual level, while their pixel-level reconstruction quality is notably weaker.

From systematical assessments of these properties (see Sec. 5.1 for details), we provide a concise summary of various autoencoders in Fig. 2. We clearly observe that no existing autoencoders satisfy all of these desired properties, lacking at least one of them. Inspired by the recent findings of masked autoencoding [36] that it constructs a well-aligned hierarchical feature space, we propose a novel transformer-based *Variational Masked AutoEncoders (VMAEs)*. VMAE achieves both latent smoothness and high reconstruction quality through probabilistic alignment, tailored loss design, and structural enhancements, while retaining the hierarchical compression benefits of MAEs. Additionally, VMAE significantly improves computational efficiency, using only 13.4% of the parameters and 4.1% of the GFLOPs of SD-VAE. It also converges 2.7 times faster with greater training stability.

Furthermore, we incorporate our VMAEs into the LDMs framework, referred to as Latent Diffusion Models with Masked AutoEncoders (LDMAEs). Through a series of experiments, we verify that LDMAEs improve generative performance across multiple datasets and settings. These findings support our hypothesis that the three key properties contribute to effective latent diffusion.

We summarize our contributions as follows:

- We identify smooth latent space, effective perceptual compression, and high reconstruction quality as desired properties of autoencoders for LDMs with empirical evidence from thoroughly designed experiments.

- Based on our analysis, we introduce VMAEs, a novel autoencoder designed to improve all three key factors with significant model efficiency.
- Through extensive experiments, we demonstrate superior generation performance of our LDMAE, equipped with our VMAE as its autoencoder.

2. Related Work

Diffusion Models. Diffusion models generate images by reversing a forward noise process. DDPM [17] introduces a Markovian framework where Gaussian noise is incrementally added, and a neural network learns to denoise by optimizing Evidence Lower Bound, but requires many sampling steps. Score-based models [41] generalize diffusion to continuous time, enabling efficient solvers via reverse-time stochastic differential equations. ADM [8] improves scalability with architectural refinements and classifier guidance.

These standard diffusion models [8, 17, 39] operate in pixel space, which presents significant computational challenges, especially for high-resolution sampling. To address this, Latent Diffusion Models (LDMs) [31, 32] shift the diffusion process to a learned latent space, using a pre-trained autoencoder to represent data in a perceptually compressed form, removing redundant pixel-level details.

Autoencoders for Image Generation. Autoencoders [13, 20, 33, 44] encode input data into a compact latent space that preserves structure, reducing computational overhead and enabling tasks like interpolation and reconstruction.

Recent studies have explored integrating autoencoder architectures into image generation. Two main approaches are based on MAEs [13] and VAEs [20]. MAEs-based methods, like MaskGIT [1] and MAGE [24], adapt masked encoding to discrete token-based space, while MAR [25] utilizes diffusion to generate continuous tokens. VAEs-based methods sample images from a Gaussian distribution. VQ-VAE [43] and VQGAN [10] further employ discrete latent spaces with a codebook for improved generation.

LDMs [32], however, directly generate compressed latents from autoencoders, which can be decoded back into images, accelerating generation but often degrading reconstruction quality. To address this, StableDiffusion3 [11], Emu [5], and FLUX [22] introduce larger latent dimensionality to capture finer details. VA-VAE [45] adds alignment losses with vision foundation models, and DC-AE [3] aims for high spatial compression without sacrificing fidelity.

3. Preliminary

3.1. Autoencoders

We compare four classical autoencoder variants within LDMs [32] and include MAEs [13], the foundation of our model. See Appendix A for details.

The vanilla **AutoEncoders (AEs)** [33] are optimized solely by the reconstruction loss, compressing the input into a more compact latent space and reconstructing it back. **Denosing AutoEncoders (DAEs)** [44] extend AEs by introducing noise to the input data and training the model to recover the original noise-free input for robust latent features with improved generalization. **Variational AutoEncoders (VAEs)** [20] adopt a probabilistic framework, encoding input data into a latent distribution instead of a fixed vector, with KL-divergence regularizing it towards a Gaussian distribution to enable sampling. **StableDiffusion VAEs (SD-VAEs)** [11, 32] are built upon VQGAN [10], the autoencoder adopted in traditional LDMs [32]. Following the VQGAN, SD-VAEs integrate an additional adversarial network and train with perceptual loss for an improved perceptual quality. Unlike VQGAN, however, SD-VAEs omit the quantization layer entirely and instead use continuous features. In this paper, we follow the SD-VAEs settings from StableDiffusion3 [11], the state-of-the-art LDM, unless noted otherwise.

Masked AutoEncoders (MAEs) [13] were introduced as a self-supervised representation learning method using Vision Transformers (ViTs) [9]. The encoder maps a randomly masked-out image to fixed vectors, while the decoder predicts the masked regions using these latent vectors along with learnable mask tokens.

3.2. Latent Diffusion Models

Standard DMs [8, 17, 39] operate in pixel space $\mathbf{x} \sim p_{\text{data}}$ and optimize the following objective:

$$\mathcal{L}_{\text{DM}} = \mathbb{E}_{\mathbf{x}, \epsilon \sim \mathcal{N}(0,1), t} \left[\|\epsilon - \epsilon_{\theta}(\mathbf{x}_t, t)\|_2^2 \right],$$

where ϵ_{θ} is the model that predicts the original noise ϵ given the corrupted input \mathbf{x}_t at timestep t .

LDMs [32] extend this framework to the latent space $\mathbf{z} \sim \mathcal{E}(\mathbf{x})$ encoded by an encoder \mathcal{E} , instead of pixel space:

$$\mathcal{L}_{\text{LDM}} = \mathbb{E}_{\mathcal{E}(\mathbf{x}), y, \epsilon \sim \mathcal{N}(0,1), t} \left[\|\epsilon - \epsilon_{\theta}(\mathbf{z}_t, t, \tau_{\theta}(y))\|_2^2 \right],$$

where $\tau_{\theta}(y)$ projects the conditioning inputs y into the latent space, enabling multi-modal generative tasks.

4. Variational Masked AutoEncoders

We introduce our novel autoencoder, called *Variational Masked AutoEncoders (VMAEs)*, which embody the three desirable properties for LDM framework, *i.e.*, smooth latent space (Sec. 4.1), perceptual compression (Sec. 4.2), and high reconstruction quality (Sec. 4.3).

4.1. Support of Latent Space

We emphasize that *probabilistic latent space* is crucial for training LDMs, as it ensures a *smooth* latent space—a fundamental requirement of the diffusion process. An

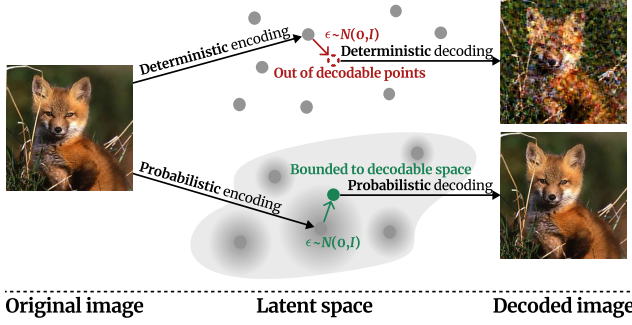


Figure 3. **Sparse latent points vs. smooth latent space.** Deterministic decoders allow only the exact points to be *decodable* due to its *sparse latent points*. In contrast, probabilistic decoders, with a *smooth latent space*, ensure that all latent points in the vicinity of an encoded point remain within the *decodable* space.

overly sparse or low-dimensional latent manifold can hinder diffusion modeling by causing inconsistent score matching [4, 40].

Intuitively, when the target space is sparse, the diffusion model needs to learn *exact* point mappings, making training highly challenging. In LDMs, latent sparsity is closely tied to decoder capability—if the decoder can reconstruct only from a limited set of discrete points, the diffusion model is constrained to generating those exact points for successful reconstruction. In contrast, in a *fully-supported* latent space—*i.e.*, where the decoder reliably reconstructs from any point—the diffusion model is allowed to flexibly generate latents within a vicinity of the target point, enabling a more robust generation process.

This phenomenon is illustrated in Fig. 3, comparing deterministic autoencoders (AEs, DAEs) with probabilistic ones (VAEs, SD-VAEs, VMAEs). Deterministic models encode inputs into fixed vectors, creating a sparse latent space where only exact encodings are *decodable*, while most other points remain inaccessible. Even slight perturbations by Gaussian noise make decoding difficult. In contrast, probabilistic ones establish a continuous latent space, ensuring that points within a noise-induced vicinity remain *decodable*, enabling stable reconstruction. See Sec. 5.1 for a detailed analysis.

Reflecting this observation, we design our VMAEs to encode input data $\mathbf{x} \sim p_{\text{data}}(\mathbf{x})$ as probabilistic distributions $q_{\phi}(\mathbf{z}|\mathbf{x})$ rather than fixed points, promoting smoother latent space. Additionally, we regularize the latent distribution toward a Gaussian prior $p(\mathbf{z})$ via KL divergence, ensuring strict adherence to the Variance Preserving (VP) probability path condition [41], *i.e.*, unit variance, by minimizing

$$\mathcal{L}_{\text{reg}} = \mathbb{E}_{p_{\text{data}}(\mathbf{x})} [\text{D}_{\text{KL}}(q_{\phi}(\mathbf{z}|\mathbf{x})|p(\mathbf{z}))]. \quad (1)$$

4.2. Perceptual Compression

Following the original LDMs [32], *perceptual compression* refers to encoding raw data into a compact latent space where

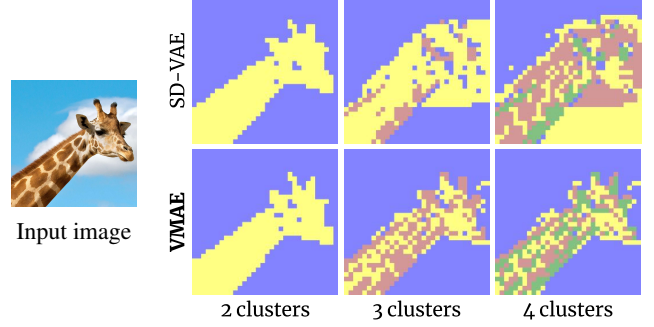


Figure 4. **Separability of latent features.** Our VMAE achieves a more highly separable latent space while maintaining strong semantic grouping, indicating *hierarchical* compression. In contrast, SD-VAE primarily exhibits semantic clustering without clearly differentiating features within each cluster.

fine-grained perceptual details are discarded, while preserving macroscopic semantics. Then, the decoder is responsible for recovering these details; that is, an ideal autoencoder encodes semantically similar features identically in the latent space, and the decoder restores the distinctiveness by adding specific perceptual details back.

However, it is inherently impossible to precisely define the boundaries between ‘semantics’ (to be preserved) and ‘details’ (to be compressed). Rather, these aspects lie on a continuous spectrum, ranging from the lowest pixel-level (no compression) to the most abstract object-level (most compressed). Intermediate levels can include parts (e.g., eyes or nose) or patterns (e.g., texture or color). Given this, we claim that it is crucial to design an autoencoder with a compression level optimized for LDM framework.

Fig. 4 (top) illustrates unsupervised clustering of the local features of SD-VAE for a giraffe image, with various number of clusters. While SD-VAE clearly distinguishes the main object (giraffe) from the background (2 clusters), it struggles to capture finer details like fur patterns (4 clusters). This suggests that its compression ability is limited to the highest object-level, renouncing to represent finer details in the latent space. This simplified latent space can facilitate the learning process of the diffusion model [2], but it may result in the loss of perceptual details at decoding, as most fine details are discarded during the encoding process. Also, as the fine-grained features are entangled in the latent space, downstream tasks like image editing, where disentangled features are particularly critical [12, 14, 27, 37] would not be performed well. Based on these intuitions, we explore an alternative autoencoder with optimal perceptual compression.

Recent studies [30, 36] have demonstrated that the encoded features of MAEs are highly distinguishable based on the perceptual visual patterns such as texture and color, learned through its *masked-part prediction* objective. Moreover, the latest analysis on MAEs [36] further confirms that

this training strategy leads the encoded patch vectors to be *hierarchically* clustered in the embedding space from the abstract objects to simpler visual patterns.

These findings pave the way to equip an autoencoder with *continuously hierarchical compression* by leveraging MAEs. That is, latent features are first compressed at the object-level, then progressively diversified down to the pattern-level hierarchically, with the degree of clustering decreasing at each stage. Fig. 4 (bottom) illustrates that our MAEs-based latent features are gradually distinguishable from high-level objects (2 clusters) to finer fur patterns (4 clusters), showcasing its hierarchical compression. This hierarchical compression meets the needs of both autoencoders and LDMs: 1) strongly clustered features at higher levels simplify diffusion model training, and 2) discriminative features across various levels create a highly disentangled, yet hierarchically structured latent space, enabling high reconstruction quality.

Considering these aspects, we incorporate the masked-part prediction objective for hierarchical perceptual compression. Specifically, the probabilistic decoder p_θ is trained to predict the masked regions \mathbf{x}_m , given only the latent variables of the visible regions, $\mathbf{z} \sim q_\phi(\mathbf{z}|\mathbf{x}_v)$, by optimizing:

$$\mathcal{L}_M = \mathbb{E}_{p_{\text{data}}(\mathbf{x}), p(\mathbf{x}_v|\mathbf{x}), p(\mathbf{x}_m|\mathbf{x}), q_\phi(\mathbf{z}|\mathbf{x}_v)} [-\log p_\theta(\mathbf{x}_m|\mathbf{z})].$$

As $\mathbf{z} \sim q_\phi(\mathbf{z}|\mathbf{x}_v)$ in masked-part prediction, the regularization for the latent distribution in Eq. (1) is modified to

$$\mathcal{L}_{\text{reg}} = \mathbb{E}_{p_{\text{data}}(\mathbf{x}), p(\mathbf{x}_v|\mathbf{x})} [\text{D}_{\text{KL}}(q_\phi(\mathbf{z}|\mathbf{x}_v)|p(\mathbf{z}))].$$

4.3. Perceptual Reconstruction

Although autoencoders solely trained to minimize the pixel-wise distance (e.g., AEs) demonstrate strong reconstruction capability, they do not necessarily recover perceptual details in the pixel space. As illustrated in Fig. 5, AE achieves an optimal reconstruction with a very low MSE. However, closer inspection of the boxed region reveals subtle differences, such as the lumped details of the ropes on the cow’s head, which affect perceptual fidelity. For VAE, neither pixel accuracy nor perceptual fidelity is achieved, resulting in blurry reconstruction, losing most perceptual details. To address the loss of perceptual details in decoding, we train our VMAEs with both perceptual and reconstruction losses.

Since the reconstruction loss can be applied only to the visible regions \mathbf{x}_v in the masked-part prediction scheme,

$$\mathcal{L}_R = \mathbb{E}_{p_{\text{data}}(\mathbf{x}), p(\mathbf{x}_v|\mathbf{x}), q_\phi(\mathbf{z}|\mathbf{x}_v)} [-\log p_\theta(\mathbf{x}_v|\mathbf{z})].$$

For perceptual loss, we use LPIPS [46] loss, implemented as a reconstruction loss at the feature level:

$$\mathcal{L}_P = \mathbb{E}_{p_{\text{data}}(\mathbf{x}), q_\phi(\mathbf{z}|\mathbf{x}_v), p_\theta(\hat{\mathbf{x}}|\mathbf{z})} \left[\sum_l w_l \|\psi_l(\mathbf{x}) - \psi_l(\hat{\mathbf{x}})\|_2^2 \right],$$

where ψ_l extracts features up to the l -th layer of a pre-trained network (e.g., VGG [38]), with w_l as its layer weight.




Figure 5. **Key factors for reconstruction quality.** Although all autoencoders achieve a comparable pixel-level reconstruction (in MSE), they differ in perceptual details especially within the boxed region (in rFID), suggesting that the autoencoders should be optimized for both pixel-level accuracy and perceptual restoration.

4.4. Model Architecture and Training Objective

Putting all ideas together, we finally propose the model architecture and training objective of our proposed method, Variational Masked AutoEncoders (VMAEs) in detail.

Model Architecture. We adopt a symmetric ViT-based encoder and decoder. At encoding, an image is spatially downsampled by patchification and fed into Transformer layers. The encoded features are then reduced to a pre-defined size. With the symmetric structure, the decoding process mirrors the encoding process in reverse order.

Training Objective. Based on our observations of latent space, we claim that the three aforementioned properties are essential for autoencoders to effectively support LDMs. To satisfy these, we design our encoder and decoder as probabilistic functions to construct a smooth latent space, regularized with a KL divergence to satisfy the VP condition (Sec. 4.1). To preserve discriminative features essential for hierarchical compression [36], we allow the predicted mean to be learnable without enforcing it to zero, while still regularizing the variance to be unit. More importantly, we incorporate the masked-part prediction loss to achieve hierarchical compression (Sec. 4.2) and the perceptual loss for perceptually accurate reconstruction (Sec. 4.3). Overall, the final training objective of our VMAE is

$$\mathcal{L}_{\text{VMAE}} = \mathcal{L}_R + \lambda_M \cdot \mathcal{L}_M + \lambda_P \cdot \mathcal{L}_P + \lambda_{\text{reg}} \cdot \mathcal{L}_{\text{reg}}, \quad (2)$$

where λ_M , λ_P and λ_{reg} are scaling hyperparameters.

5. Evaluation on Autoencoders

In this section, we first evaluate our VMAEs in the aspect of latent space, comparing with other autoencoders introduced in Sec. 3.1. All autoencoders are trained on the ImageNet-1K [7] training set with their respective optimal hyperparameters using 8 NVIDIA A100 GPUs (40GB). Elaboration on datasets and implementation details can be found in Appendix B.1.




Figure 6. **Robustness to the noise perturbation in latent space.** Probabilistic autoencoders (VAE, SD-VAE, VMAE) demonstrate greater robustness to latent perturbations, as their decoders can reconstruct from a smoother latent space compared to deterministic autoencoders (AE, DAE).

5.1. Results and Analysis

Latent Space Smoothness. As introduced in Sec. 4.1, the latent space should be smooth and continuous rather than consisting of sparse points to ensure stable and consistent score matching. As illustrated in Fig. 3, deterministic encoders (AE, DAE) project inputs into discrete and sparse latent *points*, restricting the decodable space to the discrete isolated points. In contrast, probabilistic encoders (VAE, SD-VAE, VMAE) map inputs to latent distributions, forming a continuous latent space that broadens the decodable region. A continuous *space* is preferable to discrete *points* for diffusion models to learn, as it allows generations within the vicinity of target points rather than requiring exact matches.

To assess the smoothness of the latent space, we compare the robustness of each autoencoder to a noise with varied perturbation ratios. Fig. 6 reports rFID scores of each autoencoder with varied noise levels, measuring the quality of generated images. First, the deterministic autoencoders (AE and DAE) show reasonably good performance when the noise level is extremely low, while the rFID score significantly drops when a larger level of perturbation is applied. This suggests that deterministic encoding approaches would destabilize the diffusion model’s learning process due to their limited support in the latent space. In spite of their completely deterministic design, they still show robustness for a very small noise, probably because the training process causes some level of errors when they map images to the latent space and reconstruct them, and the model learns to robustly map within this small range of noise. Second, VAE shows minimal increase in rFID across perturbation ratios, indicating that its decodable space is most widely supported. Nevertheless, its high absolute rFID scores make it less suitable for LDMs. Lastly, SD-VAE and VMAE demonstrate combined strengths, achieving the best quality (in rFID) while more robust to noise compared to AE and DAE. This




Figure 7. **Perceptual compression evaluation of autoencoders.** (a) object-level and (b) part-level are evaluated with ADE20K and CelebAMask-HQ, respectively. A lower compression degree (σ^2) signifies stronger compression, whereas higher semantic disentanglement (ρ) reflects the amount of semantic information retained after compression. Our VMAE is positioned closest to the ideal autoencoder, indicating the most effective perceptual compression.

indicates that a probabilistic latent space indeed results in a larger decodable space which is more desirable for LDMs. Comparing the two, our VMAE consistently outperforms the strongest baseline, SD-VAE across all noise levels. This conclusion is consistent with the actual generation quality, demonstrated in Tab. 3.

Perceptual Compression. As detailed in Sec. 4.2, perceptual compression consists of two key properties: how much perceptual detail is compressed during encoding and how well semantic information is preserved after encoding. For the former, the more perceptual details are discarded during the compression, the smaller the variance among the features within a same cluster will be. To quantify this, we measure the average *feature variance* within each semantic cluster, denoted by σ^2 , based on the ground truth segmentation map. In the context of LDMs, an effective autoencoder should exhibit a higher degree of compression than the raw image; that is, a smaller σ^2 is desired.

For the latter, on the other hand, if the compression effectively retains the semantics, the semantic clusters should be highly distinguishable (disentangled) from each other in the latent space. Specifically, we first construct an affinity graph of latent features from each image and compute the mean intra-cluster edges (u_{intra}) and mean inter-cluster edges (u_{inter}) based on the ground truth. Then, we quantify the semantic disentanglement by their ratio $\rho = u_{\text{intra}}/u_{\text{inter}}$. In the context of LDMs, it is desired to maintain semantically distinguishable features; that is, $\rho > 1$.

To consider various compression levels, as proposed in Sec. 4.2, we use ADE20K for object-level evaluation and CelebAMask-HQ for part-level evaluation, where the *semantics* refer to the objects (e.g., chair, table) in ADE20K, while the parts in CelebAMask-HQ (e.g., nose, ears), respectively.

As shown in Fig. 7, DAE and VAE exhibit larger σ^2 than

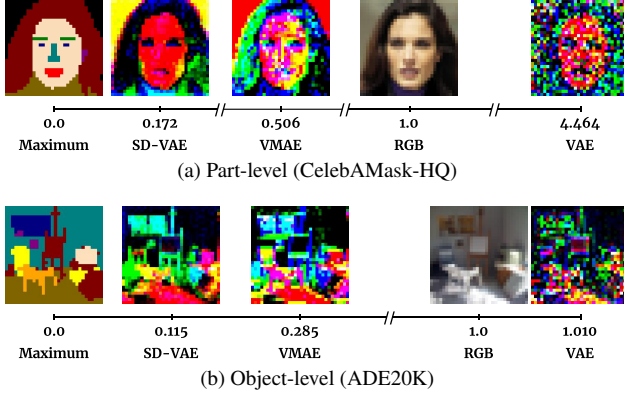


Figure 8. **Illustration of compressed latents of autoencoders, along with relative compression degrees (σ^2).** VMAE achieves a strong compression with diversified features, preserving more detailed semantics, unlike VAE with highly dispersed features or SD-VAE with excessively entangled features.

the raw image, indicating that they are unable to perform compression—a fundamental requirement for perceptual compression, while AE, SD-VAE, and our VMAE successfully achieve it. Among these, SD-VAE shows the strongest compression, indicated by the lowest σ^2 . However, its ρ is lowest among all methods, implying that its compression not just reduces perceptual details but also erases essential semantic information, resulting in excessive feature entanglement. Notably, VMAE most successfully balances and simultaneously satisfies the two criteria, *i.e.*, strong compression capability and superior semantic disentanglement.

For further qualitative analysis, we illustrate the compressed latents with respect to the degree of compression (σ_{intra}) in Fig. 8. As clearly shown in (a) part-level compression, VAE features are highly dispersed, implying that they are embedded closer to a Gaussian distribution rather than a semantically meaningful one. This observation aligns with findings in β -VAE [16]. Also, as the objects in the same class are not distinctly labeled in ADE20K, the extremely low σ^2 of SD-VAE in (b) object-level compression suggests that its latent features are not only entangled within individual objects but also struggle to discriminate different objects. This result strongly suggests that the latent features of SD-VAE are highly entangled, leading to degraded decoding performance (Tab. 1) and making them unsuitable for downstream tasks (Sec. 4.2). The superior performance of our VMAE is evident in both (a) and (b), where its features are diversified yet semantically structured, indicating its *hierarchical* compression.

Reconstruction Quality. We evaluate the reconstruction capability of autoencoders using various metrics on the ImageNet-1K test set. Tab. 1 indicates that AE and DAE excel in pixel-level reconstruction (L1, MSE, PSNR, SSIM), while they exhibit weaker perceptual reconstruction (LPIPS, rFID). VAE performs poorly across all reconstruction met-

Table 1. Reconstruction performance comparison across autoencoders on ImageNet-1K test set.

Model	L1 \downarrow	MSE \downarrow	PSNR \uparrow	SSIM \uparrow	LPIPS \downarrow	rFID \downarrow
AE	0.0218	0.0013	32.18	0.895	0.172	6.21
DAE	0.0237	0.0015	31.31	0.887	0.175	3.97
VAE	0.0281	0.0018	29.40	0.825	0.281	17.41
SD-VAE	0.0223	0.0015	29.85	0.853	<u>0.099</u>	<u>1.89</u>
VMAE (Ours)	<u>0.0221</u>	<u>0.0014</u>	<u>31.52</u>	<u>0.890</u>	0.062	0.89
Ground truth	0.0000	0.0000	∞	1.000	0.000	0.000

Table 2. Model efficiency comparisons. The numbers in parentheses indicate the relative percentage of our method’s efficiency.

Metric	AE, DAE, VAE, SD-VAE	VMAE (Ours)
Model size (MB)	319.7	42.7 (13.4%)
GFLOPS	17,331.3	703.9 (4.1%)
Training time (hr)	24	9 (37.5%)

rics, making it unsuitable for LDMs. This is possibly due to the strong influence of KL divergence loss, which forces the decoder to reconstruct images from sampled latents near a Gaussian distribution, making reconstruction more challenging. In contrast, SD-VAE achieves strong perceptual reconstruction while maintaining comparable pixel fidelity, owing to the additional reconstruction loss, *i.e.*, adversarial loss, which alleviates the strong influence of latent regularization. Our VMAE surpasses SD-VAE in both perceptual detail recovery and pixel-level accuracy, benefiting from its hierarchically structured latent features, *i.e.*, diversified representations that facilitate more precise decoding. In our VMAE, patch drop during the encoding likely weakens the effect of latent regularization.

Model Efficiency. We compare the model cost in Tab. 2. Our VMAE demonstrates significantly faster training and inference, with considerably smaller model size. Faster inference is particularly favorable for training LDMs, as image encoding is required at each iteration.

6. Evaluation on Overall Image Generation

We conduct experiments to evaluate the performance of our LDMAE, an LDM model equipped with the proposed VMAE, for image generation. The datasets and detailed experimental setup is provided in Appendix B.2.

6.1. Results and Analysis

Tab. 3 presents the generation performance of LDMs equipped with different autoencoders. We can analyze these results through the lens of the three desirable properties of autoencoders for LDMs.

AE and DAE, which lack smooth latent spaces, show weaker generation performance than probabilistic models like SD-VAE and VMAE, which benefit from smoother latent representations. Among deterministic models, DAE performs slightly better than AE, indicating that DAE’s latent space is actually more robust [44]. VAE achieves the



Figure 9. **Generated Samples on CelebA-HQ** (256×256). Our LDMAE successfully captures diversity in age, ethnicity, hairstyle, and other characteristics, while also preserving fine details of human faces.

Table 3. Comparison of generation performance among different autoencoders on ImageNet-1K (256×256) and CelebA-HQ (256×256). The best values are highlighted in **bold**, and the second-best values are underlined.

Model	ImageNet-1K					CelebA-HQ	
	gFID \downarrow	sFID \downarrow	IS \uparrow	Prec \uparrow	Rec \uparrow	gFID \downarrow	sFID \downarrow
AE	12.92	12.65	124.0	0.724	0.339	24.80	29.90
DAE	8.60	12.12	160.3	0.797	0.402	21.42	22.90
VAE	34.60	22.32	54.6	0.517	0.415	32.33	32.96
SD-VAE	<u>6.49</u>	<u>5.60</u>	<u>173.3</u>	<u>0.819</u>	<u>0.429</u>	<u>9.00</u>	<u>14.83</u>
VMAE	5.98	5.16	185.5	0.844	0.435	7.61	9.08

smoothest latent space, but its poor reconstruction quality results in the worst generation scores on both datasets. This highlights that latent smoothness alone is insufficient without reliable reconstruction. Notably, the gap between VAE and deterministic models (AE, DAE) is much smaller on CelebA-HQ than on ImageNet-1K. This is likely because the ImageNet-1K dataset is often poorly preprocessed, leaving objects off-center or smaller, with large, noisy backgrounds. Such background noise reduces sensitivity to fine reconstruction errors, especially in perceptual metrics like FID. In contrast, CelebA-HQ images are tightly cropped, with large, centered faces and minimal background, making them more sensitive to small errors in the predicted latent codes, which can be resolved by a sufficiently smooth latent space. SD-VAE balances all three properties well, achieving solid performance on both datasets. VMAE further improves each property and consistently achieves the best scores across all metrics. Moreover, VMAE’s superior pixel-level reconstruction quality, while not directly reflected in FID and IS, likely enhances the perceptual quality of generated images. Figs. 1 and 9 show generated samples from ImageNet-1K and CelebA-HQ, demonstrating high-resolution results with balanced quality across diverse objects and faces.

6.2. Ablation Study

To evaluate the impact of each loss term on reconstruction and generation quality, we conduct an ablation study on ImageNet-1K (256×256). Tab. 4 presents the results across

Table 4. **Ablation study on the effect of each loss term.** We report reconstruction quality using rFID, PSNR, LPIPS, and SSIM and assess generation quality using gFID on ImageNet-1K (256×256).

Losses	PSNR	SSIM	rFID	LPIPS	gFID
Baseline (MSE)	32.18	0.906	6.21	0.172	12.92
+ Masking loss (\mathcal{L}_M)	32.01	0.913	4.66	0.130	8.92
+ Latent regularizer (\mathcal{L}_{reg})	31.14	0.881	1.59	0.112	6.32
+ Perceptual loss (\mathcal{L}_P)	31.52	0.889	0.89	0.062	5.98

five metrics: rFID, PSNR, LPIPS, SSIM, and gFID.

Starting only with the MSE reconstruction loss, same as the standard AE, the baseline focuses solely on pixel-level fidelity. Consequently, it excels in PSNR and SSIM but underperforms in perceptual metrics (rFID, LPIPS). As analyzed in Sec. 5.1, the AE lacks latent smoothness, ultimately leading to high gFID. When the masking loss \mathcal{L}_M is added, the model becomes similar to vanilla MAE. Although this results in some loss of pixel-level reconstruction quality, it yields a more aligned latent space, thereby improving perceptual reconstruction. As a result, the gFID also improves substantially, although latent smoothness remains inadequate. Adding the latent regularizer \mathcal{L}_{reg} loss leads to a slight drop in PSNR and SSIM but confers a substantial gain in perceptual reconstruction quality, which carries over to better generative performance. Finally, adding the perceptual loss \mathcal{L}_P further enhances perceptual reconstruction quality, not only slightly recovering pixel-level reconstruction but also significantly improving rFID and LPIPS, with a modest improvement in gFID.

7. Summary

We propose three key properties for autoencoders in LDMs: latent smoothness, perceptual compression, and high-quality reconstruction. By analyzing existing autoencoders, we introduce Variational Masked AutoEncoders (VMAEs) to address their limitations while balancing these properties. Extensive experiments show that VMAEs enhance generative performance with a stable training process and highly efficient inference. Our findings provide valuable insights for designing effective autoencoders for LDMs.

Acknowledgments

This work was supported by Samsung Electronics (IO240512-09881-01), Youlchon Foundation, NRF grants (RS-2021-NR05515, RS-2024-00336576, RS-2023-0022663) and IITP grants (RS-2022-II220264, RS-2024-00353131) by the government of Korea.

References

- [1] Huiwen Chang, Han Zhang, Lu Jiang, Ce Liu, and William T Freeman. MaskGIT: Masked generative image transformer. In *Proc. of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2022. 3
- [2] Hao Chen, Yujin Han, Fangyi Chen, Xiang Li, Yidong Wang, Jindong Wang, Ze Wang, Zicheng Liu, Difan Zou, and Bhiksha Raj. Masked autoencoders are effective tokenizers for diffusion models. *arXiv:2502.03444*, 2025. 4
- [3] Junyu Chen, Han Cai, Junsong Chen, Enze Xie, Shang Yang, Haotian Tang, Muyang Li, Yao Lu, and Song Han. Deep compression autoencoder for efficient high-resolution diffusion models. *arXiv:2410.10733*, 2024. 3
- [4] Jaewoong Choi, Junho Lee, Changyeon Yoon, Jung Ho Park, Geonho Hwang, and Myungjoo Kang. Do not escape from the manifold: Discovering the local coordinates on the latent space of gans. *arXiv preprint arXiv:2106.06959*, 2021. 4
- [5] Xiaoliang Dai, Ji Hou, Chih-Yao Ma, Sam Tsai, Jialiang Wang, Rui Wang, Peizhao Zhang, Simon Vandenhende, Xiao-fang Wang, Abhimanyu Dubey, et al. Emu: Enhancing image generation models using photogenic needles in a haystack. *arXiv:2309.15807*, 2023. 2, 3
- [6] Mostafa Dehghani, Josip Djolonga, Basil Mustafa, Piotr Padlewski, Jonathan Heek, Justin Gilmer, Andreas Peter Steiner, Mathilde Caron, Robert Geirhos, Ibrahim Alabdulmohsin, et al. Scaling vision transformers to 22 billion parameters. In *Proc. of the International Conference on Machine Learning (ICML)*, 2023. ii
- [7] Jia Deng, Wei Dong, Richard Socher, Li-Jia Li, Kai Li, and Li Fei-Fei. ImageNet: A large-scale hierarchical image database. In *Proc. of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2009. 5, i, ii
- [8] Prafulla Dhariwal and Alexander Nichol. Diffusion models beat gans on image synthesis. In *Advances in Neural Information Processing Systems (NeurIPS)*, 2021. 1, 3
- [9] Alexey Dosovitskiy. An image is worth 16x16 words: Transformers for image recognition at scale. *arXiv:2010.11929*, 2020. 3, i
- [10] Patrick Esser, Robin Rombach, and Bjorn Ommer. Taming transformers for high-resolution image synthesis. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 12873–12883, 2021. 2, 3, i, ii
- [11] Patrick Esser, Sumith Kulal, Andreas Blattmann, Rahim Entezari, Jonas Müller, Harry Saini, Yam Levi, Dominik Lorenz, Axel Sauer, Frederic Boesel, et al. Scaling rectified flow transformers for high-resolution image synthesis. In *Proc. of the International Conference on Machine Learning (ICML)*, 2024. 1, 3, i
- [12] Jaehoon Hahm, Junho Lee, Sunghyun Kim, and Joonseok Lee. Isometric representation learning for disentangled latent space of diffusion models. *arXiv preprint arXiv:2407.11451*, 2024. 4
- [13] Kaiming He, Xinlei Chen, Saining Xie, Yanghao Li, Piotr Dollár, and Ross Girshick. Masked autoencoders are scalable vision learners. In *Proc. of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2022. 3, i
- [14] Amir Hertz, Kfir Aberman, and Daniel Cohen-Or. Delta denoising score. In *Proc. of the IEEE/CVF Conference on International Conference on Computer Vision (ICCV)*, 2023. 4
- [15] Martin Heusel, Hubert Ramsauer, Thomas Unterthiner, Bernhard Nessler, and Sepp Hochreiter. Gans trained by a two time-scale update rule converge to a local nash equilibrium. *Advances in Neural Information Processing Systems (NIPS)*, 2017. ii
- [16] Irina Higgins, Loic Matthey, Arka Pal, Christopher Burgess, Xavier Glorot, Matthew Botvinick, Shakir Mohamed, and Alexander Lerchner. β -VAE: Learning basic visual concepts with a constrained variational framework. In *Proc. of the International Conference on Learning Representations (ICLR)*, 2017. 7
- [17] Jonathan Ho, Ajay Jain, and Pieter Abbeel. Denoising diffusion probabilistic models. In *Advances in Neural Information Processing Systems (NeurIPS)*, 2020. 1, 3
- [18] Tero Karras, Timo Aila, Samuli Laine, and Jaakko Lehtinen. Progressive growing of GANs for improved quality, stability, and variation. In *Proc. of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2018. ii
- [19] Tero Karras, Miika Aittala, Timo Aila, and Samuli Laine. Elucidating the design space of diffusion-based generative models, 2022. 1
- [20] Diederik P Kingma. Auto-encoding variational bayes. *arXiv:1312.6114*, 2013. 1, 3, i
- [21] Tuomas Kynkäänniemi, Tero Karras, Samuli Laine, Jaakko Lehtinen, and Timo Aila. Improved precision and recall metric for assessing generative models. *Advances in Neural Information Processing Systems (NeurIPS)*, 2019. ii
- [22] Black Forest Labs. Flux. <https://github.com/black-forest-labs/flux>, 2023. 2, 3
- [23] Cheng-Han Lee, Ziwei Liu, Lingyun Wu, and Ping Luo. MaskGAN: Towards diverse and interactive facial image manipulation. In *Proc. of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2020. i
- [24] Tianhong Li, Huiwen Chang, Shlok Mishra, Han Zhang, Dina Katabi, and Dilip Krishnan. MAGE: Masked generative encoder to unify representation learning and image synthesis. In *Proc. of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2023. 3
- [25] Tianhong Li, Yonglong Tian, He Li, Mingyang Deng, and Kaiming He. Autoregressive image generation without vector quantization. *Advances in Neural Information Processing Systems (NeurIPS)*, 2024. 3
- [26] Cheng Lu, Yuhao Zhou, Fan Bao, Jianfei Chen, Chongxuan Li, and Jun Zhu. Dpm-solver: A fast ode solver for diffusion probabilistic model sampling in around 10 steps. In *Advances*

- in *Neural Information Processing Systems (NeurIPS)*, 2022. [1](#)
- [27] Hyelin Nam, Gihyun Kwon, Geon Yeong Park, and Jong Chul Ye. Contrastive denoising score for text-guided latent diffusion image editing. In *Proc. of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2024. [4](#)
- [28] Charlie Nash, Jacob Menick, Sander Dieleman, and Peter W Battaglia. Generating images with sparse representations. *arXiv:2103.03841*, 2021. [ii](#)
- [29] Alexander Quinn Nichol and Prafulla Dhariwal. Improved denoising diffusion probabilistic models. In *Proc. of the International Conference on Machine Learning (ICML)*, 2021. [1](#)
- [30] Namuk Park, Wonjae Kim, Byeongho Heo, Taekyung Kim, and Sangdoo Yun. What do self-supervised vision transformers learn? *arXiv:2305.00729*, 2023. [4](#)
- [31] William Peebles and Saining Xie. Scalable diffusion models with transformers. In *Proc. of the IEEE/CVF Conference on International Conference on Computer Vision (ICCV)*, 2023. [1](#), [3](#), [ii](#)
- [32] Robin Rombach, Andreas Blattmann, Dominik Lorenz, Patrick Esser, and Björn Ommer. High-resolution image synthesis with latent diffusion models. In *Proc. of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2022. [1](#), [3](#), [4](#), [i](#)
- [33] D. E. Rumelhart, G. E. Hinton, and R. J. Williams. *Learning internal representations by error propagation*, page 318–362. MIT Press, 1986. [3](#), [i](#)
- [34] Seyedmorteza Sadat, Jakob Buhmann, Derek Bradley, Otmar Hilliges, and Romann M Weber. LiteVAE: Lightweight and efficient variational autoencoders for latent diffusion models. *Advances in Neural Information Processing Systems (NeurIPS)*, 2024. [2](#)
- [35] Tim Salimans, Ian Goodfellow, Wojciech Zaremba, Vicki Cheung, Alec Radford, Xi Chen, and Xi Chen. Improved techniques for training GANs. In *Advances in Neural Information Processing Systems (NIPS)*. Curran Associates, Inc., 2016. [ii](#)
- [36] Jeongwoo Shin, Inseo Lee, Junho Lee, and Joonseok Lee. Self-guided masked autoencoder. In *Advances in Neural Information Processing Systems (NeurIPS)*, 2024. [2](#), [4](#), [5](#)
- [37] Zitao Shuai, Chenwei Wu, Zhengxu Tang, Bowen Song, and Liyue Shen. Latent space disentanglement in diffusion transformers enables precise zero-shot semantic editing. *arXiv:2411.08196*, 2024. [4](#)
- [38] Karen Simonyan and Andrew Zisserman. Very deep convolutional networks for large-scale image recognition. *arXiv:1409.1556*, 2014. [5](#)
- [39] Jiaming Song, Chenlin Meng, and Stefano Ermon. Denoising diffusion implicit models. In *Proc. of the International Conference on Learning Representations (ICLR)*, 2021. [1](#), [3](#)
- [40] Yang Song and Stefano Ermon. Generative modeling by estimating gradients of the data distribution. In *Advances in Neural Information Processing Systems (NeurIPS)*, 2019. [4](#)
- [41] Yang Song, Jascha Sohl-Dickstein, Diederik P. Kingma, Abhishek Kumar, Stefano Ermon, and Ben Poole. Score-based generative modeling through stochastic differential equations. In *Proc. of the International Conference on Learning Representations (ICLR)*, 2021. [1](#), [3](#), [4](#)
- [42] Christian Szegedy, Vincent Vanhoucke, Sergey Ioffe, Jon Shlens, and Zbigniew Wojna. Rethinking the inception architecture for computer vision. In *Proc. of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2016. [ii](#)
- [43] Aaron Van Den Oord, Oriol Vinyals, et al. Neural discrete representation learning. In *Advances in Neural Information Processing Systems (NIPS)*, 2017. [3](#)
- [44] Pascal Vincent, Hugo Larochelle, Yoshua Bengio, and Pierre-Antoine Manzagol. Extracting and composing robust features with denoising autoencoders. In *Proc. of the International Conference on Machine Learning (ICML)*, 2008. [3](#), [7](#), [i](#)
- [45] Jingfeng Yao and Xinggang Wang. Reconstruction vs. generation: Taming optimization dilemma in latent diffusion models. *arXiv:2501.01423*, 2025. [2](#), [3](#), [ii](#)
- [46] Richard Zhang, Phillip Isola, Alexei A Efros, Eli Shechtman, and Oliver Wang. The unreasonable effectiveness of deep features as a perceptual metric. In *Proc. of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2018. [5](#), [i](#)
- [47] Bolei Zhou, Hang Zhao, Xavier Puig, Tete Xiao, Sanja Fidler, Adela Barriuso, and Antonio Torralba. Semantic understanding of scenes through the ade20k dataset. *International Journal of Computer Vision*, 127:302–321, 2019. [i](#)