

# TIDE : Temporal-Aware Sparse Autoencoders for Interpretable Diffusion Transformers in Image Generation

Victor Shea-Jay Huang<sup>\*1,2</sup>, Le Zhuo<sup>\*1,2</sup>, Yi Xin<sup>2</sup>, Zhaokai Wang<sup>2</sup>, Fu-Yun Wang<sup>1</sup>, Yuchi Wang<sup>1</sup>, Renrui Zhang<sup>1</sup>, Peng Gao<sup>2</sup>, Hongsheng Li<sup>†1,2,3</sup>

<sup>1</sup>Multimedia Laboratory (MMLab), The Chinese University of Hong Kong

<sup>2</sup>Shanghai AI Laboratory

<sup>3</sup>CPII under InnoHK

jeix782@gmail.com, hsli@ee.cuhk.edu.hk

## Abstract

Diffusion Transformers (DiTs) are a powerful yet underexplored class of generative models compared to U-Net-based diffusion architectures. We propose TIDE—Temporal-aware sparse autoencoders for Interpretable Diffusion transformers—a framework designed to extract sparse, interpretable activation features across timesteps in DiTs. TIDE effectively captures temporally-varying representations and reveals that DiTs naturally learn hierarchical semantics (e.g., 3D structure, object class, and fine-grained concepts) during large-scale pretraining. Experiments show that TIDE enhances interpretability and controllability while maintaining reasonable generation quality, enabling applications such as safe image editing and style transfer.

## Introduction

Recent advancements in diffusion models have revolutionized text-to-image generation, with Diffusion Transformers (DiTs) (Peebles and Xie 2022) emerging as a scalable alternative to traditional U-Net backbones (Ronneberger, Fischer, and Brox 2015) by leveraging transformer-based architectures. Improvements in sampling efficiency (Yi, Li et al. 2024) and model scaling (Esser, Kulal et al. 2024; Zhuo, Du et al. 2024; Gao, Zhuo et al. 2024; Xin, Yan et al. 2025) have further elevated their performance to a new level. However, a critical question remains: how do these models function internally? While interpretability techniques have been extensively developed for U-Net frameworks (Kwon, Jeong, and Uh 2023; Toker et al. 2024), similar efforts for DiTs are absent. This is primarily due to the fact that most existing methods are tailored to convolutional architectures and depend on spatially structured feature maps, which are not directly available in the token-based representation of DiTs. More broadly, interpretability research on diffusion models lags significantly behind that of GANs (Goodfellow et al. 2014; Bau et al. 2018; Ling, Kreis et al. 2021), despite the growing prominence of diffusion models in generative modeling. This gap largely stems from diffusion mod-

els’ reliance on latent space transformations across various denoising levels, which complicates efforts to analyze their internal mechanisms. In particular, little attention has been paid to understanding latent activations within these models, leaving a significant void in our ability to interpret their processes.

Feature interpretability, which aims to uncover how neural networks encode functionality through internal representations (Elhage, Nanda et al. 2021), has seen notable success in language models, especially through the use of sparse autoencoders (SAEs) (Ng et al. 2011). These tools disentangle polysemantic neurons into sparse, human-interpretable features (Bricken et al. 2023; Marks et al. 2024; Huang et al. 2025). However, existing SAE methods face significant limitations when applied to diffusion models, particularly at the level of activation layers. Diffusion models present unique challenges due to their temporal dynamics: unlike language models, which process static tokens in an auto-regressive manner, DiTs iteratively refine latent representations across timesteps, integrating multimodal inputs from text encoders. This iterative nature introduces difficulties such as long token sequences, high token dimensionality, and time-varying activation patterns, especially for high-resolution images. Similarly, although SAEs have proven effective in vision-language models (Daujotas 2024a) and convolutional networks (Gorton 2024), they encounter distinct challenges in DiTs, including:

1. Capturing time-varying activation patterns across diffusion steps;
2. Balancing sparsity and reconstruction fidelity during iterative denoising;
3. Aligning sparse features with both visual and text-conditioned semantics.

These challenges underscore the need for new interpretability methods tailored to the temporal and multimodal complexities of DiTs with SAE.

To address these issues, we present Temporal-aware SAEs for Interpretable Diffusion transformERs (TIDE), a framework designed to capture temporal variations in activation patterns and extract sparse, interpretable features throughout the denoising process. To ensure high-fidelity feature learning under sparsity constraints, we introduce

<sup>\*</sup>These authors contributed equally.

<sup>†</sup>Corresponding author

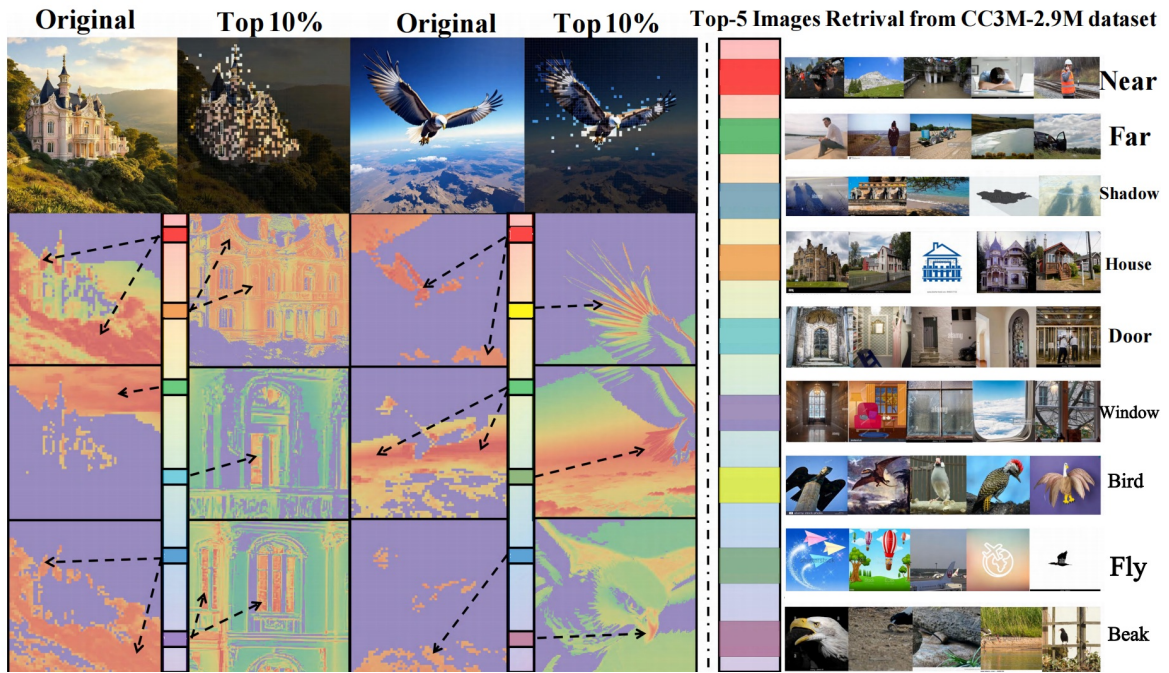


Figure 1: TIDE effectively encodes interpretable features across different levels (class, semantic, and 3D) in pre-trained diffusion transformer. This validates that the diffusion model inherently captures and organizes these multi-level concepts by large-scale generative pre-training, enabling it to perform various downstream diffusion tasks effectively.

specialized training strategies, including progressive sparsity scheduling and random sampling augmentation. Additionally, we explore TIDE’s scaling laws and validate its performance using diffusion loss, demonstrating that TIDE achieves advanced results compared to other SAE-based methods. Moreover, TIDE effectively learns multi-level features (e.g., 3D, semantics, and class), as shown in Fig. 1, highlighting how diffusion models inherently organize such concepts through large-scale generative pre-training. This finding illustrates their ability to be adapted to various downstream tasks, such as depth estimation (Gui, Schusterbauer et al. 2024; Ke, Obukhov et al. 2024), semantic segmentation (Tian, Aggarwal et al. 2024), and image classification (Li et al. 2023; Mukhopadhyay et al. 2023). Additionally, we propose a temporal analysis that isolates and interprets the roles of individual timesteps in the generation process, offering deeper insights into the latent representations of text-conditioned image synthesis. By tailoring SAEs to the dynamics of DiTs, TIDE advances model interpretability, enhances understanding of generative mechanisms, and lays the foundation for more trustworthy and controllable generative systems.

This paper presents the first systematic study of SAEs in DiT-based diffusion models. Our contributions are fourfold:

- **TIDE Framework:** A paradigm that adapts SAEs to DiT via temporal-aware architecture and progressive sparsity scheduling, coupled with training strategies to extract interpretable features from DiT’s latent space while preserving generation quality and balancing sparsity with denoising efficacy.

- **Scaling Law and Evaluation:** We analyze the scaling law of TIDE, highlighting its ease of training, and compare its downstream loss with other SAEs to demonstrate the superior performance of TIDE.
- **Interpretation:** Using TIDE, we demonstrated that diffusion models learn hierarchical features through pre-training, enabling efficient transfer to target domains via fine-tuning. Furthermore, through a decoupled analysis of sampling timesteps, we verify and extend existing understanding about DiT’s Coarse-to-Fine generation.
- **Potential Applications:** Using PixArt- $\alpha$  (Chen et al. 2023) and other diffusion models, we demonstrate the practical utility of TIDE-learned sparse features in tasks like safe image editing and style transfer.

## Related Work

### Sparse Autoencoders for Interpretability

Mechanistic interpretability (Elhage, Nanda et al. 2021) aims to understand how neural networks produce outputs by identifying features encoded in neurons (Olah et al. 2020). Sparse autoencoders (SAEs)(Olshausen and Field 1997) address the polysemantic nature of neurons caused by superposition, where networks encode more features than dimensions(?Sharkey, Braun, and beren 2022), by enforcing sparsity, activating only a small subset of neurons per input. SAEs have proven effective in uncovering human-interpretable features (Huben et al. 2024; Bricken et al. 2023), decomposing activations in large language models

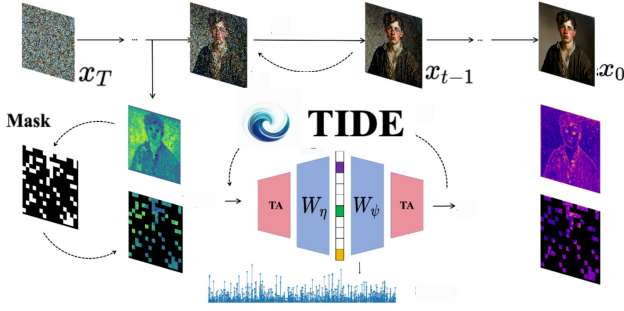


Figure 2: Overview of our training process: For each DiT, we train its TIDE (Embed SAE within the TA (temporal-aware) architecture) individually, extracting activation layers from different timesteps each time. The training of the SAE is conducted using both with and without random sampling.

(LLMs)(Templeton 2024; Yun et al. 2021), and extracting interpretable concepts in vision-language models such as CLIP(Daujotas 2024b). They have also facilitated image editing (Daujotas 2024a), enabled concept naming (Rao, Mahajan et al. 2024), and disabled specific generative features (Bricken et al. 2023; Duan, et al. 2024; Wen, Wang et al. 2025), demonstrating their potential to mitigate superposition and enhance interpretability across diverse model architectures.

### Interpretability of Diffusion Models

Diffusion models achieve text-to-image generation via iterative denoising, evolving from U-Net architectures (Rombach, Blattmann et al. 2022a; Podell et al. 2023) to scalable designs like Diffusion Transformer (DiT) (Peebles and Xie 2022), with improved sampling efficiency (Song and Ermon 2019; Song, Meng, and Ermon 2020; Lu et al. 2022). Interpretability efforts focus on bottleneck layers (Kwon, Jeong, and Uh 2023; Park et al. 2023) and cross-attention (Tang, Liu et al. 2022), enabling manipulation of attributes (Basu et al. 2023) and text encoders (Toker et al. 2024). SAEs, applied to CLIP (Radford et al. 2021; Fry 2024) and diffusion models (Ijishakin, Ang et al. 2024; Surkov et al. 2024; Kim, Thomas, and Ghadiyaram 2024; Yin et al. 2025; An et al. 2025), help identify features, reduce hallucinations (Jiang et al. 2024), and analyze activations, linking features to visual semantics. For more details on related work, please refer to the supplementary material.

## Preliminaries

### Sparse Autoencoder

To encode the extracted activations into a sparse latent representation, we employ a  $k$ -sparse autoencoder. The encoder and decoder are defined as follows:

$$\begin{aligned} l &= \text{ReLU}(W_{\text{enc}}a + b_{\text{enc}}), \\ \hat{a} &= W_{\text{dec}}l, \end{aligned} \quad (1)$$

where the encoder weights  $W_{\text{enc}} \in \mathbb{R}^{n \times f}$ , encoder bias  $b_{\text{enc}} \in \mathbb{R}^n$ , and decoder weights  $W_{\text{dec}} \in \mathbb{R}^{f \times n}$  define the pa-

rameters of the autoencoder. Here,  $a \in \mathbb{R}^f$  is the input vector,  $l \in \mathbb{R}^n$  is the latent feature representation, and  $\hat{a} \in \mathbb{R}^f$  is the reconstructed vector.

### TopK Activation

We employ the  $k$ -sparse autoencoder (Makhzani and Frey 2013), which uses the TopK activation function to enforce sparsity. This function selects the  $k$  largest latent activations and sets the rest to zero. The function is defined as:

$$z = \text{TopK}(l), \quad (2)$$

where  $l \in \mathbb{R}^n$  is the latent feature representation. The decoder remains unchanged.

As shown in (Gao, la Tour et al. 2024), TopK activation offers several advantages: it directly enforces sparsity by controlling the number of active units, eliminating the need for  $\ell_1$  regularization; simplifies model comparison by avoiding hyperparameter tuning; improves sparsity-reconstruction trade-offs, especially in larger models; and enhances feature interpretability by zeroing out small activations.

### Diffusion Transformer

Diffusion Transformers (DiT)(Peebles and Xie 2022) are transformer-based diffusion models operating in latent space, offering competitive image generation performance. Unlike traditional U-Net-based models, DiT leverages transformers to process latent patches, enhancing scalability and efficiency. In the latent diffusion framework (LDM)(Rombach, Blattmann et al. 2022a), images are first encoded into latents via a variational autoencoder (VAE) (Kingma et al. 2013), and then a transformer models the denoising process, reducing computation while maintaining quality.

The forward diffusion process follows a Markov chain, sequentially adding Gaussian noise. At timestep  $t$ , the noisy data  $x_t$  is formulated as:

$$x_t = \sqrt{\bar{\alpha}_t}x_0 + \sqrt{1 - \bar{\alpha}_t}\epsilon_t, \quad \epsilon_t \sim \mathcal{N}(0, \mathbf{I}), \quad (3)$$

where  $\bar{\alpha}_t$  defines the noise schedule. The reverse process, modeled by the transformer, predicts noise  $\epsilon_\theta(x_t)$  to iteratively denoise  $x_t$ . The training objective minimizes the mean squared error between predicted and true noise:

$$L_{\text{simple}}(\theta) = \mathbb{E}_{x_0, \epsilon_t, t} [\|\epsilon_\theta(x_t) - \epsilon_t\|^2]. \quad (4)$$

DiT's scalable architecture and strong performance on benchmarks like ImageNet position it as an advanced model for high-resolution image generation.

## Methodology

### Where to Train Sparse Autoencoders (SAEs)

Recent research on the mechanistic interpretability of diffusion models (Basu et al. 2023; Basu, Rezaei et al. 2024) has revealed that distinct cross-attention blocks are responsible for generating particular visual elements, including style and objects. Using these insights, since diffusion models share similar architectures, we selected a representative model,

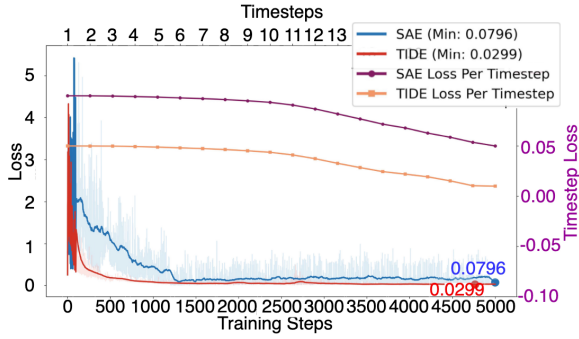


Figure 3: TIDE integrates timestep-dependent modulation into the original SAE architecture, achieving significantly faster convergence and enhanced performance.

**PixArt-XL-2-1024** (Chen et al. 2023), to extract activation layers for training SAEs. A dedicated SAE was trained for each activation layer. The training process is in Fig. 2.

### Activations in DiTs

To enable the use of latent activations from diffusion models for training SAEs, activations are extracted from specific transformer layers during both the forward diffusion and reverse denoising processes.

Given an image  $\mathbf{x}_0$  and a text prompt  $\mathbf{p}$ , the image is encoded into the latent space  $\mathbf{z}_0$  using a VAE encoder. Noise  $\epsilon$  is added to simulate the diffusion process, producing a noisy latent  $\mathbf{z}_t$  at timestep  $t$ . Simultaneously, the prompt is embedded into  $\mathbf{e}_p$  via a text encoder. The noisy latent  $\mathbf{z}_t$  and text embedding  $\mathbf{e}_p$  are fed into the transformer, where hooks capture activations from specified layers during the forward pass. These activations are then used to train SAEs.

This approach leverages intermediate representations from the generation pipeline for downstream tasks. The activation extraction process is detailed in algorithm below:

---

#### Algorithm 1: Extract Activations from DiTs

---

**Require:** Image  $\mathbf{x}_0$ , Text Prompt  $\mathbf{p}$ , Time step  $t$ , DiTs  $\mathcal{T}$ , Target layer  $\ell$

**Ensure:** Activations  $\mathbf{a}_\ell$

- 1: Encode image:  $\mathbf{z}_0 \leftarrow f_{\text{enc}}(\mathbf{x}_0)$
  - 2: Sample noise:  $\epsilon \sim \mathcal{N}(0, \mathbf{I})$
  - 3: Add noise:  $\mathbf{z}_t \leftarrow \sqrt{\bar{\alpha}_t} \mathbf{z}_0 + \sqrt{1 - \bar{\alpha}_t} \epsilon$
  - 4: Encode text:  $\mathbf{e}_p \leftarrow f_{\text{text}}(\mathbf{p})$
  - 5: Register hook at layer  $\ell$ :  $g_\ell : \mathbf{a}_\ell = \phi_\ell(\mathbf{u}_\ell)$
  - 6: Forward pass:  $\mathbf{h}_{\text{out}} \leftarrow \mathcal{T}(\mathbf{z}_t, \mathbf{e}_p, t)$
  - 7: Extract activations:  $\mathbf{a}_\ell \leftarrow g_\ell(\mathbf{u}_\ell)$
  - 8: **return**  $\mathbf{a}_\ell$
- 

### TIDE: Temporal-Aware SAE

In our training process, we observed that directly reconstructing the activation layers of DiTs using a SAE yielded relatively favorable results. However, the convergence was still noticeably slow. Inspired by the timestep sampling

mechanism in diffusion models, we introduce temporal information into the activation layers to enhance the representation of timestep-dependent features.

The inclusion of temporal information is motivated by the fact that the token distribution varies significantly across different timesteps. For instance, at smaller timesteps (low  $t$ ), the amount of information present is substantially different compared to larger timesteps (high  $t$ ), where the input contains significant noise. To effectively disentangle these timestep-specific characteristics, it is crucial to inform the SAE of the corresponding timestep during reconstruction.

To address these timestep-dependent activation variations, we propose the Temporal-Aware SAE (TIDE), which incorporates a timestep-dependent modulation operation:

$$x_{\text{mod}} = x \cdot (1 + \text{scale}(t)) + \text{shift}(t), \quad (5)$$

where both  $\text{scale}(t)$  and  $\text{shift}(t)$  are functions of the current timestep  $t$ , modeled via a lightweight MLP initialized from the pre-trained adaptive Layer Normalization in DiT. This modulation allows the model to adaptively align activations across different diffusion timesteps.

As shown in Fig. 3, this modulation adjusts the activations by applying timestep-specific scaling and shifting, allowing the model to adapt to the varying distributions of activation layers across different timesteps.

Building upon the Top-k SAE framework, the addition of this modulation mechanism enables TIDE to effectively handle timestep-dependent activation distributions without requiring substantial modifications to the SAE parameters.

To prevent dead latents—latent variables that cease to activate during training—we employ careful initialization and revival strategies. Specifically, the encoder is initialized as the transpose of the decoder, and any inactive latents are periodically revived throughout the training process.

## Experiments

### Training Setup

Each SAE/TIDE was carried out using three distinct strategies, with 40 hours of training on 8 A100 80G GPUs:

- **Direct Training:** Using 80K prompt-image pairs extracted from CC3M (Sharma, Ding et al. 2018) without augmentation.
- **Augmented Training:** To mitigate overly dense local features and reduce overfitting, we randomly sample 1/16 of the tokens at the token level for each activation layer. Furthermore, we use a consistent set of 80K pairs to control for the effects of the sampling operation.
- **TIDE Training:** Incorporating a *Temporal-Aware Architecture* to enhance the model’s ability to capture temporal dynamics within the activation features.

**Evaluation** The performance of TIDE was evaluated using cosine similarity, mean squared error (MSE), and sparsity loss (see the supplementary material). These metrics also guided the training process, ensuring that the TIDEs(SAEs) struck a balance between sparsity and reconstruction accuracy. A validation set of 500 prompt-image pairs is sampled from the CC3M dataset and is not used

Train	$MSE$	$S_{cos}$	$MSE^{val}$	$S_{cos}^{val}$
SAE	$3.3e-3$	0.968	$1.23e-2$	0.935
SAE(sample)	$2.5e-3$	0.970	$1.26e-2$	0.962
TIDE	$3.1e-3$	0.970	$1.02e-2$	0.944
TIDE(sample)	$2.5e-3$	0.972	$1.02e-2$	0.964

Table 1: SAE ( $d \times 16d$ , sparsity 5%, Avg of 28 layers) training Results (w/o Sampling Strategy and TA Architecture).

during training. For TIDE (SAE) models trained solely on the training data, this validation set is considered out-of-distribution (OOD).

### SAEs / TIDEs Exhibit Superior Reconstruction

We evaluate the effect of token-level sampling during training and the integration of the temporal-aware (TA) architecture into PixArt’s 28 DiT activation layers. As shown in Tab. 1, sampling only 1/16 of tokens substantially improves training-set reconstruction (cols. 1–2) with only a marginal increase in MSE. More importantly, it yields a clear boost in cosine similarity on the OOD validation set (col. 4), indicating stronger generalization. In contrast, training without sampling provides limited reconstruction gains and fails to enhance OOD alignment.

This result reveals two key insights: **(1)** Token-level sampling improves learning efficiency by focusing on more informative local features, while also mitigating the impact of global noise—leading to better generalization. **(2)** The negligible change in MSE indicates that the norm of the activation distribution remains stable, implying that the reconstruction process preserves essential structural information and does not adversely affect downstream image generation quality.

Finally, incorporating the TA architecture yields consistent improvements in both MSE and cosine similarity, enabling more accurate reconstructions at the  $10^{-3}$  scale. This benefit stems from a sufficiently expressive latent space, well-chosen sparsity, and the architectural consistency with diffusion time-embedding dynamics.

### Scaling Laws and Downstream Loss of TIDEs

Given the broad capabilities of advanced diffusion models—distinct from large language models (LLMs)—diffusion models exhibit significantly higher feature dimensionality and process a large number of tokens in each operation. Consequently, we hypothesize that accurately capturing the state of the model requires a large number of sparse features. To address this, we propose two primary approaches for determining the size of the autoencoder and the token budget:

As shown in Fig. 4(a)(b), increasing the number of latent dimensions from  $d = 4608$  to  $16d = 73728$ , along with a higher top-k selection, leads to progressively better reconstruction of the activation layers. This improvement is reflected in a MSE reaching the order of  $10^{-3}$  and a cosine similarity approaching 0.97. Compared to other diffusion U-Net-based SAEs, as well as most advanced LLM-based SAEs, the TIDEs trained using our strategy consistently achieve advanced performance.

We use the downstream diffusion loss to evaluate the reconstruction quality of TIDE. Specifically, the activation reconstructed by TIDE are used to replace the original activation in DiT, and the diffusion loss is then calculated. As shown in Fig. 4(c)(d), the diffusion loss of TIDE’s reconstructed features is comparable to that of the original model and surpasses other types of SAEs, all while achieving advanced performance. This highlights TIDE’s effectiveness in preserving features crucial to the generation task.

### Diffusion Really Learned Features

Recent studies have demonstrated the versatility of diffusion models across a range of downstream tasks, including depth estimation, classification, and segmentation. This raises two fundamental questions:

1. **Do diffusion models inherently encode the knowledge necessary for these tasks?**
2. **Can TIDE validate the existence of such representations, thereby demonstrating their potential to support downstream applications?**

To address these questions, we conducted a series of well-designed experiments to investigate the internal feature representations learned by diffusion models. For each target feature of interest, we selected 1K images containing identical target elements from a curated dataset optimized to minimize background interference and irrelevant semantic content. From these images, we extracted activations from the penultimate layer of the DiT model and applied average pooling to aggregate the latent features. Based on these pooled features, we computed similarity scores between the input image tokens and identified the top-k tokens with the highest similarity. The results, as illustrated in Fig. 1, demonstrate that TIDE effectively exploits the internal representations of DiT to extract multi-level features, which can be categorized as follows:

- **3D-Level:** TIDE captures features like “far,” “near,” and “shadow” that reflect foreground-background relationships. These features show highly specific activation patterns, and images with similar activations consistently share such 3D characteristics—highlighting the model’s ability to encode depth-related concepts useful for tasks like depth estimation.
- **Class-Level:** TIDE learns class-specific features and their subcomponents. For example, “Rococo” activates features like “house,” “window,” and “door,” while “eagle” activates “bird,” “fly,” and “beak.” These activations are localized, reflecting a clear hierarchical structure in the learned representations.
- **Semantic-Level:** At a finer granularity, TIDE identifies detailed semantic features and their spatial positions—e.g., pinpointing the “beak” of an eagle—demonstrating its capacity for extracting precise, localized semantic information.

As shown in Fig.1, these results confirm that TIDE effectively encodes interpretable features across multiple levels—3D, class, and semantic—within pre-trained DiTs. This




Figure 4: (a)(b) Scaling laws of convergence loss with fixed latents  $n$  under MSE and cosine similarity loss ( $L_{cos} = 1 - S_{cos}$ ). (c)(d) Comparison of TIDE and other activation functions: For 73728-d latents, TIDE achieves better trade-offs in diffusion loss against both cosine similarity and MSE.



Figure 5: By manipulating the latent space of the TIDE, we can achieve various concept transformations, such as erasing Rococo style, increasing age, and altering the orientation and shape of architecture. The extent of feature modification intensifies as the number of altered tokens increases.

highlights the hierarchical structure of features learned during large-scale generative pre-training, demonstrating that diffusion models inherently capture and organize multi-level concepts. Such representations support diverse downstream tasks. For example, class- and semantic-level features enable classification (Li et al. 2023; Mukhopadhyay et al. 2023) and segmentation, including zero-shot methods like Diff-Seg (Tian, Aggarwal et al. 2024). Likewise, the ability to learn 3D concepts allows these models to perform depth estimation, as shown in methods like Marigold (Ke, Obukhov et al. 2024) and DepthFM (Gui, Schusterbauer et al. 2024).

### Potential Application

Based on our previous analysis, we modify the activation layer of the penultimate DiT block, which strongly influences the model output. For each timestep, a consistent directional vector exists between the latent feature and its target, highlighting the role of Self-Attention Embedding. By operating in the TIDE latent space instead of directly editing activations, we improve interpretability and avoid out-of-distribution (OOD) issues. We further apply CFG Renormalization (Lin, Liu et al. 2024) to maintain stable feature norms.

The interpretable latent space supports flexible editing

strategies, such as erasing or inverting top-K features to remove fine details, scaling activation indices for smooth feature transitions, or replacing latent features for global modifications. These operations yield diverse visual effects, as illustrated in Fig. 5.

### Ablation Study

To understand the effectiveness of our proposed TIDE module, we conduct comprehensive ablation experiments across four axes: (1) **hyper-parameter configurations**, (2) **DiT layer positioning**, (3) **generalization across diffusion backbones** and (4) **safe image editing**.

**Evaluation Metrics.** We evaluate performance using three complementary metrics:

- **Fréchet Inception Distance (FID)** (Heusel, Ramsauer et al. 2017), using the FID-30K protocol on MSCOCO (Lin et al. 2014), to measure generation quality.
- **Compositional Alignment**, via T2I-CompBench (Huang et al. 2023), using AlignScore to evaluate semantic consistency with compositional prompts.
- **Safety**, using the I2P benchmark (Schramowski et al. 2023), reporting Attack Success Rate (ASR) to assess robustness against unsafe prompt injection.

**Hyper-parameter Analysis** We analyze the influence of key TIDE hyper-parameters, including sparsity ratio, latent dimensionality, and token sampling. As shown in Table 2, all TIDE variants have slightly higher FID than the baseline, but the default setting (5% sparsity, token sampling, 16d latent) yields the smallest FID increase (+0.15) and minimal drops in AlignScore, indicating negligible performance loss.

Reducing sparsity to 3% further increases FID and weakens alignment, while disabling token sampling leads to overfitting, reflected in worse FID and semantic alignment. Increasing latent dimension improves expressiveness: 32d achieves the best Complex score (0.407), though still below baseline. Overall, results show a clear sparsity–capacity trade-off, and suggest that moderate sparsity with adequate latent dimensions minimizes degradation.

**Impact of DiT Layer Selection** We study the effect of inserting TIDE at different DiT layers. Table 2 shows that

Setting	Configuration	ASR(%) ↑	FID ↓	ΔFID ↓	AlignScore ↑					
					Attribute Binding			Object Relationship		Complex
					Color	Shape	Texture	Spatial	Non-Spatial	
Hyper-parameters	Baseline (no SAE/TIDE)	–	7.30	–	0.688	0.549	0.708	0.210	0.318	0.412
	SAE (5% sparsity)	–	7.93	+0.63	0.602	0.479	0.652	0.148	0.254	0.352
	TIDE (5%, 16d, default)	–	7.45	+0.15	0.647	0.514	0.679	0.185	0.271	0.392
	TIDE (3% sparsity)	–	7.60	+0.30	0.589	0.443	0.636	0.153	0.238	0.348
	TIDE (latent dim = 4d)	–	7.97	+0.67	0.562	0.438	0.621	0.147	0.213	0.335
	TIDE (latent dim = 8d)	–	7.88	+0.58	0.585	0.467	0.649	0.169	0.229	0.342
	TIDE (latent dim = 32d)	–	7.41	+0.11	0.672	0.505	0.696	0.192	0.262	0.407
DiT Layers	Baseline (no SAE/TIDE)	–	7.30	–	0.688	0.549	0.708	0.210	0.318	0.412
	SAE (Block 27)	–	7.34	+0.04	0.663	0.531	0.652	0.182	0.294	0.382
	TIDE (Block 4, early)	–	10.49	+3.19	0.495	0.337	0.475	0.102	0.152	0.141
	TIDE (Block 10, mid)	–	8.92	+1.62	0.607	0.441	0.554	0.133	0.256	0.244
	TIDE (Block 27, penultimate)	–	7.64	+0.34	0.672	0.537	0.667	0.195	0.301	0.352
	TIDE (Block 28, final)	–	7.34	+0.04	0.672	0.529	0.657	0.193	0.309	0.349
Generalization	SDXL (Podell et al. 2023)	–	8.24	0.00	0.638	0.544	0.528	0.212	0.318	0.417
	SDXL + SAE	–	8.44	+0.20	0.583	0.504	0.451	0.167	0.295	0.348
	SDXL + TIDE	–	8.38	+0.14	0.607	0.528	0.495	0.181	0.309	0.373
	FLUX-dev (Labs 2024)	–	10.09	0.00	0.612	0.583	0.480	0.208	0.287	0.401
	FLUX-dev + SAE	–	10.18	+0.09	0.489	0.441	0.432	0.178	0.216	0.314
	FLUX-dev + TIDE	–	10.13	+0.04	0.549	0.505	0.453	0.192	0.247	0.356
Safety	SDv1.4 (Rombach, Blattmann et al. 2022b)	17.80	16.40	0.00	0.310	0.336	0.347	0.344	0.350	0.339
	ESD (Gandikota et al. 2024)	2.84	18.08	+1.68	0.315	0.34	0.351	0.347	0.355	0.343
	SA (Heng and Soh 2024)	2.78	24.96	+8.56	0.319	0.344	0.355	0.354	0.359	0.348
	RECE (Gong, Chen et al. 2024)	0.80	17.02	+0.62	0.321	0.337	0.348	0.345	0.351	0.34
	SAFREE (Yoon, Yu et al. 2024)	1.48	19.88	+3.48	0.316	0.341	0.342	0.348	0.356	0.344
	CS (Kim and Ghadiyaram 2025)	0.42	18.59	+2.19	0.327	0.345	0.351	0.368	0.369	0.357
	TIDE (ours)	0.64	19.72	+3.32	0.317	0.345	0.357	0.352	0.366	0.33

Table 2: The use of TIDE achieves better editing performance without compromising the baseline quality. All results are evaluated on the MS-COCO dataset using FID (↓), ASR (↑) on I2P, and AlignScore (↑) under the T2I-CompBench benchmark. We highlight the best value in blue, and the second-best in green.

deeper layers offer better trade-offs: injecting at the penultimate block (Block 27) produces the smallest FID increase (+0.34) and strong semantic alignment (e.g., Texture = 0.537), making it the most effective insertion location.

By contrast, early-layer injection (e.g., Block 4) significantly harms performance, yielding high FID (+3.19) and lower alignment. These findings confirm that deeper layers carry richer semantics and better tolerate sparse interpretability modules without damaging image quality.

**Generalization Across Diffusion Backbones** To evaluate model-agnostic applicability, we deploy TIDE on two additional diffusion models: SDXL and Flux-dev. As reported in Table 2, TIDE introduces only minor FID increases (8.24→8.38 on SDXL, 10.09→10.13 on Flux-dev), while alignment scores stay close to baseline or slightly improve.

This demonstrates that TIDE generalizes smoothly across architectures, maintaining stable generation quality with low overhead. Although TIDE does not directly improve raw generation metrics, its value lies in adding interpretability with minimal performance cost.

**Safety Evaluation** We further assess TIDE on the I2P benchmark for NSFW-to-safe image editing. As shown in Table 2, TIDE attains a low Attack Success Rate (ASR = 0.64%), competitive with CS (0.42%) and surpassing SAFREE (1.48%), while maintaining good generation quality (FID = 19.72).

Unlike previous experiments, this setting evaluates real image editing, where TIDE must both preserve visual qual-

ity and sanitize unsafe content. Our results show that TIDE offers strong safety protection with only minor fidelity loss.

These ablations demonstrate that:

- TIDE introduces minimal performance degradation, with the smallest increases in FID and smallest drops in AlignScore across hyper-parameters, layers, and backbones.
- Moderate sparsity and 16–32d latent space offer the best trade-off between interpretability and performance.
- Deeper DiT layers (especially Block 27) are the optimal points for interpretability injection.
- TIDE generalizes well across diffusion backbones, confirming its robustness and portability.
- In safe editing, TIDE filters NSFW content while preserving quality and defense.

## Conclusion

TIDE introduces a sparse, temporally-aware interpretability framework for Diffusion Transformers, enabling transparent analysis of their internal representations. By uncovering multi-level semantic features, TIDE supports practical applications like image editing and demonstrates strong generalization across models and tasks. While current performance in complex editing scenarios remains limited by sparsity constraints, future work will explore more flexible sparse mechanisms to improve expressiveness. Overall, TIDE provides a solid foundation for building more trustworthy, controllable, and interpretable generative systems.

## Acknowledgments

This study was supported in part by the National Key R&D Program of China Project 2022ZD0161100, in part by the Centre for Perceptual and Interactive Intelligence, a CUHK-led InnoCentre under the InnoHK initiative of the Innovation and Technology Commission of the Hong Kong Special Administrative Region Government, in part by NSFC-RGC Project N.CUHK498/24, and in part by the Guangdong Basic and Applied Basic Research Foundation (No. 2023B1515130008, XW).

## References

- An, J.; et al. 2025. Agfsync: Leveraging ai-generated feedback for preference optimization in text-to-image generation. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 39, 1746–1754.
- Basu, S.; Rezaei, K.; et al. 2024. On Mechanistic Knowledge Localization in Text-to-Image Generative Models. In *Forty-first International Conference on Machine Learning*.
- Basu, S.; Zhao, N.; Morariu, V. I.; Feizi, S.; and Manjunatha, V. 2023. Localizing and editing knowledge in text-to-image generative models. In *The Twelfth International Conference on Learning Representations*.
- Bau, D.; Zhu, J.-Y.; Strobel, H.; Zhou, B.; Tenenbaum, J. B.; Freeman, W. T.; and Torralba, A. 2018. Gan dissection: Visualizing and understanding generative adversarial networks. *arXiv preprint arXiv:1811.10597*.
- Bricken, T.; Templeton, A.; Batson, J.; Chen, B.; et al. 2023. Towards Monosemanticity: Decomposing Language Models With Dictionary Learning. *Transformer Circuits Thread*.
- Chen, J.; et al. 2023. PixArt- $\alpha$ : Fast Training of Diffusion Transformer for Photorealistic Text-to-Image Synthesis.
- Daujotas, G. 2024a. Case Study: Interpreting, Manipulating, and Controlling CLIP With Sparse Autoencoders.
- Daujotas, G. 2024b. Interpreting and Steering Features in Images.
- Duan, Z.; et al. 2024. A Non-negative VAE: the Generalized Gamma Belief Network. *arXiv:2408.03388*.
- Elhage, N.; Nanda, N.; et al. 2021. A Mathematical Framework for Transformer Circuits. *Anthropic*.
- Esser, P.; Kulal, S.; et al. 2024. Scaling rectified flow transformers for high-resolution image synthesis.
- Fry, H. 2024. Towards Multimodal Interpretability: Learning Sparse Interpretable Features in Vision Transformers.
- Gandikota, R.; Orgad, H.; Belinkov, Y.; Materzyńska, J.; and Bau, D. 2024. Unified Concept Editing in Diffusion Models. In *Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision (WACV)*, 5111–5120.
- Gao, L.; la Tour, T. D.; et al. 2024. Scaling and evaluating sparse autoencoders. *arXiv preprint arXiv:2406.04093*.
- Gao, P.; Zhuo, L.; et al. 2024. Lumina-T2X: Scalable Flow-based Large Diffusion Transformer for Flexible Resolution Generation. In *The Thirteenth International Conference on Learning Representations*.
- Gong, C.; Chen, K.; et al. 2024. Reliable and efficient concept erasure of text-to-image diffusion models. In *European Conference on Computer Vision*, 73–88. Springer.
- Goodfellow, I.; et al. 2014. Generative adversarial nets. *Advances in neural information processing systems*, 27.
- Gorton, L. 2024. The Missing Curve Detectors of InceptionV1: Applying Sparse Autoencoders to InceptionV1 Early Vision. *arXiv preprint arXiv:2406.03662*.
- Gui, M.; Schusterbauer, J.; et al. 2024. Depthfm: Fast monocular depth estimation with flow matching. *arXiv preprint arXiv:2403.13788*.
- Heng, A.; and Soh, H. 2024. Selective amnesia: A continual learning approach to forgetting in deep generative models. *Advances in Neural Information Processing Systems*, 36.
- Heusel, M.; Ramsauer, H.; et al. 2017. Gans trained by a two time-scale update rule converge to a local nash equilibrium. *Advances in neural information processing systems*, 30.
- Huang, K.; Sun, K.; Xie, E.; Li, Z.; and Liu, X. 2023. T2i-compbench: A comprehensive benchmark for open-world compositional text-to-image generation.
- Huang, X.; Wang, X.; et al. 2025. Medical mllm is vulnerable: Cross-modality jailbreak and mismatched attacks on medical multimodal large language models. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 39, 3797–3805.
- Huben, R.; Cunningham, H.; Smith, L. R.; Ewart, A.; and Sharkey, L. 2024. Sparse Autoencoders Find Highly Interpretable Features in Language Models. In *The Twelfth International Conference on Learning Representations*.
- Ijishakin, A.; Ang, M. L.; et al. 2024. H-Space Sparse Autoencoders. In *Neurips Safe Generative AI Workshop 2024*.
- Jiang, N.; et al. 2024. Interpreting and Editing Vision-Language Representations to Mitigate Hallucinations. *arXiv preprint arXiv:2410.02762*.
- Ke, B.; Obukhov, A.; et al. 2024. Repurposing diffusion-based image generators for monocular depth estimation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 9492–9502.
- Kim, D.; and Ghadiyaram, D. 2025. Concept steerers: Leveraging k-sparse autoencoders for controllable generations. *arXiv preprint arXiv:2501.19066*.
- Kim, D.; Thomas, X.; and Ghadiyaram, D. 2024. *Revelio*: Interpreting and leveraging semantic information in diffusion models. *arXiv preprint arXiv:2411.16725*.
- Kingma, D. P.; et al. 2013. Auto-encoding variational bayes.
- Kwon, M.; Jeong, J.; and Uh, Y. 2023. Diffusion Models Already Have A Semantic Latent Space. In *The Eleventh International Conference on Learning Representations*.
- Labs, B. F. 2024. FLUX. <https://github.com/black-forest-labs/flux>.
- Li, A. C.; Prabhudesai, M.; Duggal, S.; Brown, E.; and Pathak, D. 2023. Your diffusion model is secretly a zero-shot classifier. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 2206–2217.

- Lin, T.-Y.; Maire, M.; Belongie, S.; Hays, J.; Perona, P.; Ramanan, D.; Dollár, P.; and Zitnick, C. L. 2014. Microsoft coco: Common objects in context. In *European conference on computer vision*, 740–755. Springer.
- Lin, Z.; Liu, W.; et al. 2024. STIV: Scalable Text and Image Conditioned Video Generation. *arXiv preprint arXiv:2412.07730*.
- Ling, H.; Kreis, K.; et al. 2021. Editgan: High-precision semantic image editing. *Advances in Neural Information Processing Systems*, 34: 16331–16345.
- Lu, C.; Zhou, Y.; Bao, F.; Chen, J.; Li, C.; and Zhu, J. 2022. Dpm-solver: A fast ode solver for diffusion probabilistic model sampling in around 10 steps. 35: 5775–5787.
- Makhzani, A.; and Frey, B. 2013. K-sparse autoencoders. *arXiv preprint arXiv:1312.5663*.
- Marks, S.; Rager, C.; Michaud, E. J.; Belinkov, Y.; Bau, D.; and Mueller, A. 2024. Sparse feature circuits: Discovering and editing interpretable causal graphs in language models. *arXiv preprint arXiv:2403.19647*.
- Mukhopadhyay, S.; Gwilliam, M.; Agarwal, V.; Padmanabhan, N.; Swaminathan, A.; Hegde, S.; Zhou, T.; and Shrivastava, A. 2023. Diffusion models beat gans on image classification. *arXiv preprint arXiv:2307.08702*.
- Ng, A.; et al. 2011. Sparse autoencoder. *CS294A Lecture notes*, 72(2011): 1–19.
- Olah, C.; Cammarata, N.; Schubert, L.; Goh, G.; Petrov, M.; and Carter, S. 2020. Zoom in: An introduction to circuits. *Distill*, 5(3): e00024–001.
- Olshausen, B. A.; and Field, D. J. 1997. Sparse coding with an overcomplete basis set: A strategy employed by V1? *Vision Research*, 37: 3311–3325.
- Park, Y.-H.; Kwon, M.; Choi, J.; Jo, J.; and Uh, Y. 2023. Understanding the latent space of diffusion models through the lens of riemannian geometry. *Advances in Neural Information Processing Systems*, 36: 24129–24142.
- Peebles, W. S.; and Xie, S. 2022. Scalable diffusion models with transformers. volume 4172.
- Podell, D.; English, Z.; Lacey, K.; Blattmann, A.; Dockhorn, T.; Müller, J.; Penna, J.; and Rombach, R. 2023. Sdxl: Improving latent diffusion models for high-resolution image synthesis. *arXiv preprint arXiv:2307.01952*.
- Radford, A.; et al. 2021. Learning transferable visual models from natural language supervision. In *International conference on machine learning (ICLR)*, 8748–8763. PMLR.
- Rao, S.; Mahajan, S.; et al. 2024. Discover-then-Name: Task-Agnostic Concept Bottlenecks via Automated Concept Discovery. *arXiv preprint arXiv:2407.14499*.
- Rombach, R.; Blattmann, A.; et al. 2022a. High-resolution image synthesis with latent diffusion models. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 10684–10695.
- Rombach, R.; Blattmann, A.; et al. 2022b. High-Resolution Image Synthesis With Latent Diffusion Models. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 10684–10695.
- Ronneberger, O.; Fischer, P.; and Brox, T. 2015. U-net: Convolutional networks for biomedical image segmentation. In *MICCAI 2015: 18th international conference, Munich, Germany, October 5-9, 2015, proceedings, part III 18*, 234–241. Springer.
- Schramowski, P.; Brack, M.; Deiseroth, B.; and Kersting, K. 2023. Safe latent diffusion: Mitigating inappropriate degeneration in diffusion models. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 22522–22531.
- Sharkey, L.; Braun, D.; and beren. 2022. Taking Features Out of Superposition with Sparse Autoencoders. AI Alignment Forum.
- Sharma, P.; Ding, N.; et al. 2018. Conceptual Captions: A Cleaned, Hypernymed, Image Alt-text Dataset For Automatic Image Captioning. In *Proceedings of ACL*.
- Song, J.; Meng, C.; and Ermon, S. 2020. Denoising diffusion implicit models. *arXiv preprint arXiv:2010.02502*.
- Song, Y.; and Ermon, S. 2019. Generative modeling by estimating gradients of the data distribution. 32.
- Surkov, V.; Wendler, C.; Terekhov, M.; Deschenaux, J.; West, R.; and Gulcehre, C. 2024. Unpacking sdxl turbo: Interpreting text-to-image models with sparse autoencoders. *arXiv preprint arXiv:2410.22366*.
- Tang, R.; Liu, L.; et al. 2022. What the daam: Interpreting stable diffusion using cross attention. *arXiv preprint arXiv:2210.04885*.
- Templeton, A. 2024. *Scaling monosemanticity: Extracting interpretable features from claude 3 sonnet*. Anthropic.
- Tian, J.; Aggarwal, L.; et al. 2024. Diffuse attend and segment: Unsupervised zero-shot segmentation using stable diffusion. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 3554–3563.
- Toker, M.; Orgad, H.; Ventura, M.; Arad, D.; and Belinkov, Y. 2024. Diffusion Lens: Interpreting Text Encoders in Text-to-Image Pipelines. *arXiv preprint arXiv:2403.05846*.
- Wen, T.; Wang, Y.; et al. 2025. Beyond Matryoshka: Revisiting Sparse Coding for Adaptive Representation. In *Forty-second International Conference on Machine Learning*.
- Xin, Y.; Yan, J.; et al. 2025. Lumina-mgpt 2.0: Stand-alone autoregressive image modeling. *arXiv preprint arXiv:2507.17801*.
- Yi, M.; Li, A.; et al. 2024. Towards Understanding the Working Mechanism of Text-to-Image Diffusion Model.
- Yin, Y.; et al. 2025. Towards precise scaling laws for video diffusion transformers. In *Proceedings of the Computer Vision and Pattern Recognition Conference*, 18155–18165.
- Yoon, J.; Yu, S.; et al. 2024. Safree: Training-free and adaptive guard for safe text-to-image and video generation. *arXiv preprint arXiv:2410.12761*.
- Yun, Z.; Chen, Y.; Olshausen, B. A.; and LeCun, Y. 2021. Transformer visualization via dictionary learning: contextualized embedding as a linear superposition of transformer factors. *arXiv preprint arXiv:2103.15949*.
- Zhuo, L.; Du, R.; et al. 2024. Lumina-next: Making lumina-t2x stronger and faster with next-dit.