arXiv:2512.16636v1 [cs.CV] 18 Dec 2025

REGLUE Your Latents with Global and Local Semantics for Entangled Diffusion

Giorgos Petsangourakis[1,2] Christos Sgouropoulos[1] Bill Psomas[3]
Theodoros Giannakopoulos[1] Giorgos Sfikas[2] Ioannis Kakogeorgiou[1]

[1]IIT, National Centre for Scientific Research "Demokritos"
[2]University of West Attica [3]VRG, FEE, Czech Technical University in Prague

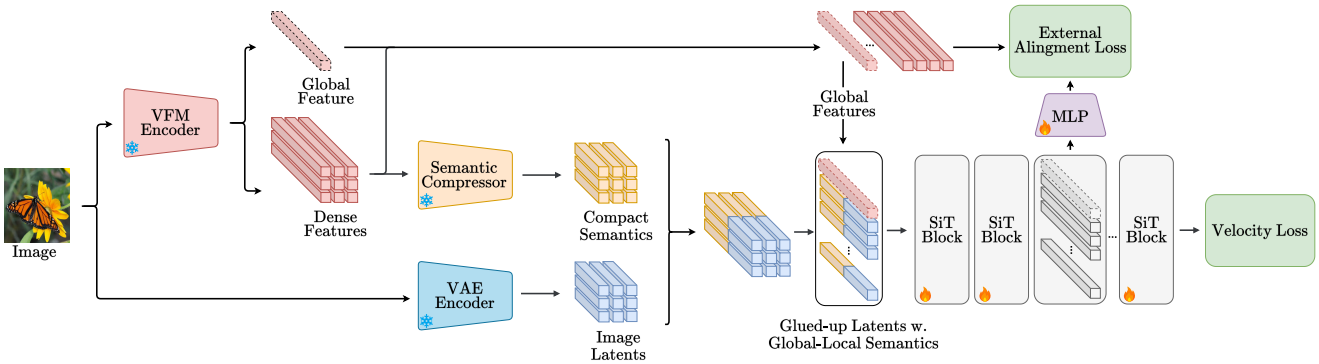

Figure 1. **Overview of** REGLUE, Representation Entanglement with Global–Local Unified Encoding. The encoder of a vision foundation model (VFM) provides (i) *local* patch-level features and (ii) a *global* image-level feature (*i.e.* [CLS]). A lightweight semantic compressor (pre-trained offline) maps the patch features to *compact spatial semantics*. In parallel, a frozen VAE encoder produces *image latents*. We concatenate the latents, compressed semantics, and the global token and feed them to a SiT backbone [44], which *jointly models* all modalities with a velocity objective. Diffusion noise injection is omitted from the illustration. At a selected block, an MLP head applies an external alignment [73] to match hidden SiT features to clean VFM targets. At sampling, we decode the (jointly) generated VAE latents.

Abstract

Latent diffusion models (LDMs) achieve state-of-the-art image synthesis, yet their reconstruction-style denoising objective provides only indirect semantic supervision: high-level semantics emerge slowly, requiring longer training and limiting sample quality. Recent works inject semantics from Vision Foundation Models (VFMs) either externally via representation alignment or internally by jointly modeling only a narrow slice of VFM features inside the diffusion process, under-utilizing the rich, nonlinear, multi-layer spatial semantics available.

We introduce REGLUE (Representation Entanglement with Global–Local Unified Encoding), a unified latent diffusion framework that jointly models (i) VAE image latents, (ii) compact local (patch-level) VFM semantics, and (iii) a global (image-level) [CLS] token within a single SiT backbone. A lightweight convolutional semantic compressor nonlinearly aggregates multi-layer VFM features into a low-dimensional, spatially structured representation, which is entangled with the VAE latents in the diffusion process. An external alignment loss further regularizes internal representations toward frozen VFM targets. On ImageNet 256×256, REGLUE consistently improves FID and accelerates convergence over SiT-B/2 and SiT-XL/2 baselines, as well as over REPA, ReDi, and REG. Extensive experiments show that (a) spatial VFM semantics are crucial, (b) non-linear compression is key to unlocking their full benefit, and (c) global tokens and external alignment act as complementary, lightweight enhancements within our global–local–latent joint modeling framework. The code is available at https://github.com/giorgospets/reglue.

1. Introduction

Latent Diffusion Models (LDMs) have become a dominant choice for high-quality image synthesis, largely because they shift the generative paradigm to modeling a compact VAE latent space [56]. However, training such models is *challenging*: under a single denoising objective, the model

must concurrently learn high-level semantics (*what* to generate, *e.g*. objects, layout, relations) and low-level visual details (*how* to generate it, *e.g*. fine-grained appearance) [71]. The reconstruction-style denoising loss provides semantic supervision only *indirectly*, so semantic structure emerges slowly, limiting image quality and convergence speed [73].

To address these challenges, recent works leverages semantic representations from strong, pretrained Vision Foundation Models (VFMs), accelerating convergence and improving image quality [33, 67, 73]. `REPA` [73] proposes a representation *alignment* objective to distill VFM features as an external teacher to the diffusion model. `REG` [67] and `ReDi` [33] go a step further: they *jointly model* the image latent and the semantic signal inside the diffusion process. Regarding this signal, `REG` uses a single *image-level* (global) representation (*i.e*., `[CLS]`), while `ReDi` employs *linearly* PCA-projected *patch-level* (local) VFM features.

Both `REG` and `ReDi` expose only a *narrow*, *low-capacity* semantic slice of the VFM to the diffusion model, underutilizing the rich, nonlinear, multi-layer, and spatial semantics available. `REG`'s `[CLS]` offers informative image-level guidance, but is inherently non-spatial; fine-grained semantics are recovered via an external alignment loss as in `REPA` [73], which distills spatial semantics and boosts generative performance. In contrast, `ReDi` explicitly models patch-level semantics but its linear PCA projection restricts the representations to a low-dimensional linear subspace, limiting the richness and non-linear spatial information.

We argue that a critical design choice for effectively leveraging VFM features is to model spatial semantics within the diffusion model, while preserving their nonlinear, multi-layer information via compact learned representations produced by a *semantic compressor*. Specifically, jointly modeling patch-level features with VAE latents provides spatial guidance critical for capturing fine-grained structure and yields larger gains than either external feature alignment alone (`REPA`) or jointly modeling a single image-level token (`REG`). These signals remain complementary: an alignment loss and a global `[CLS]` token can be added as orthogonal auxiliary signals, but the primary improvements stem from spatial joint modeling of compressed patch features within the diffusion model (see Table 1).

In our work, we first train a lightweight convolutional *semantic compressor* that maps nonlinear, multi-layer VFM features into a compact, spatially structured, semantics-preserving representation. Then, we introduce a unified diffusion modeling approach that *jointly* models: (i) these compact patch-level (local) semantic features, (ii) an image-level (global) representation, and (iii) the VAE image latents. During training, we also apply a feature-alignment auxiliary loss that aligns diffusion internal features with VFM teacher representations, further improving image synthesis performance. The overview is shown in Figure 1.

In summary, our contributions are:

1. We introduce `REGLUE` (*Representation Entanglement with Global–Local Unified Encoding*), a unified diffusion framework that jointly models image-level (global) and patch-level (local) VFM semantics with VAE latents, significantly boosting generative performance.
2. We propose a lightweight *semantic compressor* that aggregates multi-layer VFM features and maps them to a compact, semantics-preserving space. This compact representation enables significant gains via patch-level joint modeling with VAE latents, strongly improving synthesized image quality.
3. We show that, in `REGLUE`, (a) patch-level (local) semantics, (b) image-level (global) semantics, and (c) REPA-style representation alignment act synergistically, delivering substantial gains in image quality and training convergence, while keeping the diffusion model's parameters and inference-time compute essentially unchanged.
4. On ImageNet 256×256 generation benchmark, `SiT-XL/2+REGLUE` reaches the 1M-step performance of `ReDi` and `REG` using less than 30% and 80% of their iterations, respectively.

2. Related Work

Latent-variable generative modeling. Latent-variable models like Variational Autoencoders and Diffusion Denoising Probabilistic models are core building blocks of modern generative pipelines [32, 56]. There are often underappreciated connections across these families [6, 51]: diffusion can be viewed as a hierarchical VAE with a fixed latent dimension and a frozen encoder [43], and both are trained via variational approximations [5]. Normalizing flows [30, 76] have likewise been analyzed through their links to diffusion [1, 74, 78]. VAEs are also closely related to probabilistic PCA, replacing the linear projection with a neural decoder [20, 64]. Analyzing models under a unified framework has repeatedly yielded progress: Scalable Interpolant Transformers (`SiT`) [44], alongside DiT [50] and Lightning DiT [71], an indispensable component of state-of-the-art generative frameworks [33, 59, 67, 71, 73]. Our work espouses an analogous rationale, focusing on a holistic, joint modeling of tokenizer (VAE latents) and VFM semantics within a single framework.

Representation alignment with VFM features. Latent diffusion typically uses a VAE first stage to obtain compact image latents, which are optimized for reconstruction rather than semantics; this can slow semantic emergence and curb the ability to represent abstract concepts [2, 36]. To mitigate this, works enhance the *first stage* by aligning or reshaping the latent space: VA-VAE [71] adds a VFM alignment loss, TexTok [75] injects text description embeddings, MAETok [9] argues for discriminative (non-variational) latents, and FA-VAE [45] separates low/high

frequencies via wavelets. Complementarily, the *second stage* (the denoiser) can be aligned to VFMs: REPA [73] distills mid-block features to VFM targets and accelerates convergence; DDT [66] decouples encoder/decoder; REPA-E [37] pursues end-to-end VAE+denoiser alignment; and SVG [59] focuses on few-step generation. Our approach is complementary: instead of exposing a narrow semantic slice or relying solely on external representation alignment, we *jointly model* global and local VFM semantics with VAE latents via a frozen, lightweight semantic compressor, entangling them within a single diffusion backbone and thereby accelerating convergence and improving generative quality.

Joint feature generative modeling. A complementary line of work directly *models* foundation features as observed variables, conceptually related to multimodal diffusion that learns joint spaces across modalities (*e.g.* text–image–video–audio) rather than distilling from them. Examples include CoDi [62] for any-to-any generation across modalities, and video methods that entangle appearance with motion signals like VideoJam [8]. Closer to our setting, `REG` [67] proposes a SiT-based model that incorporates compressed tokens with the addition of also modeling a VFM `[CLS]` token; Representation Diffusion (ReDI) [33] takes this approach one step further and models spatially referenced high-level features on equal grounds to VAE latents by a Diffusion Transformer. We argue that these prior works model VFM in a suboptimal manner, either by discarding useful spatial cues [67] or working with constrained, linear projections of high-level semantics.

Representation learning. Supervised convolutional networks [23] and, more recently, Vision Transformers (ViTs) [18] have established strong baselines for transferable visual features, but the field has largely shifted toward pretraining regimes that reduce or remove manual labels. Early self-supervised work relies on hand-crafted pretext tasks (*e.g.* patch permutation [47] or rotation prediction [21]), while modern approaches favor contrastive objectives [11, 48] and self-distillation [7, 22], yielding highly transferable features at scale. Transformer-era enables masked image modeling (MIM): from BEiT [4] and MAE [24] to hybrid variants like iBOT [81] and AttMask [27]. In parallel, vision–language pretraining learns joint embeddings from web-scale image–text pairs [58]. CLIP [54] popularizes this paradigm by aligning images and captions with a contrastive objective, yielding strong zero-shot recognition [16], retrieval [31, 53], and segmentation [61] performance. SigLIP [77] refines the recipe by replacing the softmax with independent sigmoid losses, and SigLIP-2 [65] further improves transfer via stronger data and training strategies.

3. Method

3.1. Preliminaries

Scalable Interpolant Transformer (SiT). We adopt the `SiT` [44] framework, built on the stochastic–interpolant formulation [40, 44]. Given a clean image $\mathbf{x}_*$ and a pretrained VAE encoder $\mathcal{E}_z$ that produces image latents $\mathbf{z}_* \in \mathbb{R}^{D_z \times H_z \times W_z}$, we consider the continuous-time stochastic interpolant:

$$\mathbf{z}_t \;=\; \alpha_t \mathbf{z}_* + \sigma_t \boldsymbol{\epsilon}_z, \qquad \boldsymbol{\epsilon}_z \sim \mathcal{N}(\mathbf{0}, \mathbf{I}), \;\; t \in [0,1], \quad (1)$$

where $\alpha_0 = \sigma_1 = 1$ and $\alpha_1 = \sigma_0 = 0$, with α_t decreasing and σ_t increasing in t. `SiT` adopts a Transformer-based architecture with $\mathcal{K}$ stacked blocks, parameterizes the velocity field $\mathbf{v}_\theta(\mathbf{z}_t, t)$ and is trained with the standard velocity objective:

$$\mathbb{E}_{\mathbf{z}_*, \boldsymbol{\epsilon}_z, t}\Big[\big\| \mathbf{v}_\theta(\mathbf{z}_t, t) \;-\; \dot{\alpha}_t \mathbf{z}_* \;-\; \dot{\sigma}_t \boldsymbol{\epsilon}_z \big\|^2 \Big]. \quad (2)$$

Unless stated otherwise, we adopt the linear schedule $\alpha_t = 1 - t$ and $\sigma_t = t$, which yields constant derivatives $\dot{\alpha}_t = -1$ and $\dot{\sigma}_t = 1$.

Vision Foundation Model (VFM). We denote by $\mathcal{E}_v(\cdot)$ a pretrained VFM (*e.g.* DINOv2), which provides *patch-level* semantic features $\mathbf{f}_*^{(\ell)} \in \mathbb{R}^{D_f \times H_f \times W_f}$ for each layer $\ell \in \{1, 2, \ldots, \mathcal{L}\}$, and a global *image-level* representation $\mathbf{cls}_* \in \mathbb{R}^{D_f}$. $H_f \times W_f$ denotes the patch grid size and D_f the feature dimensionality. In ViT-based encoders H_f, W_f, D_f remain fixed across layers.

3.2. Global–local representation entanglement

Our aim is to *jointly model* (i) VAE latents, (ii) local semantics (patch-level VFM features), and (iii) global semantics (*i.e.* image-level `[CLS]` token) within a single `SiT` model. We reuse the notation from Sec. 3.1. Given an image $\mathbf{x}_*$, $\mathcal{E}_z(\mathbf{x}_*) = \mathbf{z}_* \in \mathbb{R}^{D_z \times H_z \times W_z}$ denotes VAE latents, $\mathbf{cls}_* \in \mathbb{R}^{D_f}$ the global VFM token, and $\mathbf{s}_* \in \mathbb{R}^{D_s \times H_z \times W_z}$ patch-level *compressed* VFM features derived from $\{\mathbf{f}_*^{(\ell)}\}_{\ell=1}^{\mathcal{L}}$ and aligned to the VAE latents' grid (cf. Sec. 3.3 for a detailed description of the compressor).

Forward process. We first adopt a shared schedule (α_t, σ_t) to inject noise for all three entangled modalities:

$$\begin{aligned} \mathbf{z}_t &= \alpha_t \mathbf{z}_* + \sigma_t \boldsymbol{\epsilon}_z, & \boldsymbol{\epsilon}_z &\sim \mathcal{N}(\mathbf{0}, \mathbf{I}), \\ \mathbf{s}_t &= \alpha_t \mathbf{s}_* + \sigma_t \boldsymbol{\epsilon}_s, & \boldsymbol{\epsilon}_s &\sim \mathcal{N}(\mathbf{0}, \mathbf{I}), \\ \mathbf{cls}_t &= \alpha_t \mathbf{cls}_* + \sigma_t \boldsymbol{\epsilon}_{\text{cls}}, & \boldsymbol{\epsilon}_{\text{cls}} &\sim \mathcal{N}(\mathbf{0}, \mathbf{I}), \end{aligned} \quad (3)$$

with independent noise terms and $t \in [0, 1]$.

Velocity objective. SiT parameterizes a velocity field $\mathbf{v}_\theta(\mathbf{z}_t, \mathbf{s}_t, \mathbf{cls}_t, t)$ over the joint state $(\mathbf{z}_t, \mathbf{s}_t, \mathbf{cls}_t)$ and is trained with the multimodal velocity loss:

$$\begin{aligned}\mathcal{L}_{\mathrm{v}} = \mathbb{E}\Big[& \left\|\mathbf{v}_\theta^z(\mathbf{z}_t, \mathbf{s}_t, \mathbf{cls}_t, t) - \dot{\alpha}_t \mathbf{z}_* - \dot{\sigma}_t \boldsymbol{\epsilon}_z\right\|_2^2 \\ & + \lambda_s \left\|\mathbf{v}_\theta^s(\mathbf{z}_t, \mathbf{s}_t, \mathbf{cls}_t, t) - \dot{\alpha}_t \mathbf{s}_* - \dot{\sigma}_t \boldsymbol{\epsilon}_s\right\|_2^2 \\ & + \lambda_{\mathrm{cls}} \left\|\mathbf{v}_\theta^{cls}(\mathbf{z}_t, \mathbf{s}_t, \mathbf{cls}_t, t) - \dot{\alpha}_t \mathbf{cls}_* - \dot{\sigma}_t \boldsymbol{\epsilon}_{\mathrm{cls}}\right\|_2^2\Big],\end{aligned} \tag{4}$$

where $\mathbf{v}_\theta^z$, $\mathbf{v}_\theta^s$ and $\mathbf{v}_\theta^{cls}$ correspond to the predictions for the VAE latent, global VFM token, and patch-level VFM features velocity; λ_s and λ_{cls} are weighting coefficients.

Tokenization and fusion. The SiT diffusion backbone operates on a sequence of tokens with a shared width D, so we aim to bring the different modalities to a common dimensional space. Concretely, we first *patchify*[1] $\mathbf{z}_t$ and $\mathbf{s}_t$, forming $\mathbf{z}'_t = \mathrm{patch}(\mathbf{z}_t) \in \mathbb{R}^{N \times D_{z'}}$ and $\mathbf{s}'_t = \mathrm{patch}(\mathbf{s}_t) \in \mathbb{R}^{N \times D_{s'}}$ (where $N{=}H_z W_z / p^2$ is the number of $p \times p$ patches). We then project each modality to the model width D with linear embedding layers:

$$\tilde{\mathbf{z}}_t = \mathbf{z}'_t \mathbf{W}_z, \quad \tilde{\mathbf{s}}_t = \mathbf{s}'_t \mathbf{W}_s, \quad \tilde{\mathbf{cls}}_t = \mathbf{cls}_t \mathbf{W}_{\mathrm{cls}}$$

where $\mathbf{W}_z \in \mathbb{R}^{D_{z'} \times D}$, $\mathbf{W}_s \in \mathbb{R}^{D_{s'} \times D}$, and $\mathbf{W}_{\mathrm{cls}} \in \mathbb{R}^{D_f \times D}$ are learned embedding matrices.

To combine and jointly model $\tilde{\mathbf{z}}_t$ (VAE latents) and $\tilde{\mathbf{s}}_t$ (local semantics), there are two straightforward options: (i) concatenate image latents and semantic features along the *sequence* dimension and pass $2N$ patch tokens through SiT, or (ii) *merge channel-wise* and keep a single grid of N tokens. We adopt (ii), avoiding the $2\times$ longer sequence and quadratic self-attention overhead of (i). The global [CLS] is inherently a single token; thus, we always keep it as a separate token, adding a negligible throughput overhead. Finally, the input sequence to the SiT Transformer is:

$$\mathbf{h}_t^0 = \Big[\underbrace{\tilde{\mathbf{cls}}_t}_{1\times D}\,;\, \underbrace{(\tilde{\mathbf{z}}_t + \tilde{\mathbf{s}}_t)}_{N\times D}\Big] \in \mathbb{R}^{(1+N)\times D}, \tag{5}$$

where "[;]" denotes concatenation along token dimension.

Prediction heads. Let $\mathbf{h}_t^{\mathcal{K}} \in \mathbb{R}^{(1+N)\times D}$ denote the hidden sequence after the last ($\mathcal{K}$-th) SiT block. We obtain the per-modality velocity predictions as:

$$\begin{aligned}\mathbf{v}_\theta^z &= \mathrm{unpatch}\big(\mathbf{h}_t^{\mathcal{K}}[1{:}N]\,\mathbf{W}_{\mathrm{dec}}^z\big), \\ \mathbf{v}_\theta^s &= \mathrm{unpatch}\big(\mathbf{h}_t^{\mathcal{K}}[1{:}N]\,\mathbf{W}_{\mathrm{dec}}^s\big), \\ \mathbf{v}_\theta^{\mathrm{cls}} &= \mathbf{h}_t^{\mathcal{K}}[0]\,\mathbf{W}_{\mathrm{dec}}^{\mathrm{cls}}.\end{aligned} \tag{6}$$

[1]Partitioning a tensor $x \in \mathbb{R}^{C\times H\times W}$ into non-overlapping $p \times p$ tiles (here $p{=}2$; denoted SiT/2), flattening each tile, and stacking them row-major into a sequence of patch tokens, resulting in $x' \in \mathbb{R}^{(HW/p^2)\times(Cp^2)}$.

where $\mathbf{W}_{\mathrm{dec}}^z \in \mathbb{R}^{D\times D_{z'}}$, $\mathbf{W}_{\mathrm{dec}}^s \in \mathbb{R}^{D\times D_{s'}}$ are linear prediction heads, $\mathrm{unpatch}(\cdot)$ reshapes the N patch tokens back to spatial tensors of shape $D_z \times H_z \times W_z$ and $D_s \times H_z \times W_z$, respectively; $\mathbf{W}_{\mathrm{dec}}^{\mathrm{cls}} \in \mathbb{R}^{D\times D_f}$ projects the global token.

External representation alignment. At a selected Transformer block $k \in \{1, \ldots, \mathcal{K}\}$, we encourage SiT hidden tokens to stay close to *clean*, frozen VFM targets via a lightweight projector $\phi : \mathbb{R}^D \to \mathbb{R}^{D_f}$ and a cosine loss, following [73]. Let $\mathbf{h}_t^k \in \mathbb{R}^{(1+N)\times D}$ be the hidden sequence at block k. We form the target token sequence by concatenating the global token with patch-level VFM features:

$$\mathbf{y}_* \;=\; \big[\mathbf{cls}_*\,;\,\tilde{\mathbf{f}}_*^{(\mathcal{L})}\big] \in \mathbb{R}^{(1+N)\times D_f}. \tag{7}$$

where $\tilde{\mathbf{f}}_*^{(\mathcal{L})}$ denotes flattened spatial dimension to tokens.

The representation alignment loss is

$$\mathcal{L}_{\mathrm{REPA}} = -\,\mathbb{E}\left[\frac{1}{N+1}\sum_{n=1}^{N+1} \mathrm{sim}\Big(\mathbf{y}_*^{[n]},\ \phi\big(\mathbf{h}_t^k\big)^{[n]}\Big)\right], \tag{8}$$

where we apply alignment on the [CLS] position and the N semantic patch tokens and sim is cosine similarity.

Total objective. We train with the multimodal velocity loss in Eq. (4) and the auxiliary alignment loss in Eq. (8):

$$\mathcal{L}_{\mathrm{total}} \;=\; \mathcal{L}_{\mathrm{v}} \;+\; \lambda_{\mathrm{rep}}\,\mathcal{L}_{\mathrm{REPA}}. \tag{9}$$

Sampling. To generate new samples, we employ the reverse-time SDE Euler–Maruyama [60] using $\mathbf{v}_\theta$ to obtain $(\mathbf{z}_0, \mathbf{s}_0, \mathbf{cls}_0)$ from Gaussian noise, and reconstruct the image via the (frozen) VAE decoder.

3.3. A lightweight spatial semantic compressor

Our goal is to jointly model VAE latents with VFM semantics. Naïvely fusing patch-level VFM features with image latents substantially widens the dimensionality to $D_z{+}D_f$ with $D_f \gg D_z$, biasing SiT capacity toward the representation modality and hurting generative quality (a phenomenon also addressed in ReDi [33] via PCA).

Convolutional autoencoder. Instead of a linear subspace, we introduce a *nonlinear*, lightweight *semantic compressor* $\mathcal{E}_\psi$ that preserves spatial structure while re-balancing dimensionality. We instantiate $\mathcal{E}_\psi$ as a shallow convolutional autoencoder, pretrain it once to reconstruct VFM features and then keep it frozen. Our compressor finally projects compressed features of dimension $D_s \ll \sum_\ell D_\ell$.

Multi-layer aggregation. Since leveraging representations across depths has shown benefits to many scene understanding tasks [10, 12, 39, 41, 55, 79], we typically aggregate the multi-layer patch-level VFM features by channel-wise concatenation:

$$\mathbf{f}_* = [\mathbf{f}_*^{(1)}, \mathbf{f}_*^{(2)}, \dots, \mathbf{f}_*^{(\mathcal{L})}] \in \mathbb{R}^{(\mathcal{L}\cdot D_f)\times H_f \times W_f}, \qquad (10)$$

where "[,]" denotes concatenation along the channels.

Training and inference. A lightweight convolutional autoencoder $(\mathcal{E}_\psi, \mathcal{D}_\psi)$ is trained (offline) to reconstruct $\mathbf{f}_*$:

$$\min_{\psi} \mathbb{E}\Big[\|\mathcal{D}_\psi(\mathcal{E}_\psi(\mathbf{f}_*)) - \mathbf{f}_*\|_2^2 \Big], \qquad (11)$$

We then *freeze* $\mathcal{E}_\psi$, produce compact spatial semantics $\mathcal{E}_\psi(\mathbf{f}_*) \in \mathbb{R}^{D_s \times H_f \times W_f}$, and spatially resample them to the VAE latent grid $(H_z \times W_z)$ resulting in $\mathbf{s}_* \in \mathbb{R}^{D_s \times H_z \times W_z}$, typically using bilinear resampling.

4. Experiments

4.1. Setup

Implementation details. We strictly follow the standard training protocols of `SiT` [44]. Our experiments are conducted on the ImageNet [15] dataset. Following the ADM preprocessing pipeline [17], all images are center-cropped and resized to 256×256 resolution. Each image is then encoded into a latent representation $\mathbf{z}_* \in \mathbb{R}^{4\times 32\times 32}$ using the pre-trained `SD-VAE-FT-EMA` [56]. Our main experiments are based on `SiT-B/2` models, which use a 2×2 patch size and are trained for 400K steps. To assess the impact of our approach at larger scales and longer training, we additionally train `SiT-XL/2` models for 1M steps. We maintain a batch size of 256 for all experiments.

Unless stated otherwise, for semantic feature extraction, we employ `DINOv2-B` [14, 49], concatenating features from blocks 9-12 to obtain a 3072-channel (768×4) feature map at 16×16 spatial resolution. Our lightweight convolutional autoencoder $(\mathcal{E}\psi, \mathcal{D}\psi)$, with a hidden layer of 256, is pretrained for 25 epochs using a reconstruction loss (Eq. 11) to compress these maps. The encoder $\mathcal{E}\psi$ produces a compact 16-channel, 16×16 latent representation.

For external representation alignment we follow the formulation in Sec. 3. We apply the external alignment using the local and global representations derived from the VFMs last layer, with layer k of the `SiT` backbone. We use k = 4 for `SiT-B/2` and k = 8 for `SiT-XL/2`. Regarding the total objective coefficients we set $\lambda_s = 1$, $\lambda_{cls} = 0.03$, and $\lambda_{rep} = 0.5$. More implementation details regarding the semantic compressor and `SiT` are provided in Appendix Sec. B.1 and Sec. B.2.

Evaluation. In order to evaluate image generation quality, we report a standard set of quantitative metrics. These include Fréchet Inception Distance (FID) [25] for perceptual quality, sFID [46] for spatial coherence, Inception Score (IS) [57] for diversity, as well as Precision (Pre.) and Recall (Rec.) [34] to measure sample fidelity and distribution coverage, respectively. All metrics are computed using 50,000 generated samples, following the standard ADM evaluation suite [17]. For all experiments, we use Euler–Maruyama SDE sampling with 250 steps. When using Classifier-Free Guidance (CFG) [26], we set CFG scale at $w = 2.8$ and guidance interval to $[0, 0.9]$, following [35].

4.2. Analyzing VFM semantics for generation

We leverage our framework to investigate how diffusion modeling in `SiT-B/2` benefits from: (i) *which* semantics are modeled (global vs. local), (ii) *how* local semantics are compressed (linear vs. non-linear), and (iii) the *degree to which* an external representation-alignment objective contributes to generation quality. The corresponding design choices and their impact on FID are shown in Table 1.

Local (patch-level) outperform global (image-level) semantics. We first focus on representations modeled *directly* by the diffusion model, without any external alignment. We find that patch-level semantics clearly outperform global-only signals. Relying solely on the global `[CLS]` token (setting **(c)**) attains 25.7 FID, whereas modeling patch-level features (setting **(d)**, `ReDi`, linear PCA) improves to 21.4 FID. Both settings substantially outperform the baseline `SiT-B/2` backbone (33.0 FID, setting **(a)**), but the gap between **(c)** and **(d)** underscores that fine-grained *spatial* semantics are pivotal for improved generative modeling.

Non-linear compression unlocks local guidance. Replacing the linear PCA used in `ReDi` with our lightweight *non-linear* semantic compressor boosts patch-level joint modeling: setting **(i)** reaches 14.3 FID without any alignment loss, an absolute 7.1 FID reduction over `ReDi` (setting **(d)**). Notably, this also surpasses the state-of-the-art `REG` baseline (setting **(h)**, 15.2 FID), even though `REG` combines global modeling with external alignment. Moreover, enriching local guidance by aggregating multi-layer VFM patch-level features before compression (setting **(l)**) further reduces FID to 13.3, *without* using any global token or external alignment. The results indicate that rich spatial semantics are a crucial signal, and non-linear compression is key to unlocking their full benefit.

External alignment under local and global modeling. We analyze how `REPA` behaves when applied to *local* and/or *global* semantics. When the backbone does not

Table 1. **Impact of VFM semantics on `SiT-B/2` for improved generation. Results at 400K training steps.** (a) Baseline `SiT-B/2` in **Gray** , (b) `REPA` in **Purple**, (d) `ReDi` in **Light Green**, (h) `REG` in **Yellow**, and (i-n) `REGLUE` (ours) in **Light Cyan**. (✓) denotes novel components proposed in our work. While the nonlinear patch-level semantics alone yield substantial gains, the other listed components provide additional improvements. †Setting (n) uses a stronger DINOv3-B VFM; for fairness and consistency with prior work, all other experiments adopt DINOv2-B as the default VFM.

	Local (Patch-level)			Global (Image-level)	External Alignment		FID
	Non-linear	Linear	Multi-Layer	[CLS]	Patch-Level	[CLS]	
(a)							33.0
(b)					✓		24.4
(c)				✓			25.7
(d)		✓					21.4
(e)		✓			✓		18.8
(f)				✓		✓	33.7
(g)				✓	✓		15.5
(h)				✓	✓	✓	15.2
(i)	(✓)						14.3
(j)	(✓)				✓		14.1
(k)	(✓)			✓	✓	✓	13.7
(l)	(✓)		(✓)				13.3
(m)	(✓)		(✓)	✓	✓	✓	**12.9**
(n)	(✓)		(✓)	✓	✓	✓	**12.3**†

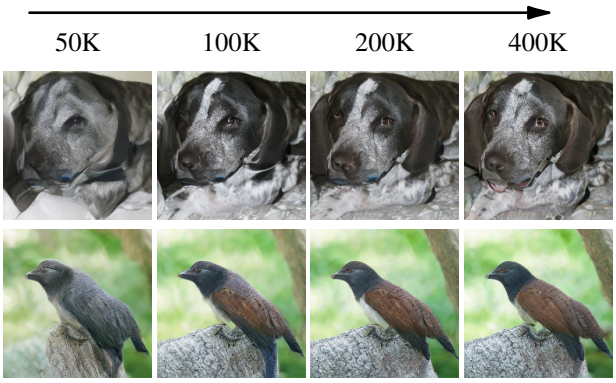

Figure 2. **`REGLUE` fast convergence.** Qualitative evolution of `SiT-B/2+REGLUE` at 50K/100K/200K/400K training steps. All use identical noise, the same sampling schedule/step count, and no classifier-free guidance. `REGLUE` achieves high fidelity early.

jointly model patch tokens, only aligning *local* VFM features already provides a strong boost: the original `REPA` configuration (setting **(b)**) improves the default `SiT-B/2` baseline from 33.0 **(a)** to 24.4 FID. `REPA` also improves the patch-only setting of `ReDi` **(d)** to **(e)** (from 21.4 to 18.8 FID). A similar pattern appears when starting from a model that only jointly models the global `[CLS]` token: adding *local-only* external alignment (setting **(g)**) reduces FID from 25.7 **(c)** to 15.5, and including the global component in the alignment as well (setting **(h)**) further improves it to 15.2. This indicates that local patch alignment is the dominant source of improvement, while global alignment provides a smaller, complementary gain. In contrast, aligning *only* the global information without any local alignment (setting **(f)**) degrades performance from 25.7 **(c)** to 33.7 FID, suggesting that alignment on global features alone is unstable without spatial anchors. Finally, in our setting, adding `REPA` on top of non-linear patch-level modeling improves FID from 14.3 **(i)** to 14.1 **(j)**, showing that once strong spatial semantics are jointly modeled, external alignment acts as a mild but consistent performance complement.

`REGLUE`: Joint local-global-latent modeling. Building on these observations, we progressively add the global token and alignment to compressed local modeling. Adding REPA-style alignment on top of **(i)** yields setting **(j)** with 14.1 FID (a modest but consistent gain), indicating that external supervision *complements* joint local modeling. Incorporating the global `[CLS]` and aligning both local and global signals (setting **(k)**) further improves to 13.7 FID. Finally, aggregating multi-layer patch features (from the last four VFM blocks) before compression (setting **(m)**) forms our final `REGLUE` unified setting, achieving 12.9 FID. To study the dependence of `REGLUE` on the underlying VFM, we also examine a stronger VFM, DINOv3-B. As reported in setting **(n)**, DINOv3-B yields the best result (FID 12.3), improving over DINOv2-B (FID 12.9) and indicating that `REGLUE` can effectively exploit a more powerful semantic encoder. Nevertheless, to remain consistent with prior work

Table 2. **Conditional and unconditional generation.** Comparison of SiT-B/2 with REPA, ReDi, REG, and REGLUE on ImageNet 256×256 without classifier-free guidance (CFG). We report parameter count, iterations, and FID.

Model	#Params	Iter.	FID↓
(a) Conditional Generation			
SiT-B/2	130M	400K	33.0
+ REPA	130M	400K	24.4
+ ReDi	130M	400K	21.4
+ REG	132M	400K	15.2
+ REGLUE (ours)	132M	300K	**14.5**
+ REGLUE (ours)	132M	400K	**12.9**
(b) Unconditional Generation			
SiT-B/2	130M	400K	59.8
+ ReDi	130M	400K	43.6
+ REG	132M	400K	29.7
+ REGLUE (ours)	132M	400K	**28.7**

and enable fair comparison, we adopt DINOv2-B as our default VFM in all main experiments. For a more in-depth analysis of our compressor, see Sec. 3.3.

4.3. Enhancing diffusion models

Accelerating convergence. Table 2(a) reports conditional ImageNet 256 × 256 results without classifier-free guidance (no CFG) with a SiT-B/2 backbone. REGLUE reaches 14.5 FID at 300K steps surpassing REG (15.2 at 400K) with 25% *fewer* iterations, and further improves to 12.9 at 400K. At 400K, REGLUE reduces FID by 60.9% vs. vanilla SiT-B/2 (33.0), 47.1% vs. REPA (24.4), and 39.7% vs. ReDi (21.4). In Figure 2, we show visual examples demonstrating that REGLUE achieves high-fidelity generations early in training. Moving to a larger backbone and more training steps, in Table 3 we show conditional ImageNet 256 × 256 results (no CFG) with a SiT-XL/2 backbone. At 200K steps, REGLUE achieves 4.6 FID, outperforming REG (5.0) and substantially surpassing REPA (11.1) and ReDi (12.5). Notably, REGLUE reaches 2.7 FID at 700K, matching REG's 1M performance (2.7) with 30% fewer iterations. At 1M, REGLUE sets the best score (2.5) vs. REG (2.7), ReDi (5.1), and REPA (6.4).

Unconditional generation. We evaluate our method in the unconditional setting and summarize the results in Table 2(b). The findings are consistent with the analysis in the previous section. Our REGLUE achieves 52%, 34.2%, and 3.4% improvements over SiT-B/2, ReDi, and REG, respectively, demonstrating the effectiveness of nonlinear compression and joint local–global feature modeling. Remarkably, even in this more challenging unconditional setting, REGLUE (28.7 FID) substantially outperforms the conditional SiT-B/2 baseline (33.0 FID).

Table 3. **Conditional generation.** Comparison of SiT-XL/2 with REPA, ReDi, REG, and REGLUE on ImageNet 256 × 256 without classifier-free guidance (CFG) under comparable settings. We report parameter count, iterations, and FID.

Model	#Params	Iter.	FID↓
SiT-XL/2	675M	7M	8.3
+ REPA	675M	200K	11.1
+ ReDi	675M	200K	12.5
+ REG	677M	200K	5.0
+ REGLUE (ours)	677M	200K	**4.6**
+ REPA	675M	400K	7.9
+ ReDi	675M	400K	7.5
+ REG	677M	400K	3.4
+ REGLUE (ours)	677M	400K	**3.2**
+ ReDi	675M	700K	5.6
+ REGLUE (ours)	677M	700K	**2.7**
+ REPA	675M	1M	6.4
+ ReDi	675M	1M	5.1
+ REG	677M	1M	2.7
+ REGLUE (ours)	677M	1M	**2.5**

Table 4. **Comparison with state-of-the-art.** Quantitative results on ImageNet 256 × 256 with classifier-free guidance (CFG). REPA, ReDi, REG and REGLUE employ an SiT-XL/2 model.

Model	Epochs	FID↓	sFID↓	IS↑	Pre.↑	Rec.↑
Autoregressive Models						
VAR	350	1.80	-	365.4	0.83	0.57
MagViTv2	1080	1.78	-	319.4	0.83	0.57
MAR	800	1.55	-	303.7	0.81	0.62
Latent Diffusion Models						
LDM	200	3.60	-	247.7	0.87	0.48
U-ViT-H/2	240	2.29	5.68	263.9	0.82	0.57
DiT-XL/2	1400	2.27	4.60	278.2	0.83	0.57
MaskDiT	1600	2.28	5.67	276.6	0.80	0.61
SD-DiT	480	3.23	-	-	-	-
SiT-XL/2	1400	2.06	4.50	270.3	0.82	0.59
FasterDiT	400	2.03	4.63	264.0	0.81	0.60
MDT	1300	1.79	4.57	283.0	0.81	0.61
Leveraging VFM's Representations						
ReDi	800	1.61	4.66	295.1	0.78	0.64
REPA	800	1.42	4.70	305.7	0.80	0.65
REG	800	1.36	4.25	299.4	0.77	0.66
REG	80	1.86	4.49	321.4	0.76	0.63
REGLUE (ours)	80	1.61	4.25	313.1	0.78	0.64
REG	160	1.59	4.36	304.6	0.77	0.65
REGLUE (ours)	160	1.53	4.26	320.4	0.78	0.65

State-of-the-art comparison. Table 4 reports quantitative results on ImageNet with classifier-free guidance. REGLUE improves over REG at matched epochs and closes the gap to longer-trained baselines. At 80 epochs, REGLUE lowers FID to 1.61 vs 1.86 for REG. At 160 epochs, it further improves to 1.53 vs 1.59. Although trained for 5× fewer epochs than the 800-epoch variants (REPA, ReDi, REG), the 160-epoch REGLUE remains competitive with models that leverage VFM representations and are

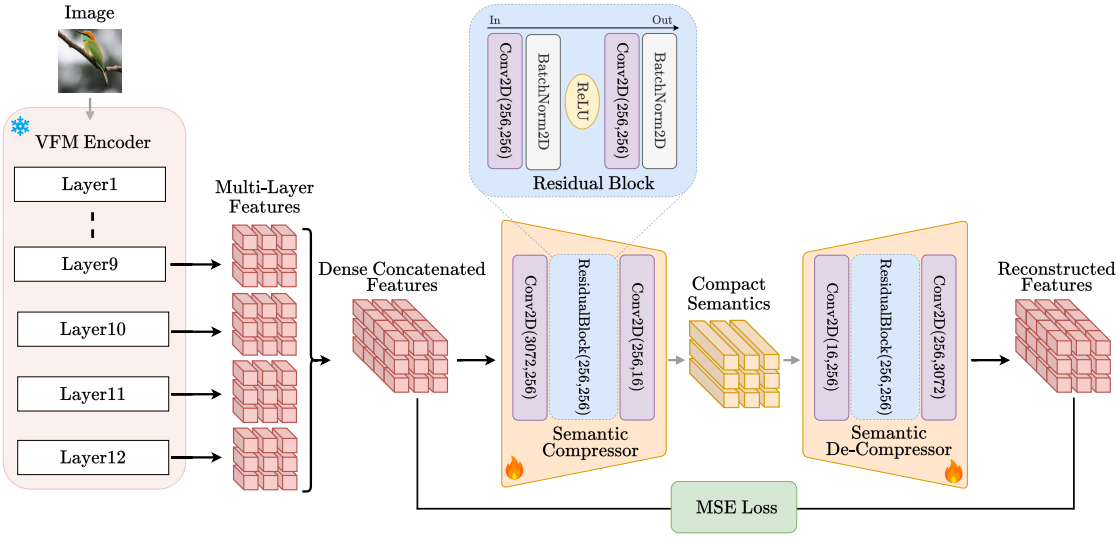

Figure 3. **Semantic compressor architecture and training.** The representations from the last four layers of the vision foundation model (VFM) encoder are concatenated and passed to the compression model, which projects them into a compact 16-channel semantic representation. In our default configuration (corresponding to the middle row of Table 5), the compressor maps the dense concatenated VFM features through an input layer Conv2D(3072, 256), a middle ResidualBlock(256, 256), and an output layer Conv2D(256, 16), where 256 is the hidden dimensionality. The semantic de-compressor then reconstructs the compact semantics back to their original dimensionality. The model is trained using an MSE loss between the dense concatenated features and their reconstructed counterparts.

trained for substantially longer (REPA, FID 1.42; REG, FID 1.36). Classifier-free guidance ablations are presented in Appendix Table 10. We provide qualitative results of generated images in Appendix Sec. E.

4.4. Semantic compressor impact

As we highlight in Sec. 3.3, the channel dimensionality of VFM representations is substantially higher than that of image latents, which can lead to degraded performance when fused naïvely. To mitigate this, ReDi [33] employs linear PCA to project the representations into an low-dimensional latent space; however, as we show in Table 1, this design choice is suboptimal. In contrast, we show that our non-linear CNN-based semantic compressor (Fig. 3) can substantially improve generation quality. In this section, we systematically examine its main design choices: the compression dimensionality (Fig. 5), the compressor capacity (Tab. 5), the set of VFM layers used as input (Tab. 6), and quantify their effect on both sample quality and efficiency. We further measure how much semantic information is preserved under compression using downstream probing tasks (Fig. 4).

Semantic preservation under compression. Figure 4 evaluates how well compressed patch-level features retain VFM semantics via attentive probing accuracy on ImageNet [52] and how this relates to generative quality (FID). Our non-linear semantic compressor preserves semantics much better than linear PCA: with only 8 channels it

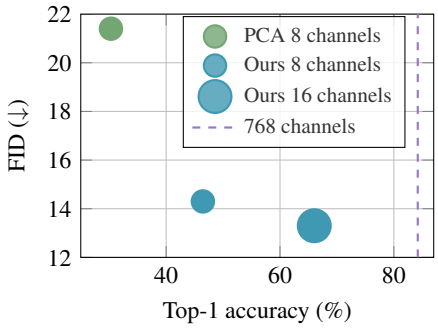

Figure 4. **Attentive probing accuracy vs. generation quality on ImageNet for different DINOv2 patch-level compression variants.** Each point shows top-1 attentive probing accuracy [52] and FID of the corresponding SiT model, with bubble area proportional to the semantic feature dimensionality. Our non-linear semantic compressors (8 and 16 channels) achieve substantially better FID at higher probing accuracy than the PCA-compressed features of ReDi, while the vertical dashed line marks the accuracy of the full 768-channel DINOv2 representation.

achieves substantially higher probing accuracy and lower FID than the PCA-compressed ReDi features, and increasing to 16 channels further improves both metrics, approaching the full 768-channel DINOv2 baseline. In contrast, the PCA-based compression in ReDi yields low probing accuracy and only modest FID gains, indicating that non-linear, spatially structured compression is key to preserving semantic information while improving generation qual-

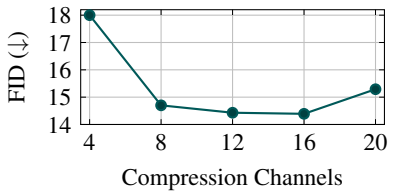

Figure 5. **Performance vs. compression channels.** Ablation of the final compression channels, in DINOv2 last layer's representation, using `SiT-B/2` trained for 400K steps without `REPA` loss.

Table 5. **Compression model size comparison.** End-to-end evaluation of compression model hidden layer sizes using the last 4 DINOv2 layers and `SiT-B/2` as the diffusion model, trained for 400K steps. Representations are compressed to 16 channels. `REPA` loss is not applied. Input and Output layers are Conv2D(*in*, *out*) layers. Each middle block is a Residual Block which comprises two 3×3 convolutional layers (stride = 1, padding = 1) with Batch Normalization and ReLU activation. Samples in Throughput column are VFM local representations.

Input Layer	Middle Block	Output Layer	#Params (M)	Throughput (samples/sec)	FID↓
(3072,128)	(128,128)	(128,16)	7.7	32,304	14.2
(3072,256)	(256,256)	(256,16)	16.6	18,578	**13.3**
(3072,512)	(512,512)	(512,16)	37.9	9,510	**13.3**

ity. Additional semantic preservation analysis with semantic segmentation experiments on Cityscapes [13] are presented in Appendix Sec. A.1.

Semantic compression dimensionality. In Figure 5, we examine how the number of compression channels affects generative performance. Aggressive compression (i.e., 4-channels) removes too much information, leading to degraded FID. Performance improves as we increase the number of channels up to 16, but degrades again at 20 channels. This suggests an optimal intermediate subspace: the compressed features preserve essential information to guide generation, yet are compact enough to stay balanced with the 4-channel image latents and not dominate the model's capacity. We therefore adopt 16 channels as the default in our main `REGLUE` configuration.

Lightweight compressor design. Our aim is to keep the semantic compressor lightweight while preserving essential information. In Table 5, we present an ablation over different hidden-layer widths. Reducing the hidden dimensionality from 256 to 128 channels degrades FID, indicating that overly constrained bottlenecks limit the ability to preserve essential semantic information. Increasing the hidden size from 256 to 512 channels yields no further improvement in FID, while doubling the model size and significantly reducing throughput, making this configuration inefficient. Pushing the capacity further (e.g., 1024 hidden size) leads to unstable compressor training. Moreover, increasing the model depth, via additional convolutional layers or residual blocks, exhibits the same instability. Overall, a shallow compressor with 256 hidden size offers the best balance between stability, efficiency, and generative performance.

Table 6. **The effect of multi-layer features.** We compare the generation performance without using `REPA` loss across different sets of Dino-V2 layers.The results are reported at 400K training steps. In all runs VFM patch tokens are compressed to 16 channels and `SiT-B/2` is used. The input layer of the compressor is changing accordingly, denoted by CNN In Dim column.

Dino Layers	CNN In Dim	FID↓
12	768	14.3
3, 6, 9, 12	3072	16.9
9, 10, 11, 12	3072	**13.3**

Choosing VFM layers for compression. In Table 6, we study how the choice of VFM layers fed into the semantic compressor affects generation performance. Motivated by [49], we compare three configurations of DINOv2-B patch features: (i) using only the last layer (12), (ii) using four intermediate layers (3, 6, 9, 12), and (iii) using the last four layers (9–12). In all cases, the compressed features are mapped to 16 channels and fed into the `SiT-B/2` diffusion model. Using only the final layer yields 14.3 FID, while including shallow intermediate layers (i.e., 3 & 6) degrades performance to 16.9 FID, indicating that early-layer features do not provide useful semantic guidance to the generation. In contrast, aggregating the last four layers (9–12) leads to 13.3 FID, suggesting that jointly compressing semantically rich, deeper VFM features provides the most beneficial signal for our framework. For more analysis about our compressor, see Appendix. Ablations about other loss variants of the compressor can be found in Appendix Table 7 and more experimentations with different VFMs in Table 8.

5. Conclusion

We introduced `REGLUE`, a unified generative model for latent diffusion that enables efficient entanglement of reconstruction-optimized and semantics-optimized image representations. By *jointly* modeling VAE latents with VFM patch-level & global semantics, coupled with *lightweight compression and aggregation* components, we have shown that `REGLUE` improves generation FID and accelerate convergence on ImageNet baselines by significant margins.

Acknowledgements We thank our colleague Theodoros Kouzelis for fruitful discussions. Bill was supported by the EU Horizon Europe programme MSCA PF RAVIOLI (No. 101205297). AWS resources were provided by the National Infrastructures for Research and Technology GRNET and funded by the EU Recovery and Resiliency Facility.

References

[1] Michael Samuel Albergo and Eric Vanden-Eijnden. Building normalizing flows with stochastic interpolants. In *ICLR*, 2023.

[2] Mahmoud Assran, Quentin Duval, Ishan Misra, Piotr Bojanowski, Pascal Vincent, Michael Rabbat, Yann LeCun, and Nicolas Ballas. Self-supervised learning from images with a joint-embedding predictive architecture. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 15619–15629, 2023.

[3] Fan Bao, Shen Nie, Kaiwen Xue, Yue Cao, Chongxuan Li, Hang Su, and Jun Zhu. All are worth words: A vit backbone for diffusion models. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 22669–22679, 2023.

[4] Hangbo Bao, Li Dong, Songhao Piao, and Furu Wei. Beit: Bert pre-training of image transformers. *arXiv preprint arXiv:2106.08254*, 2021.

[5] Christopher M Bishop. *Pattern recognition and machine learning*. Springer, 2006.

[6] Sam Bond-Taylor, Adam Leach, Yang Long, and Chris G Willcocks. Deep generative modelling: A comparative review of VAEs, GANs, normalizing flows, energy-based and autoregressive models. *IEEE transactions on pattern analysis and machine intelligence*, 44(11):7327–7347, 2021.

[7] Mathilde Caron, Hugo Touvron, Ishan Misra, Hervé Jegou, Julien Mairal, Piotr Bojanowski, and Armand Joulin. Emerging properties in self-supervised vision transformers. In *ICCV*, 2021.

[8] Hila Chefer, Uriel Singer, Amit Zohar, Yuval Kirstain, Adam Polyak, Yaniv Taigman, Lior Wolf, and Shelly Sheynin. Videojam: Joint appearance-motion representations for enhanced motion generation in video models. *arXiv preprint arXiv:2502.02492*, 2025.

[9] 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. In *ICML*, 2025.

[10] Liang-Chieh Chen, George Papandreou, Iasonas Kokkinos, Kevin Murphy, and Alan L Yuille. Deeplab: Semantic image segmentation with deep convolutional nets, atrous convolution, and fully connected crfs. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 2017.

[11] Ting Chen, Simon Kornblith, Mohammad Norouzi, and Geoffrey Hinton. A simple framework for contrastive learning of visual representations. In *International conference on machine learning*, pages 1597–1607. PmLR, 2020.

[12] Bowen Cheng, Ishan Misra, Alexander G Schwing, Alexander Kirillov, and Rohit Girdhar. Masked-attention mask transformer for universal image segmentation. In *CVPR*, 2022.

[13] Marius Cordts, Mohamed Omran, Sebastian Ramos, Timo Rehfeld, Markus Enzweiler, Rodrigo Benenson, Uwe Franke, Stefan Roth, and Bernt Schiele. The cityscapes dataset for semantic urban scene understanding. In *CVPR*, 2016.

[14] Timothée Darcet, Maxime Oquab, Julien Mairal, and Piotr Bojanowski. Vision transformers need registers. *arXiv preprint arXiv:2309.16588*, 2023.

[15] Jia Deng, Wei Dong, Richard Socher, Li-Jia Li, Kai Li, and Li Fei-Fei. Imagenet: A large-scale hierarchical image database. In *CVPR*, pages 248–255, 2009.

[16] Jia Deng, Wei Dong, Richard Socher, Li-Jia Li, Kai Li, and Li Fei-Fei. Imagenet: A large-scale hierarchical image database. In *2009 IEEE Conference on Computer Vision and Pattern Recognition*, pages 248–255, 2009.

[17] Prafulla Dhariwal and Alexander Quinn Nichol. Diffusion models beat GANs on image synthesis. In *NeurIPS*, 2021.

[18] Alexey Dosovitskiy. An image is worth 16x16 words: Transformers for image recognition at scale. *arXiv preprint arXiv:2010.11929*, 2020.

[19] Shanghua Gao, Pan Zhou, Ming-Ming Cheng, and Shuicheng Yan. Mdtv2: Masked diffusion transformer is a strong image synthesizer. *arXiv preprint arXiv:2303.14389*, 2023.

[20] Benyamin Ghojogh, Ali Ghodsi, Fakhri Karray, and Mark Crowley. Factor analysis, probabilistic principal component analysis, variational inference, and variational autoencoder: Tutorial and survey. *arXiv preprint arXiv:2101.00734*, 2021.

[21] Spyros Gidaris, Praveer Singh, and Nikos Komodakis. Unsupervised representation learning by predicting image rotations. *arXiv preprint arXiv:1803.07728*, 2018.

[22] Jean-Bastien Grill, Florian Strub, Florent Altché, Corentin Tallec, Pierre Richemond, Elena Buchatskaya, Carl Doersch, Bernardo Avila Pires, Zhaohan Guo, Mohammad Gheshlaghi Azar, et al. Bootstrap your own latent-a new approach to self-supervised learning. *Advances in neural information processing systems*, 33:21271–21284, 2020.

[23] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 770–778, 2016.

[24] Kaiming He, Xinlei Chen, Saining Xie, Yanghao Li, Piotr Dollár, and Ross Girshick. Masked autoencoders are scalable vision learners. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 16000–16009, 2022.

[25] 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*, 30, 2017.

[26] Jonathan Ho and Tim Salimans. Classifier-free diffusion guidance. *arXiv preprint arXiv:2207.12598*, 2022.

[27] Ioannis Kakogeorgiou, Spyros Gidaris, Bill Psomas, Yannis Avrithis, Andrei Bursuc, Konstantinos Karantzalos, and Nikos Komodakis. What to hide from your students: Attention-guided masked image modeling. In *European Conference on Computer Vision*, pages 300–318. Springer, 2022.

[28] Efstathios Karypidis, Ioannis Kakogeorgiou, Spyros Gidaris, and Nikos Komodakis. DINO-foresight: Looking into the future with DINO. In *The Thirty-ninth Annual Conference on Neural Information Processing Systems*, 2025.

[29] Diederik P. Kingma and Jimmy Ba. Adam: A method for stochastic optimization, 2017.

[30] Durk P Kingma and Prafulla Dhariwal. Glow: Generative flow with invertible 1x1 convolutions. *Advances in neural information processing systems*, 31, 2018.

[31] Giorgos Kordopatis-Zilos, Vladan Stojnić, Anna Manko, Pavel Suma, Nikolaos-Antonios Ypsilantis, Nikos Efthymiadis, Zakaria Laskar, Jiri Matas, Ondrej Chum, and Giorgos Tolias. Ilias: Instance-level image retrieval at scale. In *Proceedings of the Computer Vision and Pattern Recognition Conference*, pages 14777–14787, 2025.

[32] Theodoros Kouzelis, Ioannis Kakogeorgiou, Spyros Gidaris, and Nikos Komodakis. EQ-VAE: Equivariance regularized latent space for improved generative image modeling. In *Forty-second International Conference on Machine Learning*, 2025.

[33] Theodoros Kouzelis, Efstathios Karypidis, Ioannis Kakogeorgiou, Spyros Gidaris, and Nikos Komodakis. Boosting generative image modeling via joint image-feature synthesis. In *NeurIPS*, 2025.

[34] 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*, 32, 2019.

[35] Tuomas Kynkäänniemi, Miika Aittala, Tero Karras, Samuli Laine, Timo Aila, and Jaakko Lehtinen. Applying guidance in a limited interval improves sample and distribution quality in diffusion models. In *NeurIPS*, 2024.

[36] Yann LeCun. A path towards autonomous machine intelligence version 0.9. 2, 2022-06-27. *Open Review*, 62(1):1–62, 2022.

[37] Xingjian Leng, Jaskirat Singh, Yunzhong Hou, Zhenchang Xing, Saining Xie, and Liang Zheng. REPA-E: Unlocking VAE for end-to-end tuning with latent diffusion transformers. In *ICCV*, 2025.

[38] Tianhong Li, Yonglong Tian, He Li, Mingyang Deng, and Kaiming He. Autoregressive image generation without vector quantization. In *NeurIPS*, 2024.

[39] Tsung-Yi Lin, Piotr Dollár, Ross Girshick, Kaiming He, Bharath Hariharan, and Serge Belongie. Feature pyramid networks for object detection. In *CVPR*, 2017.

[40] Yaron Lipman, Ricky T. Q. Chen, Heli Ben-Hamu, Maximilian Nickel, and Matthew Le. Flow matching for generative modeling. In *The Eleventh International Conference on Learning Representations*, 2023.

[41] Jonathan Long, Evan Shelhamer, and Trevor Darrell. Fully convolutional networks for semantic segmentation. In *CVPR*, 2015.

[42] Ilya Loshchilov and Frank Hutter. Decoupled weight decay regularization, 2019.

[43] Calvin Luo. Understanding diffusion models: A unified perspective. *arXiv preprint arXiv:2208.11970*, 2022.

[44] Nanye Ma, Mark Goldstein, Michael S Albergo, Nicholas M Boffi, Eric Vanden-Eijnden, and Saining Xie. Sit: Exploring flow and diffusion-based generative models with scalable interpolant transformers. In *ECCV*, page 23–40, 2024.

[45] Tejaswini Medi, Hsien-Yi Wang, Arianna Rampini, and Margret Keuper. FAVAE-effective frequency aware latent tokenizer. In *NeurIPS 2025 Workshop: Reliable ML from Unreliable Data*, 2025.

[46] Charlie Nash, Jacob Menick, Sander Dieleman, and Peter W Battaglia. Generating images with sparse representations. *arXiv preprint arXiv:2103.03841*, 2021.

[47] Mehdi Noroozi and Paolo Favaro. Unsupervised learning of visual representations by solving jigsaw puzzles. In *European conference on computer vision*, pages 69–84. Springer, 2016.

[48] Aaron van den Oord, Yazhe Li, and Oriol Vinyals. Representation learning with contrastive predictive coding. *arXiv preprint arXiv:1807.03748*, 2018.

[49] Maxime Oquab, Timothée Darcet, Théo Moutakanni, Huy V. Vo, Marc Szafraniec, Vasil Khalidov, Pierre Fernandez, Daniel HAZIZA, Francisco Massa, Alaaeldin El-Nouby, Mido Assran, Nicolas Ballas, Wojciech Galuba, Russell Howes, Po-Yao Huang, Shang-Wen Li, Ishan Misra, Michael Rabbat, Vasu Sharma, Gabriel Synnaeve, Hu Xu, Herve Jegou, Julien Mairal, Patrick Labatut, Armand Joulin, and Piotr Bojanowski. DINOv2: Learning robust visual features without supervision. *Transactions on Machine Learning Research*, 2024.

[50] William Peebles and Saining Xie. Scalable diffusion models with transformers. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 4195–4205, 2023.

[51] Simon JD Prince. *Understanding deep learning*. MIT press, 2023.

[52] Bill Psomas, Dionysis Christopoulos, Eirini Baltzi, Ioannis Kakogeorgiou, Tilemachos Aravanis, Nikos Komodakis, Konstantinos Karantzalos, Yannis Avrithis, and Giorgos Tolias. Attention, please! revisiting attentive probing for masked image modeling. *arXiv preprint arXiv:2506.10178*, 2025.

[53] Bill Psomas, George Retsinas, Nikos Efthymiadis, Panagiotis Filntisis, Yannis Avrithis, Petros Maragos, Ondrej Chum, and Giorgos Tolias. Instance-level composed image retrieval. In *The Thirty-ninth Annual Conference on Neural Information Processing Systems*, 2025.

[54] Alec Radford, Jong Wook Kim, Chris Hallacy, Aditya Ramesh, Gabriel Goh, Sandhini Agarwal, Girish Sastry, Amanda Askell, Pamela Mishkin, Jack Clark, et al. Learning transferable visual models from natural language supervision. In *ICML*, 2021.

[55] René Ranftl, Alexey Bochkovskiy, and Vladlen Koltun. Vision transformers for dense prediction. In *CVPR*, 2021.

[56] Robin Rombach, Andreas Blattmann, Dominik Lorenz, Patrick Esser, and Björn Ommer. High-resolution image synthesis with latent diffusion models. In *CVPR*, pages 10684–10695, 2022.

[57] Tim Salimans, Ian Goodfellow, Wojciech Zaremba, Vicki Cheung, Alec Radford, and Xi Chen. Improved techniques for training gans. *Advances in neural information processing systems*, 29, 2016.

[58] Christoph Schuhmann, Romain Beaumont, Richard Vencu, Cade Gordon, Ross Wightman, Mehdi Cherti, Theo

Coombes, Aarush Katta, Clayton Mullis, Mitchell Wortsman, Patrick Schramowski, Srivatsa Kundurthy, Katherine Crowson, Ludwig Schmidt, Robert Kaczmarczyk, and Jenia Jitsev. Laion-5b: An open large-scale dataset for training next generation image-text models. In *NeurIPS*, pages 25278–25294, 2022.

[59] Minglei Shi, Haolin Wang, Wenzhao Zheng, Ziyang Yuan, Xiaoshi Wu, Xintao Wang, Pengfei Wan, Jie Zhou, and Jiwen Lu. Latent diffusion model without variational autoencoder. *arXiv preprint arXiv:2510.15301*, 2025.

[60] Yang Song, Jascha Sohl-Dickstein, Diederik P Kingma, Abhishek Kumar, Stefano Ermon, and Ben Poole. Score-based generative modeling through stochastic differential equations. *arXiv preprint arXiv:2011.13456*, 2020.

[61] Vladan Stojnić, Yannis Kalantidis, Jiří Matas, and Giorgos Tolias. Lposs: Label propagation over patches and pixels for open-vocabulary semantic segmentation. In *Proceedings of the Computer Vision and Pattern Recognition Conference*, pages 9794–9803, 2025.

[62] Zineng Tang, Ziyi Yang, Chenguang Zhu, Michael Zeng, and Mohit Bansal. Any-to-any generation via composable diffusion. *Advances in Neural Information Processing Systems*, 36:16083–16099, 2023.

[63] Keyu Tian, Yi Jiang, Zehuan Yuan, Bingyue Peng, and Liwei Wang. Visual autoregressive modeling: Scalable image generation via next-scale prediction. *arXiv preprint arXiv:2404.02905*, 2024.

[64] Michael E Tipping and Christopher M Bishop. Probabilistic principal component analysis. *Journal of the Royal Statistical Society Series B: Statistical Methodology*, 61(3):611–622, 1999.

[65] Michael Tschannen, Alexey Gritsenko, Xiao Wang, Muhammad Ferjad Naeem, Ibrahim Alabdulmohsin, Nikhil Parthasarathy, Talfan Evans, Lucas Beyer, Ye Xia, Basil Mustafa, Olivier Hénaff, Jeremiah Harmsen, Andreas Steiner, and Xiaohua Zhai. Siglip 2: Multilingual vision-language encoders with improved semantic understanding, localization, and dense features. *arXiv preprint arXiv:2502.14786*, 2025.

[66] Shuai Wang, Zhi Tian, Weilin Huang, and Limin Wang. DDT: Decoupled diffusion transformer. *arXiv preprint arXiv:2504.05741*, 2025.

[67] Ge Wu, Shen Zhang, Ruijing Shi, Shanghua Gao, Zhenyuan Chen, Lei Wang, Zhaowei Chen, Hongcheng Gao, Yao Tang, Jian Yang, et al. Representation entanglement for generation: Training diffusion transformers is much easier than you think. *NeurIPS*, 2025.

[68] Lihe Yang, Bingyi Kang, Zilong Huang, Xiaogang Xu, Jiashi Feng, and Hengshuang Zhao. Depth anything: Unleashing the power of large-scale unlabeled data. In *CVPR*, 2024.

[69] Lihe Yang, Bingyi Kang, Zilong Huang, Zhen Zhao, Xiaogang Xu, Jiashi Feng, and Hengshuang Zhao. Depth anything v2. In *NeurIPS*, 2024.

[70] Jingfeng Yao, Cheng Wang, Wenyu Liu, and Xinggang Wang. Fasterdit: Towards faster diffusion transformers training without architecture modification. In *NeurIPS*, 2024.

[71] Jingfeng Yao, Bin Yang, and Xinggang Wang. Reconstruction vs. generation: Taming optimization dilemma in latent diffusion models. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2025.

[72] Lijun Yu, Jose Lezama, Nitesh Bharadwaj Gundavarapu, Luca Versari, Kihyuk Sohn, David Minnen, Yong Cheng, Agrim Gupta, Xiuye Gu, Alexander G Hauptmann, Boqing Gong, Ming-Hsuan Yang, Irfan Essa, David A Ross, and Lu Jiang. Language model beats diffusion - tokenizer is key to visual generation. In *The Twelfth International Conference on Learning Representations*, 2024.

[73] Sihyun Yu, Sangkyung Kwak, Huiwon Jang, Jongheon Jeong, Jonathan Huang, Jinwoo Shin, and Saining Xie. Representation alignment for generation: Training diffusion transformers is easier than you think. In *ICLR*, 2025.

[74] Mohsen Zand, Ali Etemad, and Michael Greenspan. Diffusion models with deterministic normalizing flow priors. *Transactions of Machine Learning Research*, 2024.

[75] Kaiwen Zha, Lijun Yu, Alireza Fathi, David A Ross, Cordelia Schmid, Dina Katabi, and Xiuye Gu. Language-guided image tokenization for generation. In *CVPR*, pages 15713–15722, 2025.

[76] Shuangfei Zhai, Ruixiang Zhang, Preetum Nakkiran, David Berthelot, Jiatao Gu, Huangjie Zheng, Tianrong Chen, Miguel Angel Bautista, Navdeep Jaitly, and Josh Susskind. Normalizing flows are capable generative models. In *ICML*, 2025.

[77] Xiaohua Zhai, Basil Mustafa, Alexander Kolesnikov, and Lucas Beyer. Sigmoid Loss for Language Image Pre-Training . In *ICCV*, pages 11941–11952, 2023.

[78] Qinsheng Zhang and Yongxin Chen. Diffusion normalizing flow. In *NeurIPS*, pages 16280–16291, 2021.

[79] Hengshuang Zhao, Jianping Shi, Xiaojuan Qi, Xiaogang Wang, and Jiaya Jia. Pyramid scene parsing network. In *CVPR*, 2017.

[80] Hongkai Zheng, Weili Nie, Arash Vahdat, and Anima Anandkumar. Fast training of diffusion models with masked transformers. *arXiv preprint arXiv:2306.09305*, 2023.

[81] Jinghao Zhou, Chen Wei, Huiyu Wang, Wei Shen, Cihang Xie, Alan Yuille, and Tao Kong. ibot: Image bert pre-training with online tokenizer. *International Conference on Learning Representations (ICLR)*, 2022.

[82] Rui Zhu, Yingwei Pan, Yehao Li, Ting Yao, Zhenglong Sun, Tao Mei, and Chang Wen Chen. Sd-dit: Unleashing the power of self-supervised discrimination in diffusion transformer. In *CVPR*, pages 8435–8445, 2024.

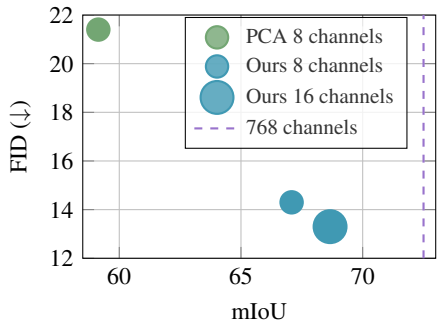

Figure 6. **Semantic segmentation performance mIoU vs generation quality for different DINOv2 patch-level compression variants.** Each point shows the segmentation mIoU on Cityscapes [13] using a DPT [55] head on frozen features following implementation from [28, 68, 69] and the FID on ImageNet of the corresponding `SiT` model. Bubble area is proportional to feature dimensionality. Our non-linear semantic compressors (8 and 16 channels) achieve substantially better FID at higher mIoU than the PCA-compressed features of `ReDi`. The vertical dashed line indicates the mIoU of the full 768-channel DINOv2 representation.

A. Additional Experimental Results

A.1. Semantic preservation under compression

To assess how well our compressed patch-level features preserve vision foundation model (VFM) semantics, we perform an additional experiment on *semantic segmentation*. Figure 6 shows semantic segmentation on Cityscapes [13] measured by mIoU (using DPT [55] head on frozen features) versus ImageNet FID for different DINOv2 patch–level compression variants. At the same 8-channel compression, our nonlinear compressor achieves 67.1 mIoU/14.3 FID, notably better than 59.1/21.4 of PCA (+8.0 mIoU/−7.1 FID). Increasing to 16 channels further improves to 68.7 mIoU/13.3 FID. Despite having 96× or 48× less channels (respectively) than the original 768-channel DINOv2 representation (vertical dashed line at 72.5 mIoU), our compressed variants effectively retain most of the semantics while substantially improving generative fidelity, indicating that the learned nonlinear compressor preserves semantic representations and is a better fit than linear PCA in joint semantics-VAE latents modeling.

A.2. Semantic compressor auxiliary objectives

We further investigate how the compressor's training objectives impact downstream generation. Starting from our MSE-only autoencoder, we (i) switch to a variational formulation with a KL term and (ii) add an adversarial (GAN) loss. In all experiments, we compress the *last* VFM layer to 16 channels and, during diffusion training, model only local semantics without external alignment (similar to setting (i) in Table 1). We follow [56] for the VAE/GAN setup: KL weight 10^{-6}; a lightweight two-layer discriminator applied from the start; all other hyperparameters identical to the MSE baseline. Table 7 shows that plain MSE yields the best performance (14.3 FID and 6.7 sFID). Adding KL noticeably degrades performance (17.2 FID, 7.1 sFID), while including the GAN term provides no gains and slightly worsens overall performance. Overall, exploring these well-established additions does not yield any further improvements in our setting.

Table 7. **Semantic compressor loss variants.** ImageNet 256×256 comparison without classifier-free guidance (CFG) using `SiT-B/2`. We compare the impact of different auxiliary objectives in our semantic compressor. We use DINOv2-B last layer representations compressed to 16 channels. For `REGLUE`, we follow setting (i) in Table 1 (main paper).

Training Objectives	FID↓	sFID↓	Precision↑	Recall↑
MSE	**14.3**	6.7	0.62	0.65
MSE+KL	17.2	7.1	0.63	0.64
MSE+GAN	14.4	6.9	0.62	0.65

Table 8. **`REGLUE` with different VFMs.** ImageNet 256×256 comparison without CFG using `SiT-B/2`. For `REGLUE`, we follow setting (m/n) in Table 1 (main paper).

VFM	FID↓	sFID	Precision↑	Recall↑
DINOv2-B	12.9	5.8	0.67	0.63
DINOv3-B	**12.3**	5.8	0.67	0.63
CLIP-L	18.1	7.1	0.63	0.62

A.3. Impact of VFM

To evaluate `REGLUE` across different VFMs, we experiment with three encoders: DINOv2-B, DINOv3-B, and CLIP-L. For each backbone, we concatenate the last four layers and adapt the compressor's input projection to the corresponding embedding size (*e.g.*, 4×768=3072 for DINOv2-B, 4×1024=4096 for CLIP-L). All compressors are trained for 25 epochs with a target compression of 16 channels, and the downstream `SiT-B/2` generator is trained for 400K steps in every setting for fair comparison. Table 8 reports FID, sFID, precision, and recall. DINOv3-B delivers the best generation quality (lowest FID), DINOv2-B is a close second, while CLIP-L lags behind. As already discussed in the main paper, to remain consistent with prior work [33, 67, 73], we adopt DINOv2-B as our default VFM.

A.4. Detailed benchmark

We provide a detailed evaluation of `SiT-XL/2+REGLUE` with more training iterations and additional metrics. Table 9 demonstrates the performance, reporting FID, sFID, inception score, precision, and recall. Notably, `REGLUE` reaches 7.8 FID at 100K steps, already surpassing the vanilla `SiT-XL/2` baseline at 7M steps (8.3 FID). It con-

tinues to improve substantially, reaching 3.2 at 400K, 2.6 at 750K, and 2.5 at 1M steps.

Table 9. **Detailed evaluation** for `SiT-XL/2+REGLUE`. ImageNet 256×256 without CFG.

MODEL	#ITERS.	FID↓	sFID↓	IS↑	PREC.↑	REC.↑
`SiT-XL/2`	7M	8.3	6.3	131.7	0.68	0.67
w/ `REGLUE`	50K	20.0	6.3	64.7	0.65	0.58
w/ `REGLUE`	100K	7.8	4.7	116.0	0.64	0.57
w/ `REGLUE`	200K	4.6	4.4	148.0	0.74	0.63
w/ `REGLUE`	400K	3.2	4.3	171.6	0.75	0.63
w/ `REGLUE`	700K	2.7	4.2	185.0	0.76	0.65
w/ `REGLUE`	750K	2.6	4.1	185.6	0.76	0.65
w/ `REGLUE`	1M	2.5	4.1	188.6	0.76	0.65

A.5. Classifier-free guidance

We provide more evaluation results for classifier-free guidance scales and guidance intervals. We denote by w the CFG scale applied to the VAE latents and the VFM representations, and use VAE-Only to refer to the setting where CFG is applied exclusively to the VAE latents. We also vary the guidance *interval* $[0, \tau]$, following [35]. Table 10 presents ImageNet 256×256 results for `SiT-XL/2+REGLUE` at 800K steps.

A.6. Limited data

In Figure 7, we evaluate data efficiency by training `SiT-B/2` for 80 epochs on class-balanced ImageNet sets of 20%, 50%, and 100%. `REGLUE` consistently outperforms `REG`, with larger gains when data is scarce: -5.5 FID at 20% and -3.4 at 50%. This indicates that jointly modeling compact local and global VFM semantics improves robustness on data-limited regimes.

Table 10. **CFG ablations on `SiT-XL/2+REGLUE` (800K steps, ImageNet 256×256).** We vary the guidance interval $[0, \tau]$ and scale w. VAE-only applies CFG only to the VAE latents; otherwise CFG is applied to both VAE latents and VFM representations.

INTERVAL	w	VAE-ONLY	FID ↓	sFID ↓	IS ↑	PREC. ↑	REC. ↑
$[0, 0.85]$	2.8	False	1.55	4.30	278.20	0.77	0.66
$[0, 0.90]$	2.8	False	1.53	4.26	320.43	0.78	0.65
$[0, 0.95]$	2.8	False	2.75	4.20	395.91	0.82	0.60
$[0, 0.90]$	2.7	False	1.53	4.3	315.02	0.78	0.65
$[0, 0.90]$	2.8	False	1.53	4.3	320.43	0.78	0.65
$[0, 0.90]$	2.9	False	1.56	4.4	323.68	0.78	0.65
$[0, 0.90]$	2.8	False	1.53	4.3	320.43	0.78	0.65
$[0, 0.90]$	2.8	True	1.84	4.5	235.28	0.76	0.66

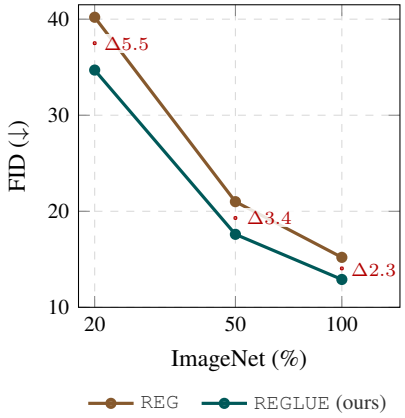

Figure 7. **Dataset pruning on ImageNet.** FID on ImageNet 256×256 for `SiT-B/2` trained for 80 epochs on class-balanced subsets (20%, 50%, 100% of ImageNet). `REGLUE` consistently outperforms `REG`, with improvements of -5.5, -3.4, and -2.3 FID at 20%, 50%, and 100%, respectively.

B. Additional Experimental Setup

B.1. Semantic compressor details

Architecture settings. The compression model is a lightweight convolutional autoencoder composed of the *semantic compressor*, which encodes the high-dimensional VFM features into a compact representation and the *semantic de-compressor*, which symmetrically decodes them back to their original space. The detailed architecture is presented in Figure 3. The semantic encoder is composed of three main components: an *input layer*, a *middle block*, and an *output layer*. The input layer is a 3×3 convolutional layer (3072→256), where 3072 corresponds to the number of input channels from the concatenated VFM features and 256 denotes the hidden size. The middle block is a residual block (Conv–BN–ReLU–Conv–BN, 256 channels, identity skip) that preserves spatial shape. The output layer is a convolutional layer (256→16) that projects the representation to 16 compressed channels. A symmetric semantic de-compressor mirrors this design (16→256→3072). The model is fully convolutional, preserves the spatial resolution, and is trained with an MSE reconstruction loss; at inference we retain only the encoder to provide compact local semantics.

Optimization settings. We train the semantic compressor for 25 epochs with an MSE reconstruction loss between the concatenated multi-layer VFM features and their decoded counterparts. We use Adam [29] with a learning rate of 1×10^{-3}, $(\beta_1, \beta_2) = (0.9, 0.999)$, batch size 4096, and no weight decay. The learning rate decays with a cosine schedule to a final value of 8.5×10^{-4}. Figure 8 plots the training curve: the loss decreases smoothly and plateaus by the final epoch, indicating stable convergence. The model is lightweight; a full run finishes in *under one hour* on 8×A100 GPUs.

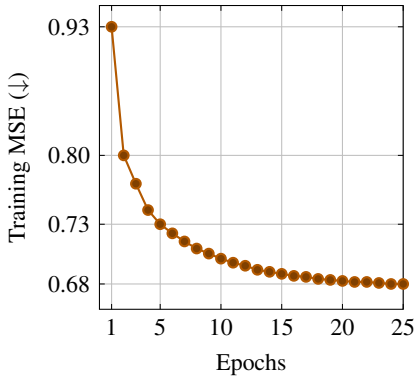

Figure 8. **Compressor training.** Training curve showing MSE loss over epochs. The compression model utilizes 4 last DINOv2-B layers. The Input layer is 256 and the compression is to 16 channels. A run finishes in *less than one hour* on 8×A100 GPUs.

B.2. SiT details

Architecture settings. We adopt the official `SiT` configurations [44]. The base `SiT-B/2` (132M params) uses 12 transformer blocks with embedding dimension 768 and 12 attention heads. The larger `SiT-XL/2` (677M params) uses 28 blocks with embedding dimension 1152 and 16 heads.

Table 11. **`SiT` optimization settings.**

Optimization	
Batch size	256
Optimizer	AdamW
Learning Rate	0.0001
(β_1, β_2)	(0.9, 0.999)
Interpolants	
α_t	$1 - t$
σ_t	t
Training objective	v-prediction
Sampler	Euler-Maruyama
Sampling steps	250

Optimization settings. We use AdamW [29, 42] with a constant learning rate of 1×10^{-4}, $(\beta_1, \beta_2) = (0.9, 0.999)$, and a batch size of 256 for both `SiT` models. To speed up training, we use mixed-precision (fp16) with gradient clipping. We also pre-compute image latent using $\mathrm{SD-VAE}$ [56]. The training objective is $\mathrm{v-prediction}$, and we use the Euler–Maruyama sampler with 250 steps, defining the interpolants as $\alpha_t = 1 - t$ and $\sigma_t = t$. We provide the optimization details in Table 11.

C. Limitations and Future Work

As shown in Table 2 and Table 3 in the main paper, `REGLUE` consistently improves sample quality under comparable training budgets and reaches or surpasses strong baselines in substantially fewer iterations. However, due to resource constraints, we restrict `SiT-XL/2+REGLUE` experiments to 1M iterations and do not explore full convergence at ultra–long schedules (*e.g.*, 4M iterations). On our available compute (8×A100 GPUs), a single 1M-iteration `SiT-XL/2` run requires roughly 7 days, making exploration of such long schedules impractical. Higher-resolution ImageNet 512×512 experiments are an interesting next step as well; in this work, we instead prioritize configurations that we consider also interesting and practical, such as limited-data regimes (see Figure 7).

Beyond scaling, our results point to several promising extensions. First, swapping DINOv2 for DINOv3 yields further gains (Table 8 and Table 1 in the main paper), suggesting that stronger VFMs could enhance `REGLUE`. Second, while we currently include the raw global `[CLS]` token, learning a compact *global* compressor (analogous to our spatial one) may better balance global–local capacity within our joint modeling framework.

D. Baseline Generative Models

We briefly summarize the baselines used in our comparisons. Autoregressive baselines include `VAR` [63], which progressively predicts fine image details from coarse inputs across multiple scales; `MagViTv2` [72], which removes lookup tables in quantization to support much larger token vocabularies; and `MAR` [38] which avoids vector quantization altogether in an autoregressive setup.

For latent diffusion models, we consider `LDM` [56], which performs diffusion in a compact latent space; `DiT` [50], a transformer-based architecture; `U-ViT-H/2` [3] a ViT-based diffusion model with skip connections; `MaskDiT` [80], which adds a mask-reconstruction auxiliary task; `MDT` [19], which employs an asymmetric masked latent modeling scheme; `SD-DiT` [82], which augments `MaskDiT` with a momentum-encoder–based discrimination loss; `SiT` [44], which recasts the `DiT` backbone within a continuous-time interpolant framework; and `FasterDiT` [70], which accelerates training through velocity-supervised objectives.

Finally, among methods that explicitly use visual representations, we include `REPA` [73], which aligns diffusion-model features with local VFM features; `ReDi` [33] which linearly compress VFM features and directly model them in the diffusion model; and finally `REG` [67] which models the global VFM representation while also applying `REPA`-style alignment.

E. Visualizations

We present uncurated, class-conditional samples from `SiT-XL/2+REGLUE` trained for 1M steps at 256×256 in Figure 9 and Figure 10. We use CFG with w=4.0. The grids illustrate both fine-grained detail (textures and object parts) and diversity within each class.

Figure 9. **Uncurated ImageNet 256 × 256 samples.** Class-conditional generations from `SiT-XL/2+REGLUE` trained for 1M steps with CFG (w=4.0). Grids illustrate great fidelity and intra-class diversity.

Figure 10. **Uncurated ImageNet 256 × 256 samples.** Class-conditional generations from `SiT-XL/2+REGLUE` trained for 1M steps with CFG (w=4.0). Grids illustrate great fidelity and intra-class diversity.