

Aligning Text to Image in Diffusion Models is Easier Than You Think

Jaa-Yeon Lee^{*1} Byunghye Cha^{*1} Jeongsol Kim² Jong Chul Ye¹

¹Kim Jaechul Graduate School of AI, KAIST

²Department of Bio and Brain Engineering, KAIST

{jaaeyeon, paulcha1025, jeongsol, jong.ye}@kaist.ac.kr

^{*}Equal contribution



Figure 1: **Representative results for image generation and image editing.** SoftREPA provides much improved text-to-image alignment by introducing a negligible size of learnable soft tokens.

Abstract

While recent advancements in generative modeling have significantly improved text-image alignment, some residual misalignment between text and image representations still remains. Some approaches address this issue by fine-tuning models in terms of preference optimization, etc., which require tailored datasets. Orthogonal to these methods, we revisit the challenge from the perspective of representation alignment—an approach that has gained popularity with the success of REpresentation Alignment (REPA) [46]. We first argue that conventional text-to-image (T2I) diffusion models, typically trained on paired image and text data (i.e., positive pairs) by minimizing score matching or flow matching losses, is suboptimal from the standpoint of representation alignment. Instead, a better

alignment can be achieved through contrastive learning that leverages existing dataset as both positive and negative pairs. To enable efficient alignment with pretrained models, we propose *SoftREPA*—a lightweight contrastive fine-tuning strategy that leverages soft text tokens for representation alignment. This approach improves alignment with minimal computational overhead by adding fewer than 1M trainable parameters to the pretrained model. Our theoretical analysis demonstrates that our method explicitly increases the mutual information between text and image representations, leading to enhanced semantic consistency. Experimental results across text-to-image generation and text-guided image editing tasks validate the effectiveness of our approach in improving the semantic consistency of T2I generative models. Project Page: <https://softrepa.github.io/>.

1 Introduction

Achieving effective alignment between modalities is essential in multimodal generative modeling. Specifically, latent diffusion models [33] enable various conditioning mechanisms, such as text, during the generation process for the semantic alignment between the two modalities. For UNet-based denoisers [33, 29, 31], image representations are updated to align with fixed text representations encoded by pre-trained CLIP models through cross-attention layers. In contrast, transformer-based denoisers [30, 8] jointly update both text and image representations by concatenating them and processing them through self-attention layers. However, there is still room for further improvement, and to mitigate the remaining misalignment between text and image representations, a more effective representation learning approach is necessary.

Recent advancements in representation alignment for Diffusion Transformers (DiT) [30] have notably enhanced their ability to learn semantically meaningful internal representations [46, 11]. Notably, REPA [46] showed that aligning the internal representations of DiT with an external pre-trained visual encoder during training significantly improves both discriminative and generative performance. To further enhance text-image alignment in text-to-image (T2I) generative models, we leverage these ideas for aligning DiT's internal representations.

Most vision-language generative foundation models have focused on improving multimodal alignment through architectural modifications. These include MLP-based projection layers (e.g., LLaVA-style visual instruction tuning) [25, 26], cross-attention-based fusion mechanisms (e.g., Flamingo, Stable Diffusion 1.5) [2?], and early fusion models that integrate text and vision features within a shared representation space (e.g., Chameleon, Stable Diffusion 3, FLUX) [37, 8]. In fact, many of these approaches have been largely adopted in T2I generative models. However, we explore an efficient yet effective way of further enhancing the representation alignment for given pre-trained models.

Specifically, we focus on the contrastive learning framework, a widely adopted strategy for multimodal alignment in vision-language representation learning, as seen in models such as CLIP [32], ALBEF [21], BLIP [22], and ALIGN [12]. One of the key contributions of this work is the introduction of *soft text tokens*, which are optimized via contrastive image-text training. These soft tokens allow the model to dynamically adapt its text representations, improving alignment with generated images without requiring full model fine-tuning. This simple yet effective approach significantly enhances text-to-image alignment in both text-to-image generation and text-guided image editing tasks. Notably, the method is highly flexible and can be seamlessly integrated with any pretrained text-to-image (T2I) generative model. From a theoretical standpoint, we show that this approach explicitly increases the mutual information between text and image representations, resulting in improved semantic consistency across modalities. We refer to our method as *SoftREPA*, short for Soft REPresentation Alignment. Our contributions can be summarized as follows.

- We propose SoftREPA, a novel text-image representation alignment method that leverages a lightweight fine-tuning strategy with soft text tokens. This approach improves text-image alignment while adding fewer than 1M additional parameters, ensuring efficiency with minimal computational overhead.
- SoftREPA is simple yet flexible so that it can be used with any pretrained T2I generative models to improve performance of image generation, editing, etc.

- We show that our method explicitly increases mutual information between image and text, leading to better semantic consistency in multi-modal representations.

2 Preliminaries

Flow Models Suppose that we have access to samples from target distribution $X_0 \sim q$ and source distribution $X_1 \sim p$. The goal of flow model is to generate X_0 starting from X_1 . Specifically, we define a velocity field $v_t(\mathbf{x})$ of a flow $\psi_t(\mathbf{x}) : [0, 1] \times \mathbb{R}^d \rightarrow \mathbb{R}^d$ that satisfies $\psi_t(X_0) = X_t$ and $\psi_1(X_0) = X_1$. Here, the ψ_t is uniquely characterized by a flow ODE:

$$d\psi_t(\mathbf{x}) = v_t(\psi_t(\mathbf{x}))dt \quad (1)$$

where the flow velocity v_t is fitted to the parameterized neural network $v_{t,\theta}$ via flow matching:

$$\mathcal{L}_{FM} = \mathbb{E}_{t \in [0,1], \mathbf{x}_t \sim p_t} \|v_t(\mathbf{x}_t) - v_{t,\theta}(\mathbf{x}_t)\|^2. \quad (2)$$

However, this is computationally expensive to solve due to the integration with respect to X_0 , so the authors in [24] proposed a conditional flow matching that has the same gradient with the original objective function:

$$\mathcal{L}_{CFM} = \mathbb{E}_{t \in [0,1], \mathbf{x}_0 \sim q} \|v_t(\mathbf{x}_t|\mathbf{x}_0) - v_{t,\theta}(\mathbf{x}_t)\|^2, \quad (3)$$

where $v_t(\mathbf{x}_t|\mathbf{x}_0)$ defines a conditional flow $\psi_t(\mathbf{x}_t|\mathbf{x}_0)$ satisfying $\psi_t(\mathbf{x}_1|\mathbf{x}_0) = \mathbf{x}_t$. In particular, the linear conditional flow defines the flow as $\mathbf{x}_t = \psi_t(\mathbf{x}_1|\mathbf{x}_0) = (1-t)\mathbf{x}_0 + t\mathbf{x}_1$. Then, the conditional velocity field is given by $v_t(\mathbf{x}_t|\mathbf{x}_0) = \dot{\psi}_t(\psi_t^{-1}(\mathbf{x}_t|\mathbf{x}_0)|\mathbf{x}_0) = \mathbf{x}_1 - \mathbf{x}_0$ [24]. Thus, the conditional flow matching loss is defined as

$$\mathbb{E}_{t \in [0,1], \mathbf{x}_0, \mathbf{x}_1 \sim \pi_{0,1}} \|(\mathbf{x}_1 - \mathbf{x}_0) - v_{t,\theta}(\mathbf{x}_t)\|^2. \quad (4)$$

Considering a marginal velocity field,

$$\begin{aligned} v_t(\mathbf{x}_t) &= \int v_t(\mathbf{x}_t|\mathbf{x}_0)p(\mathbf{x}_0|\mathbf{x}_t)d\mathbf{x}_0 = \mathbb{E}[v_t(\mathbf{x}_t|\mathbf{x}_0)|\mathbf{x}_t] \\ &= \mathbb{E}[\mathbf{x}_1 - \mathbf{x}_0|\mathbf{x}_t] = \mathbb{E}[\mathbf{x}_1|\mathbf{x}_t] - \mathbb{E}[\mathbf{x}_0|\mathbf{x}_t], \end{aligned} \quad (5)$$

the generation process via flow ODE is

$$d\mathbf{x}_t = v_t(\mathbf{x}_t)dt = (\mathbb{E}[\mathbf{x}_1|\mathbf{x}_t] - \mathbb{E}[\mathbf{x}_0|\mathbf{x}_t])dt. \quad (6)$$

Accordingly,

$$d\mathbf{x}_t = \left(\mathbb{E} \left[\frac{\mathbf{x}_t - (1-t)\mathbf{x}_0}{t} | \mathbf{x}_t \right] - \mathbb{E}[\mathbf{x}_0|\mathbf{x}_t] \right) dt = \frac{\mathbf{x}_t - \mathbb{E}[\mathbf{x}_0|\mathbf{x}_t]}{t} dt \quad (7)$$

which is the probability-flow ODE (PF-ODE) of DDIM [35, 15] when $\mathbf{x}_1 \sim p := \mathcal{N}(0, \mathbf{I}_d)$. In other words, flow models and score-based models can be used interchangeably.

Contrastive Representation Learning Contrastive learning aims to maximize the similarity between semantically related text-image pairs while pushing apart unrelated pairs in a shared representation space [6, 32]. Formally, given a batch of N image-text pairs $(\mathbf{I}_i, \mathbf{T}_i)_{i=1}^N$, contrastive learning optimizes a contrastive loss function, typically a variant of the InfoNCE loss:

$$\mathcal{L}_{CLIP} = -\frac{1}{N} \sum_{i=1}^N \log \frac{\exp(\text{sim}(\mathbf{I}_i, \mathbf{T}_i)/\tau)}{\sum_{j=1}^N \exp(\text{sim}(\mathbf{I}_i, \mathbf{T}_j)/\tau)} \quad (8)$$

where $\text{sim}(\mathbf{I}_i, \mathbf{T}_i)$ is a similarity metric such as cosine similarity or inner product between image and text embeddings, and τ denotes a temperature parameter that controls the distribution sharpness. This contrastive loss encourages image and text representations to form a joint multimodal representation space, where aligned pairs are close together, and unaligned pairs are separated. Several fundamental models have successfully applied contrastive learning for multi-modal representation learning [32, 12, 21, 22] and self-supervised learning [6].




Figure 2: **Network architecture and algorithmic concept of SoftREPA.** (a) Learnable soft tokens of each layer are prepended to the text features across the upper layers. (b) The soft tokens are optimized to contrastively match the score with positively conditioned predicted noise while repelling the score from negatively conditioned predicted noise. This process implicitly sharpens the joint probability distribution of images and text by reducing the log probability of negatively paired conditions.

Despite the success of contrastive learning in representation learning, its application in generative models remains underexplored due to key challenges. First, there is a fundamental mismatch between discriminative and generative features; conventional contrastive learning optimizes a representation space for discriminative tasks, whereas generative models focus on realistic sample synthesis. Second, text-image alignment in diffusion-based generative models is often implicit, relying on denoising objectives rather than explicit representation learning.

To bridge this gap, we propose a contrastive learning framework that aligns both representation learning and generative objectives, ensuring effective text-image representation alignment while preserving generative quality.

3 SoftREPA

3.1 Contrastive T2I Alignment Loss

Let $\{(\mathbf{x}^{(i)}, \mathbf{y}^{(i)})\}_{i=1}^n$ denotes the matched image and text pairs as a training data set. Then, contrastive text-to-image alignment in T2I model can be generally formulated as follows:

$$\mathcal{L} = -\frac{1}{n} \sum_{i=1}^n \log \frac{\exp(l(\mathbf{x}^{(i)}, \mathbf{y}^{(i)}))}{\sum_j \exp(l(\mathbf{x}^{(i)}, \mathbf{y}^{(j)}))} \quad (9)$$

where $l(\cdot, \cdot)$ represents the similarity measure. It is important to note that unlike the standard T2I model training that only consider positive pairs, i.e. $(\mathbf{x}^{(i)}, \mathbf{y}^{(i)})$, our SoftREPA employs the contrastive T2I alignment that additionally considers negative image and text pairs $(\mathbf{x}^{(i)}, \mathbf{y}^{(j)})$, $i \neq j$. This significantly improves the alignment performance.

Furthermore, we define $l(\mathbf{x}^{(i)}, \mathbf{y}^{(j)})$ as the logit from the variation of the denoising score matching loss for the case of T2I diffusion models [41].

$$l(\mathbf{x}, \mathbf{y}) = e^{-\mathbb{E}_{t, \epsilon} [\|\epsilon_{\theta}(\mathbf{x}_t, t, \mathbf{y}) - \epsilon\|^2 / \tau(t)]} \quad (10)$$

Similarly, in conditional flow matching, the logit $l(\mathbf{x}, \mathbf{y})$ can be formalized as:

$$l(\mathbf{x}, \mathbf{y}) = e^{-\mathbb{E}_{t, \epsilon} [\|v_{\theta}(\mathbf{x}_t, t, \mathbf{y}) - (\epsilon - \mathbf{x}_0)\|^2 / \tau(t)]} \quad (11)$$

Here, $\tau(t)$ represents time scheduling parameter. One might consider using the denoising score matching loss, $-\mathbb{E}_{t,\epsilon}[\|\epsilon_\theta(\mathbf{x}_t, t, \mathbf{y}) - \epsilon\|^2]$, as a similarity measure instead of relying on its logit value. However, our experiments revealed that this formulation can introduce instability during training due to the unbounded nature of the similarity measure. To address this, SoftREPA introduce an exponential function to constrain the logit values, effectively stabilizing the training process.

Intuitively, as illustrated in Fig. 2(b), training a diffusion model with this loss encourage the predicted noise to align with the true noise when conditioned on a matching text description. Conversely, when conditioned on a mismatched text, the predicted noise deviates further from the true noise. This sharpens the conditional probability distribution of the trained model and increases the distinction between different text conditions, which can enhance image-text alignment in generated images.

3.2 The Soft Tokens

To distill the contrastive score matching objective above, we introduce a learnable soft token $\mathbf{s}$, which is the only trainable parameter while pretrained model remains frozen. In particular, as illustrated in fig. 2, we introduce a set of learnable soft tokens that vary between layers and are indexed by time t . The soft token at layer k is defined using an embedding function with the number of soft tokens as m :

$$\mathbf{s}^{(k,t)} = \text{Embedding}(k, t) \in \mathbb{R}^{m \times d} \quad (12)$$

At the layer k , the text representation is updated by concatenating the time-indexed soft token with text hidden representation from the previous layer: $\hat{\mathbf{H}}_{\text{text}}^{(k-1,t)} = [\mathbf{s}^{(k,t)}; \mathbf{H}_{\text{text}}^{(k-1,t)}] \in \mathbb{R}^{(m+n) \times d}$. The soft token is prepended on text features across the layers, leading to a modified denoising function: $v_\theta(\mathbf{x}_t, t, \mathbf{y}, \mathbf{s})$. By adopting the soft tokens, we can adjust the image and text representations to get better semantic alignment. Additionally, for computational efficiency, we approximate the expectation within $l(\mathbf{x}, \mathbf{y})$ using a single Monte Carlo sample, effectively replacing it as

$$\tilde{l}(\mathbf{x}, \mathbf{y}, \mathbf{s}) = e^{-\|v_\theta(\mathbf{x}_t, t, \mathbf{y}, \mathbf{s}) - (\epsilon - \mathbf{x}_0)\|^2 / \tau(t)} \quad (13)$$

which can stabilize the computation using the same ϵ and t in the same batch. The resulting contrastive loss function with soft tokens can be defined as follows:

$$\mathcal{L}_{\text{SoftREPA}}(\mathbf{s}) = -\mathbb{E}_{(\mathbf{x}, \mathbf{y}) \sim p_{\text{data}}, t \sim U(0,1), \epsilon \sim \mathcal{N}(\mathbf{0}, \mathbf{I})} \log \left(\frac{\exp(\tilde{l}(\mathbf{x}, \mathbf{y}, \mathbf{s}))}{\sum_j \exp(\tilde{l}(\mathbf{x}, \mathbf{y}^{(j)}, \mathbf{s}))} \right) \quad (14)$$

Then, the learnable token $\mathbf{s}$ is optimized by minimizing $\mathcal{L}_{\text{SoftREPA}}(\mathbf{s})$. During inference, the trained soft tokens are leveraged with fixed text representation from the text encoder across layers and timesteps. The detailed procedure during the image generation process based on MM-DiT architecture is described in algorithm 1 from appendix B.

3.3 Relationship with Mutual Information

Although the relationship between contrastive loss and mutual information has been previously studied, we revisit it here to provide a more complete theoretical perspective. Additionally, we emphasize the critical role of the logit formulation in the score matching loss, highlighting its importance in ensuring stable and effective optimization.

According to [18], pointwise mutual information (PMI) over image $\mathbf{x}$ and text $\mathbf{y}$ pairs can be defined as follows:

$$i(\mathbf{x}, \mathbf{y}) = \log \frac{p_\theta(\mathbf{x}|\mathbf{y})}{p_\theta(\mathbf{x})} = \log \frac{p_\theta(\mathbf{x}|\mathbf{y})}{\mathbb{E}_{p(\mathbf{c})}[p_\theta(\mathbf{x}|\mathbf{c})]} \quad (15)$$

where the second equality follows from the definition of conditional expectation. Song et al. [36] and Kong et al. [18] further show that assuming the diffusion model is an optimal denoiser, the conditional likelihood $p_\theta(\mathbf{x}|\mathbf{y})$ can be described as (see appendix A):

$$p_\theta(\mathbf{x}|\mathbf{y}) = \exp(\hat{l}(\mathbf{x}, \mathbf{y})) \quad \text{where} \quad \hat{l}(\mathbf{x}, \mathbf{y}) = -\frac{1}{2} \int_0^T \lambda(t) \mathbb{E}[\|\epsilon_\theta(\mathbf{x}_t, t, \mathbf{y}) - \epsilon\|^2] dt + C. \quad (16)$$

By approximating the denominator of PMI via Monte Carlo sampling, we arrive at:

$$i(\mathbf{x}, \mathbf{y}) \approx \log \frac{\exp(\hat{l}(\mathbf{x}, \mathbf{y}))}{\frac{1}{N} \sum_c \exp(\hat{l}(\mathbf{x}, \mathbf{c}))} \quad (17)$$

COCO val5K							Efficiency	
Model	Human Preference		Text Alignment		Image Quality		PM (GB)	Latency (sec/image)
	ImageReward↑	PickScore↑	CLIP↑	HPS↑	FID↓	LPIPS↓		
SD1.5	17.72	21.47	26.4	25.08	24.59	43.80	2.621	1.526
SD1.5 +Ours	32.89	21.50	27.33	25.18	23.43	43.38	2.621	1.547
SDXL	75.06	22.38	26.76	27.35	24.69	42.05	10.45	4.059
SDXL +Ours	85.29	22.62	26.80	28.30	26.04	42.39	10.45	4.060
SD3	94.27	22.54	26.30	28.09	31.59	42.43	18.77	4.637
SD3 +Ours	108.5	22.55	26.91	28.91	36.21	42.88	18.77	4.695

GenEval							
Model	Mean↑	Single↑	Two↑	Counting↑	Colors↑	Position↑	Color Attribution↑
SD3	0.68	0.99	0.86	0.56	0.85	0.27	0.55
CaPO [20]	0.71	0.99	0.87	0.63	0.86	0.31	0.59
RankDPO [16]	0.74	1.00	0.90	0.72	0.87	0.31	0.66
SD3 + Ours	0.70	1.00	0.95	0.29	0.92	0.34	0.68

Table 1: Quantitative evaluation of T2I generation with learnable soft tokens on SD1.5, SDXL, and SD3. Generation quality is evaluated on the COCO-val 5K [23] and GenEval [9] benchmark. Peak GPU memory(PM) usage and per-image latency are measured to assess computational efficiency. ImageReward, CLIP, HPS, and LPIPS are scaled by $\times 10^2$.

This formulation closely resembles a cross-entropy loss, where the conditional log-likelihood $\hat{l}(\mathbf{x}, \mathbf{y})$ can be interpreted as logits. Now, from the fact that mutual information is the expectation of the pointwise mutual information, we have

$$I(X, Y) = \frac{1}{n} \sum_{i=1}^n \log \frac{\exp(\hat{l}(\mathbf{x}^{(i)}, \mathbf{y}^{(i)}))}{\sum_j \exp(\hat{l}(\mathbf{x}^{(i)}, \mathbf{y}^{(j)}))} + D \quad (18)$$

where $\mathbf{x}^{(i)}$ represents the i -th sample in the data, n denotes the number of samples in the data, and D indicates the constant. Notably, Eq. (18) is very similar in form to the contrastive objective in Eq. (9), where the logit is the negative of the diffusion loss. This implies that minimizing our contrastive learning objective is closely related to maximizing the mutual information between image-text pairs under the diffusion model.

4 Experiments

Implementation details. In our experiments, to confirm the flexibility of SoftREPA, we utilize various open-source T2I diffusion models, including Stable Diffusion 1.5, Stable Diffusion XL, and Stable Diffusion 3 to evaluate text and image alignment in both image generation and text-based image editing tasks. For the training of softs token, the batch size was set to 16 with less than 30,000 iterations using two A100 GPUs. For Stable Diffusion 3, we set the length of the soft token to 4 and used 5 layers to attach the soft tokens for further experiments. In the case of Stable Diffusion 1.5 and Stable Diffusion XL, soft tokens are applied only to the Down or Middle block layers of the UNet. Refer to further implementation details in appendix B.

Text to Image Generation We conducted text-to-image generation experiments on SD1.5, SDXL, and SD3, training soft tokens using the text-image paired COCO dataset [23]. The implementation details are provided in appendix B. To evaluate the generated images, we assessed human preference scores [45, 17], text-image alignment [48, 44], and image quality [47, 34] on COCO-val 5K dataset. Furthermore, we conducted evaluation on GenEval [9], compared with RankDPO [16], which is fine-tuned via preference optimization. As shown in table 1, in COCO dataset, the proposed method generally outperforms baseline approaches in both diffusion and rectified flow models regardless of the model architecture. Additionally, in GenEval, SD3 with SoftREPA mostly outperforms RankDPO [16]. The drop in the counting metric can be primarily attributed to the tendency of diffusion models to interpret text alignment in a way that encourages the generation of multiple instances of the referenced object. Further analysis and potential mitigation strategies are discussed in appendix D. In fig. 3, detailed textual descriptions are more accurately reflected in the generated images, as demonstrated in the qualitative results.

	Inversion	Method	Human Preference		Text Aligmnnet			Structure	Background Preservation		
			Image -Reward $\uparrow$	Pick -Score $\uparrow$	CLIP/ Edited $\uparrow$	CLIP/ Whole $\uparrow$	HPS $\uparrow$	Distance $\downarrow$	PSNR $\uparrow$	LPIPS/ Whole $\downarrow$	SSIM $\uparrow$
PIEBench	ddim	PnP	32.29	21.53	22.56	25.59	26.11	27.32	22.31	14.48	79.58
	ddim	PnP (Ours)	36.78	21.59	22.70	25.78	26.35	27.43	22.27	14.55	79.36
	direct	PnP	40.85	21.65	22.68	25.64	26.60	23.33	22.46	13.44	80.22
	direct	PnP (Ours)	43.32	21.70	22.80	25.83	26.76	23.41	22.40	13.53	79.93
	ddim	MasaCtrl	-13.94	21.03	21.20	24.18	23.59	27.63	22.31	14.59	80.41
	ddim	MasaCtrl (Ours)	-12.76	21.05	21.27	24.44	23.69	27.36	22.27	14.5	80.27
	direct	MasaCtrl	4.77	21.39	21.47	24.54	24.92	23.61	22.82	12.21	82.02
	direct	MasaCtrl (Ours)	4.48	21.40	21.49	24.69	24.93	23.32	22.77	12.19	81.85
	-	FlowEdit	87.70	22.16	23.19	26.72	28.04	25.48	24.37	12.55	88.58
	-	FlowEdit (Ours)	102.24	22.31	23.60	27.19	28.58	24.07	24.99	12.60	88.73
DIV2K	ddim	PnP	-9.72	21.10	-	26.07	24.76	32.39	-	18.89	-
	ddim	PnP (Ours)	-7.75	21.15	-	26.18	24.88	31.95	-	18.88	-
	direct	PnP	-6.86	21.17	-	25.96	25.09	29.15	-	17.96	-
	direct	PnP (Ours)	-5.56	21.22	-	26.15	25.17	28.82	-	17.99	-
	ddim	MasaCtrl	-60.30	20.55	-	23.37	21.44	29.70	-	17.40	-
	ddim	MasaCtrl (Ours)	-57.41	20.59	-	23.46	21.64	28.79	-	17.03	-
	direct	MasaCtrl	-42.33	20.80	-	23.65	22.7	71.16	-	31.41	-
	direct	MasaCtrl (Ours)	-41.20	20.81	-	23.68	22.77	71.11	-	31.26	-
	-	FlowEdit	38.08	21.74	-	26.05	25.53	39.68	-	15.48	-
	-	FlowEdit (Ours)	46.68	21.88	-	26.39	26.03	35.66	-	14.98	-
C2D	-	FlowEdit	93.70	20.72	-	22.51	26.68	33.41	-	19.92	-
	-	FlowEdit (Ours)	114.4	21.22	-	23.43	28.35	34.05	-	20.47	-

Table 2: Quantitative evaluation of image editing performance with the use of soft tokens on PnP [38], MasaCtrl [4], Direct inversion [14], and FlowEdit [19]. FlowEdit is based on SD3, and others are based on SD1.5 architecture. The editing methods are evaluated on PIEBench [14], DIV2K [1], and Cat2Dog(C2D). ImageReward, CLIP, HPS, LPIPS, and SSIM are scaled by $\times 10^2$ and Distance is scaled by $\times 10^3$.




Figure 3: The qualitative results of text-to-image generation comparing SD3 and SD3 with proposed method. The given text is from COCO and Pixart dataset.




Figure 4: The qualitative results of text guided image editing comparing on SD3 with the proposed method. The FlowEdit [19] is used as the editing method for both baseline and SoftREPA.




Figure 5: Quantitative comparison of baseline and SoftREPA on DIV2K and Cat2Dog editing using CLIP Score and LPIPS with various CFG scales.




Figure 6: Ablation study on the number of layers and token length for soft tokens. Metrics are normalized to: ImageReward [0.7, 1.1], CLIP score [0.25, 0.28], and LPIPS [0.42, 0.44].

Memory and Speed Comparison In table 1, we report the peak GPU memory usage (PM) and per-image latency for models with and without soft tokens, including SD1.5, SDXL, and SD3. All models were evaluated in float16 precision. The use of soft tokens introduces no change in peak memory and results in only a negligible increase in latency.

Text Guided Image Editing To further demonstrate the improved text-image alignment achieved by our method, we conducted text-guided image editing experiments using PnP [38], MasaCtrl [4] with DDIM [35], Direct Inversion [14], and FlowEdit [19]. FlowEdit employs SD3, while the other methods use SD1.5. For evaluation, we used PIEBench [14], DIV2K [1], and Cat2Dog datasets. For DIV2K editing dataset, we selected 800 high-quality images and generated detailed source and target captions using LLaVA [25]; dataset construction details are provided in appendix C. In table 2,

Model	Human Preference		Text Alignment		Image Quality	
	ImageReward↑	PickScore↑	CLIP↑	HPS↑	FID↓	LPIPS↓
SD3	90.89	22.50	26.26	27.97	56.87	42.47
SD3 + LoRA + Contrastive loss	97.72	22.57	26.42	28.44	70.33	42.71
SD3 + Soft tokens + Contrastive loss (Ours)	106.32	22.49	26.92	28.70	60.76	42.99

Table 3: Quantitative comparison of T2I generation (COCO-val 1K [23]) between soft tokens and LoRA on SD3. ImageReward, CLIP, HPS, and LPIPS are scaled by $\times 10^2$.

the quantitative results of human preference, text alignment, and structure distance show consistent enhancements with the use of soft tokens without compromising structure and background fidelity.

Considering that FlowEdit performs text-based image editing, where performance heavily depends on text-image alignment, this result indicates improved efficiency in modeling the joint text-image distribution using the proposed method. Furthermore, our method with soft tokens achieves more efficient editing by reducing the number of editing steps to 30, compared to 33 in FlowEdit. As shown in fig. 5, incorporating soft tokens across various CFG scales consistently enhances text-image alignment and image quality. Figure 4 further demonstrates that performance has been improved on both single and multi-concept editing compared to vanilla SD3. Additional results and details are provided in appendix F.

Ablation Study We conducted an ablation study on SD3 using the COCO val dataset [23] to evaluate the effects of the number of layers using soft text tokens and token length on generation quality (fig. 6). A value of 0 corresponds to vanilla SD3, and the number of layers is counted from the first. Comparisons were made using either a fixed token length of 8 or a fixed depth of 5 layers. Our findings indicate that incorporating soft tokens beyond layer 7 severely degrades image generation quality. The optimal performance, aligned with human preferences (ImageReward [45]), was observed when soft tokens were applied within layers 2 ~ 5. This aligns with the architecture, where early layers handle text-image alignment and later ones enhance fidelity [27]. Notably, CLIP scores continued to improve as more layers used soft tokens, indicating stronger adherence to text prompts, while overall image quality remained stable. Supporting qualitative results are provided in appendix G. For token length, using more than 8 tokens negatively impacted quality, suggesting overfitting to the fine-tuning data. Shorter lengths (1–4 tokens) preserved perceptual similarity (LPIPS), while moderate lengths (4–8 tokens) significantly improved text-image alignment (CLIP and ImageReward). Qualitative examples are shown in appendix G.

Comparison with LoRA To further assess the effectiveness of our design, we examined whether contrastive learning alone could achieve comparable improvements. Specifically, we conducted an ablation study replacing the soft tokens with LoRA. As shown in table 3, the combination of soft tokens and contrastive loss consistently outperforms all other variants across evaluation metrics. This result demonstrates the strong synergy between soft tokens and contrastive learning in enhancing text-image alignment. These findings validate the core motivation of SoftREPA—to explore the text feature space for improved alignment between textual and visual representations. Additional comparisons with different LoRA ranks and layer configurations are presented in appendix H.

Complementarity with Diffusion RL methods We examined the complementarity between SoftREPA and diffusion reinforcement learning (RL) methods. Since these approaches are not mutually exclusive, we tested whether SoftREPA can enhance diffusion RL frameworks. Specifically, we integrated pretrained soft tokens from SoftREPA into two representative methods—Diffusion-DPO [43], which leverages preference optimization, and DDPO [3], which applies policy optimization with an external reward—using a plug-and-play approach without additional training. As shown in table 4, incorporating SoftREPA’s soft tokens consistently improves performance. This suggests strong complementarity while diffusion RL methods align outputs via preference or reward signals, SoftREPA refines internal text-image representations through contrastive learning. Combining the two effectively unites explicit preference alignment with robust representation learning to enhance generation quality.

Method	Human Preference		Text Alignment		Image Quality	
	ImageReward $\uparrow$	PickScore $\uparrow$	CLIP $\uparrow$	HPS $\uparrow$	FID $\downarrow$	LPIPS $\downarrow$
Diffusion-DPO	22.26	21.68	26.51	25.60	54.19	43.96
Diffusion-DPO + Ours	33.10	21.64	27.41	25.55	53.72	43.64
DDPO	-8.83	21.20	26.21	23.45	54.83	43.85
DDPO + Ours	8.07	21.25	27.13	23.75	54.18	43.59

Table 4: Complementarity of SoftREPA with diffusion RL methods (Diffusion-DPO [43], DDPO [3]). Soft tokens are integrated in a plug-and-play manner during inference. Generation quality is evaluated on the COCO-val 1K dataset [23]. ImageReward, CLIP, HPS, and LPIPS are scaled by $\times 10^2$.

5 Related Works

Several methods have been proposed to address the text-image misalignment. Training free approach such as CFG++ [7] addresses this problem as the off-manifold phenomenon of conventional CFG and reformulates it to a manifold-constrained CFG to get better text-to-image alignment. Other training-free approach, attend-and-excite [5], adjust attention activations to enhance text-image consistency. DOODL [42] applies guidance to end-to-end manner. Meanwhile, training-based approaches such as Diffusion-DPO [43] introduce alignment losses during training to both enforce stronger text-conditioning and visual realism. Following Diffusion-DPO, denoising diffusion policy optimization (DDPO) [3] enhances the alignment using feedback from the pre-trained vision-language model. RankDPO [16] presents fully synthetic, ranking-based preference datasets generated by reward models instead of humans. However, these methods require either extensive full-finetuning or an additional annotation dataset. In this paper, we use the prior knowledge of diffusion model itself to get better text and image alignment, proposing a novel SoftREPA on soft tokens.

To enable efficient prompt-level fine-tuning, VPT [13] introduced learnable tokens for vision transformers while keeping the backbone frozen. CoOp [50] and CoCoOp [49] proposed prompt learning methods for CLIP-like vision-language models by optimizing prompts to distinguish between a predefined set of classes in image classification tasks. In recent works, MinorityPrompt [39] and Optical-Prompt [28] learned sample-specific prompts to improve diversity and temporal coherence during sampling steps, respectively. In contrast to these approaches, which typically tailor prompts to individual samples or specific class sets, we propose learning general-purpose, sample-independent soft tokens. These tokens are optimized over a distribution of text-image pairs, $(\mathbf{X}, \mathbf{Y})$, effectively guiding the propagation of text features to better align with the joint text-image distribution.

6 Conclusion

In this work, we introduce SoftREPA, a novel approach for improving text-image alignment in generative models. Building on the success of prior representation alignment techniques, we extend the idea of internal representation alignment to the text-image alignment without relying on external encoders. Our method effectively integrates contrastive learning and score matching loss to enhance text-image representation learning while preserving the generative capabilities of diffusion models. Also, we introduced soft text tokens, which can effectively adjust the image-text representations through contrastive T2I alignment, making it lightweight and efficient. We demonstrate that our approach enhances text-image alignment in both image generation and editing tasks while maintaining a minimal parameter overhead. Our experiments span various models, including Stable Diffusion 1.5, XL, and 3.

Limitation and Potential Negative Impacts. While soft tokens effectively enhance text-image alignment, they may risk overemphasizing textual guidance at the expense of faithfully preserving the user’s intent in the prompt. Future work could explore regularization techniques to mitigate this issue. Additionally, this work does not address potential risks such as data poisoning or adversarial attacks during soft token training. Although these concerns are beyond the scope of this paper, they merit consideration in future research. Moreover, our method inherits any potential negative impacts of the underlying backbone models.

Acknowledgments and Disclosure of Funding

This work was supported by the Institute of Information & communications Technology Planning & Evaluation (IITP) grant funded by the Korean government(MSIT) (No. RS-2024-00457882, AI Research Hub Project). This work was supported by the National Research Foundation of Korea under Grant RS-2024-00336454.

References

- [1] Eirikur Agustsson and Radu Timofte. Ntire 2017 challenge on single image super-resolution: Dataset and study. In *Proceedings of the IEEE conference on computer vision and pattern recognition workshops*, pages 126–135, 2017.
- [2] Jean-Baptiste Alayrac, Jeff Donahue, Pauline Luc, Antoine Miech, Iain Barr, Yana Hasson, Karel Lenc, Arthur Mensch, Katherine Millican, Malcolm Reynolds, et al. Flamingo: a visual language model for few-shot learning. *Advances in neural information processing systems*, 35:23716–23736, 2022.
- [3] Kevin Black, Michael Janner, Yilun Du, Ilya Kostrikov, and Sergey Levine. Training diffusion models with reinforcement learning. In *International Conference on Learning Representations*, 2024.
- [4] Mingdeng Cao, Xintao Wang, Zhongang Qi, Ying Shan, Xiaoju Qie, and Yinqiang Zheng. Masactrl: Tuning-free mutual self-attention control for consistent image synthesis and editing. In *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)*, pages 22560–22570, October 2023.
- [5] Hila Chefer, Yuval Alaluf, Yael Vinker, Lior Wolf, and Daniel Cohen-Or. Attend-and-excite: Attention-based semantic guidance for text-to-image diffusion models. *ACM transactions on Graphics (TOG)*, 42(4): 1–10, 2023.
- [6] Ting Chen, Simon Kornblith, Mohammad Norouzi, and Geoffrey Hinton. A simple framework for contrastive learning of visual representations. *arXiv preprint arXiv:2002.05709*, 2020.
- [7] Hyungjin Chung, Jeongsol Kim, Geon Yeong Park, Hyelin Nam, and Jong Chul Ye. Cfg++: Manifold-constrained classifier free guidance for diffusion models. *arXiv preprint arXiv:2406.08070*, 2024.
- [8] Patrick Esser, Sumith Kulal, Andreas Blattmann, Rahim Entezari, Jonas Müller, Harry Saini, Yam Levi, Dominik Lorenz, Axel Sauer, Frederic Boesel, et al. Scaling rectified flow transformers for high-resolution image synthesis. In *Forty-first international conference on machine learning*, 2024.
- [9] Dhruva Ghosh, Hannaneh Hajishirzi, and Ludwig Schmidt. Geneval: An object-focused framework for evaluating text-to-image alignment. *Advances in Neural Information Processing Systems*, 36:52132–52152, 2023.
- [10] Aaron Grattafiori, Abhimanyu Dubey, Abhinav Jauhri, Abhinav Pandey, Abhishek Kadian, Ahmad Al-Dahle, Aiesha Letman, Akhil Mathur, Alan Schelten, Alex Vaughan, et al. The llama 3 herd of models. *arXiv preprint arXiv:2407.21783*, 2024.
- [11] Hyeonho Jeong, Chun-Hao Paul Huang, Jong Chul Ye, Niloy Mitra, and Duygu Ceylan. Track4gen: Teaching video diffusion models to track points improves video generation. *arXiv preprint arXiv:2412.06016*, 2024.
- [12] Chao Jia, Yinfei Yang, Ye Xia, Yi-Ting Chen, Zarana Parekh, Hieu Pham, Quoc Le, Yun-Hsuan Sung, Zhen Li, and Tom Duerig. Scaling up visual and vision-language representation learning with noisy text supervision. In *International conference on machine learning*, pages 4904–4916. PMLR, 2021.
- [13] Menglin Jia, Luming Tang, Bor-Chun Chen, Claire Cardie, Serge Belongie, Bharath Hariharan, and Ser-Nam Lim. Visual prompt tuning. In *European conference on computer vision*, pages 709–727. Springer, 2022.
- [14] Xuan Ju, Ailing Zeng, Yuxuan Bian, Shaoteng Liu, and Qiang Xu. Pnp inversion: Boosting diffusion-based editing with 3 lines of code. *International Conference on Learning Representations (ICLR)*, 2024.
- [15] Tero Karras, Miika Aittala, Timo Aila, and Samuli Laine. Elucidating the design space of diffusion-based generative models. In *Proc. NeurIPS*, 2022.
- [16] Shyamgopal Karthik, Huseyin Coskun, Zeynep Akata, Sergey Tulyakov, Jian Ren, and Anil Kag. Scalable ranked preference optimization for text-to-image generation. *arXiv preprint*, 2024.

- [17] Yuval Kirstain, Adam Polyak, Uriel Singer, Shahbuland Matiana, Joe Penna, and Omer Levy. Pick-a-pic: An open dataset of user preferences for text-to-image generation. 2023.
- [18] Xianghao Kong, Ollie Liu, Han Li, Dani Yogatama, and Greg Ver Steeg. Interpretable diffusion via information decomposition. *arXiv preprint arXiv:2310.07972*, 2023.
- [19] Vladimir Kulikov, Matan Kleiner, Inbar Huberman-Spiegelglas, and Tomer Michaeli. Flowedit: Inversion-free text-based editing using pre-trained flow models. *arXiv preprint arXiv:2412.08629*, 2024.
- [20] Kyungmin Lee, Xiahong Li, Qifei Wang, Junfeng He, Junjie Ke, Ming-Hsuan Yang, Irfan Essa, Jinwoo Shin, Feng Yang, and Yinxiao Li. Calibrated multi-preference optimization for aligning diffusion models. In *Proceedings of the Computer Vision and Pattern Recognition Conference*, pages 18465–18475, 2025.
- [21] Junnan Li, Ramprasaath Selvaraju, Akhilesh Gotmare, Shafiq Joty, Caiming Xiong, and Steven Chu Hong Hoi. Align before fuse: Vision and language representation learning with momentum distillation. *Advances in neural information processing systems*, 34:9694–9705, 2021.
- [22] Junnan Li, Dongxu Li, Caiming Xiong, and Steven Hoi. Blip: Bootstrapping language-image pre-training for unified vision-language understanding and generation. In *International conference on machine learning*, pages 12888–12900. PMLR, 2022.
- [23] Tsung-Yi Lin, Michael Maire, Serge Belongie, James Hays, Pietro Perona, Deva Ramanan, Piotr Dollár, and C Lawrence Zitnick. Microsoft coco: Common objects in context. In *Computer Vision–ECCV 2014: 13th European Conference, Zurich, Switzerland, September 6–12, 2014, Proceedings, Part V 13*, pages 740–755. Springer, 2014.
- [24] 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. URL <https://openreview.net/forum?id=PqvMRDCJT9t>.
- [25] Haotian Liu, Chunyuan Li, Qingyang Wu, and Yong Jae Lee. Visual instruction tuning. *Advances in neural information processing systems*, 36:34892–34916, 2023.
- [26] Haotian Liu, Chunyuan Li, Yuheng Li, Bo Li, Yuanhan Zhang, Sheng Shen, and Yong Jae Lee. Llava-next: Improved reasoning, ocr, and world knowledge, 2024.
- [27] Benyuan Meng, Qianqian Xu, Zitai Wang, Xiaochun Cao, and Qingming Huang. Not all diffusion model activations have been evaluated as discriminative features. *Advances in Neural Information Processing Systems*, 37:55141–55177, 2025.
- [28] Hyelin Nam, Jaemin Kim, Dohun Lee, and Jong Chul Ye. Optical-flow guided prompt optimization for coherent video generation. *arXiv preprint arXiv:2411.15540*, 2024.
- [29] Suraj Patil, Pedro Cuenca, Nathan Lambert, and Patrick von Platen. Stable diffusion with diffusers. *Hugging Face Blog*, 2022. https://huggingface.co/blog/stable_diffusion.
- [30] 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.
- [31] Dustin Podell, Zion English, Kyle Lacey, Andreas Blattmann, Tim Dockhorn, Jonas Müller, Joe Penna, and Robin Rombach. Sdxl: Improving latent diffusion models for high-resolution image synthesis. *arXiv preprint arXiv:2307.01952*, 2023.
- [32] 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 *International conference on machine learning*, pages 8748–8763. PmLR, 2021.
- [33] Robin Rombach, Andreas Blattmann, Dominik Lorenz, Patrick Esser, and Björn Ommer. High-resolution image synthesis with latent diffusion models. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 10684–10695, 2022.
- [34] Maximilian Seitzer. pytorch-fid: FID Score for PyTorch. <https://github.com/mseitzer/pytorch-fid>, August 2020. Version 0.3.0.
- [35] Jiaming Song, Chenlin Meng, and Stefano Ermon. Denoising diffusion implicit models. In *International Conference on Learning Representations*, 2021. URL <https://openreview.net/forum?id=St1giarCHLP>.

- [36] Yang Song, Conor Durkan, Iain Murray, and Stefano Ermon. Maximum likelihood training of score-based diffusion models. *Advances in Neural Information Processing Systems*, 34, 2021.
- [37] Chameleon Team. Chameleon: Mixed-modal early-fusion foundation models. *arXiv preprint arXiv:2405.09818*, 2024.
- [38] Narek Tumanyan, Michal Geyer, Shai Bagon, and Tali Dekel. Plug-and-play diffusion features for text-driven image-to-image translation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 1921–1930, June 2023.
- [39] Soobin Um and Jong Chul Ye. Minorityprompt: Text to minority image generation via prompt optimization. *arXiv preprint arXiv:2410.07838*, 2024.
- [40] Rejin Varghese and M Sambath. Yolov8: A novel object detection algorithm with enhanced performance and robustness. In *2024 International Conference on Advances in Data Engineering and Intelligent Computing Systems (ADICS)*, pages 1–6. IEEE, 2024.
- [41] Pascal Vincent. A connection between score matching and denoising autoencoders. *Neural computation*, 23(7):1661–1674, 2011.
- [42] Bram Wallace, Akash Gokul, Stefano Ermon, and Nikhil Naik. End-to-end diffusion latent optimization improves classifier guidance. In *ICCV*, 2023.
- [43] Bram Wallace, Meihua Dang, Rafael Rafailov, Linqi Zhou, Aaron Lou, Senthil Purushwalkam, Stefano Ermon, Caiming Xiong, Shafiq Joty, and Nikhil Naik. Diffusion model alignment using direct preference optimization. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 8228–8238, 2024.
- [44] Xiaoshi Wu, Yiming Hao, Keqiang Sun, Yixiong Chen, Feng Zhu, Rui Zhao, and Hongsheng Li. Human preference score v2: A solid benchmark for evaluating human preferences of text-to-image synthesis. *arXiv preprint arXiv:2306.09341*, 2023.
- [45] Jiazheng Xu, Xiao Liu, Yuchen Wu, Yuxuan Tong, Qinkai Li, Ming Ding, Jie Tang, and Yuxiao Dong. Imagereward: learning and evaluating human preferences for text-to-image generation. In *Proceedings of the 37th International Conference on Neural Information Processing Systems*, pages 15903–15935, 2023.
- [46] 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 *International Conference on Learning Representations*, 2025.
- [47] Richard Zhang, Phillip Isola, Alexei A Efros, Eli Shechtman, and Oliver Wang. The unreasonable effectiveness of deep features as a perceptual metric. In *CVPR*, 2018.
- [48] SUN Zhengwentai. clip-score: CLIP Score for PyTorch. <https://github.com/taited/clip-score>, March 2023. Version 0.1.1.
- [49] Kaiyang Zhou, Jingkang Yang, Chen Change Loy, and Ziwei Liu. Conditional prompt learning for vision-language models. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 16816–16825, 2022.
- [50] Kaiyang Zhou, Jingkang Yang, Chen Change Loy, and Ziwei Liu. Learning to prompt for vision-language models. *International Journal of Computer Vision*, 130(9):2337–2348, 2022.

NeurIPS Paper Checklist

1. Claims

Question: Do the main claims made in the abstract and introduction accurately reflect the paper's contributions and scope?

Answer: [\[Yes\]](#)

Justification: We clearly claim the SoftREPA's contribution and scope in the abstract and introduction.

Guidelines:

- The answer NA means that the abstract and introduction do not include the claims made in the paper.
- The abstract and/or introduction should clearly state the claims made, including the contributions made in the paper and important assumptions and limitations. A No or NA answer to this question will not be perceived well by the reviewers.
- The claims made should match theoretical and experimental results, and reflect how much the results can be expected to generalize to other settings.
- It is fine to include aspirational goals as motivation as long as it is clear that these goals are not attained by the paper.

2. Limitations

Question: Does the paper discuss the limitations of the work performed by the authors?

Answer: [\[Yes\]](#)

Justification: We discussed the limitations of SoftREPA in the experiments, and appendix.

Guidelines:

- The answer NA means that the paper has no limitation while the answer No means that the paper has limitations, but those are not discussed in the paper.
- The authors are encouraged to create a separate "Limitations" section in their paper.
- The paper should point out any strong assumptions and how robust the results are to violations of these assumptions (e.g., independence assumptions, noiseless settings, model well-specification, asymptotic approximations only holding locally). The authors should reflect on how these assumptions might be violated in practice and what the implications would be.
- The authors should reflect on the scope of the claims made, e.g., if the approach was only tested on a few datasets or with a few runs. In general, empirical results often depend on implicit assumptions, which should be articulated.
- The authors should reflect on the factors that influence the performance of the approach. For example, a facial recognition algorithm may perform poorly when image resolution is low or images are taken in low lighting. Or a speech-to-text system might not be used reliably to provide closed captions for online lectures because it fails to handle technical jargon.
- The authors should discuss the computational efficiency of the proposed algorithms and how they scale with dataset size.
- If applicable, the authors should discuss possible limitations of their approach to address problems of privacy and fairness.
- While the authors might fear that complete honesty about limitations might be used by reviewers as grounds for rejection, a worse outcome might be that reviewers discover limitations that aren't acknowledged in the paper. The authors should use their best judgment and recognize that individual actions in favor of transparency play an important role in developing norms that preserve the integrity of the community. Reviewers will be specifically instructed to not penalize honesty concerning limitations.

3. Theory assumptions and proofs

Question: For each theoretical result, does the paper provide the full set of assumptions and a complete (and correct) proof?

Answer: [\[Yes\]](#)

Justification: We provide clear assumptions and complete proof of concept.

Guidelines:

- The answer NA means that the paper does not include theoretical results.
- All the theorems, formulas, and proofs in the paper should be numbered and cross-referenced.
- All assumptions should be clearly stated or referenced in the statement of any theorems.
- The proofs can either appear in the main paper or the supplemental material, but if they appear in the supplemental material, the authors are encouraged to provide a short proof sketch to provide intuition.
- Inversely, any informal proof provided in the core of the paper should be complemented by formal proofs provided in appendix or supplemental material.
- Theorems and Lemmas that the proof relies upon should be properly referenced.

4. Experimental result reproducibility

Question: Does the paper fully disclose all the information needed to reproduce the main experimental results of the paper to the extent that it affects the main claims and/or conclusions of the paper (regardless of whether the code and data are provided or not)?

Answer: [Yes]

Justification: We fully disclose the information about experiments for reproducibility.

Guidelines:

- The answer NA means that the paper does not include experiments.
- If the paper includes experiments, a No answer to this question will not be perceived well by the reviewers: Making the paper reproducible is important, regardless of whether the code and data are provided or not.
- If the contribution is a dataset and/or model, the authors should describe the steps taken to make their results reproducible or verifiable.
- Depending on the contribution, reproducibility can be accomplished in various ways. For example, if the contribution is a novel architecture, describing the architecture fully might suffice, or if the contribution is a specific model and empirical evaluation, it may be necessary to either make it possible for others to replicate the model with the same dataset, or provide access to the model. In general, releasing code and data is often one good way to accomplish this, but reproducibility can also be provided via detailed instructions for how to replicate the results, access to a hosted model (e.g., in the case of a large language model), releasing of a model checkpoint, or other means that are appropriate to the research performed.
- While NeurIPS does not require releasing code, the conference does require all submissions to provide some reasonable avenue for reproducibility, which may depend on the nature of the contribution. For example
 - (a) If the contribution is primarily a new algorithm, the paper should make it clear how to reproduce that algorithm.
 - (b) If the contribution is primarily a new model architecture, the paper should describe the architecture clearly and fully.
 - (c) If the contribution is a new model (e.g., a large language model), then there should either be a way to access this model for reproducing the results or a way to reproduce the model (e.g., with an open-source dataset or instructions for how to construct the dataset).
 - (d) We recognize that reproducibility may be tricky in some cases, in which case authors are welcome to describe the particular way they provide for reproducibility. In the case of closed-source models, it may be that access to the model is limited in some way (e.g., to registered users), but it should be possible for other researchers to have some path to reproducing or verifying the results.

5. Open access to data and code

Question: Does the paper provide open access to the data and code, with sufficient instructions to faithfully reproduce the main experimental results, as described in supplemental material?

Answer: [Yes]

Justification: The code will be provided.

Guidelines:

- The answer NA means that paper does not include experiments requiring code.
- Please see the NeurIPS code and data submission guidelines (<https://nips.cc/public/guides/CodeSubmissionPolicy>) for more details.
- While we encourage the release of code and data, we understand that this might not be possible, so “No” is an acceptable answer. Papers cannot be rejected simply for not including code, unless this is central to the contribution (e.g., for a new open-source benchmark).
- The instructions should contain the exact command and environment needed to run to reproduce the results. See the NeurIPS code and data submission guidelines (<https://nips.cc/public/guides/CodeSubmissionPolicy>) for more details.
- The authors should provide instructions on data access and preparation, including how to access the raw data, preprocessed data, intermediate data, and generated data, etc.
- The authors should provide scripts to reproduce all experimental results for the new proposed method and baselines. If only a subset of experiments are reproducible, they should state which ones are omitted from the script and why.
- At submission time, to preserve anonymity, the authors should release anonymized versions (if applicable).
- Providing as much information as possible in supplemental material (appended to the paper) is recommended, but including URLs to data and code is permitted.

6. Experimental setting/details

Question: Does the paper specify all the training and test details (e.g., data splits, hyper-parameters, how they were chosen, type of optimizer, etc.) necessary to understand the results?

Answer: [Yes]

Justification: We specified all the training and test details.

Guidelines:

- The answer NA means that the paper does not include experiments.
- The experimental setting should be presented in the core of the paper to a level of detail that is necessary to appreciate the results and make sense of them.
- The full details can be provided either with the code, in appendix, or as supplemental material.

7. Experiment statistical significance

Question: Does the paper report error bars suitably and correctly defined or other appropriate information about the statistical significance of the experiments?

Answer: [No]

Justification: In reporting results, we have been following existing literature in this domain, evaluating various methods and metrics over 5K images in the generation experiment, and reporting aggregate scores.

Guidelines:

- The answer NA means that the paper does not include experiments.
- The authors should answer "Yes" if the results are accompanied by error bars, confidence intervals, or statistical significance tests, at least for the experiments that support the main claims of the paper.
- The factors of variability that the error bars are capturing should be clearly stated (for example, train/test split, initialization, random drawing of some parameter, or overall run with given experimental conditions).
- The method for calculating the error bars should be explained (closed form formula, call to a library function, bootstrap, etc.)
- The assumptions made should be given (e.g., Normally distributed errors).

- It should be clear whether the error bar is the standard deviation or the standard error of the mean.
- It is OK to report 1-sigma error bars, but one should state it. The authors should preferably report a 2-sigma error bar than state that they have a 96% CI, if the hypothesis of Normality of errors is not verified.
- For asymmetric distributions, the authors should be careful not to show in tables or figures symmetric error bars that would yield results that are out of range (e.g. negative error rates).
- If error bars are reported in tables or plots, The authors should explain in the text how they were calculated and reference the corresponding figures or tables in the text.

8. Experiments compute resources

Question: For each experiment, does the paper provide sufficient information on the computer resources (type of compute workers, memory, time of execution) needed to reproduce the experiments?

Answer: [Yes]

Justification: We provide information on the computer resources.

Guidelines:

- The answer NA means that the paper does not include experiments.
- The paper should indicate the type of compute workers CPU or GPU, internal cluster, or cloud provider, including relevant memory and storage.
- The paper should provide the amount of compute required for each of the individual experimental runs as well as estimate the total compute.
- The paper should disclose whether the full research project required more compute than the experiments reported in the paper (e.g., preliminary or failed experiments that didn't make it into the paper).

9. Code of ethics

Question: Does the research conducted in the paper conform, in every respect, with the NeurIPS Code of Ethics <https://neurips.cc/public/EthicsGuidelines>?

Answer: [Yes]

Justification: We reviewed the NeurIPS Code of Ethics.

Guidelines:

- The answer NA means that the authors have not reviewed the NeurIPS Code of Ethics.
- If the authors answer No, they should explain the special circumstances that require a deviation from the Code of Ethics.
- The authors should make sure to preserve anonymity (e.g., if there is a special consideration due to laws or regulations in their jurisdiction).

10. Broader impacts

Question: Does the paper discuss both potential positive societal impacts and negative societal impacts of the work performed?

Answer: [Yes]

Justification: We mentioned the societal impact of our work in the conclusion.

Guidelines:

- The answer NA means that there is no societal impact of the work performed.
- If the authors answer NA or No, they should explain why their work has no societal impact or why the paper does not address societal impact.
- Examples of negative societal impacts include potential malicious or unintended uses (e.g., disinformation, generating fake profiles, surveillance), fairness considerations (e.g., deployment of technologies that could make decisions that unfairly impact specific groups), privacy considerations, and security considerations.

- The conference expects that many papers will be foundational research and not tied to particular applications, let alone deployments. However, if there is a direct path to any negative applications, the authors should point it out. For example, it is legitimate to point out that an improvement in the quality of generative models could be used to generate deepfakes for disinformation. On the other hand, it is not needed to point out that a generic algorithm for optimizing neural networks could enable people to train models that generate Deepfakes faster.
- The authors should consider possible harms that could arise when the technology is being used as intended and functioning correctly, harms that could arise when the technology is being used as intended but gives incorrect results, and harms following from (intentional or unintentional) misuse of the technology.
- If there are negative societal impacts, the authors could also discuss possible mitigation strategies (e.g., gated release of models, providing defenses in addition to attacks, mechanisms for monitoring misuse, mechanisms to monitor how a system learns from feedback over time, improving the efficiency and accessibility of ML).

11. Safeguards

Question: Does the paper describe safeguards that have been put in place for responsible release of data or models that have a high risk for misuse (e.g., pretrained language models, image generators, or scraped datasets)?

Answer: [NA]

Justification: This paper poses no such risks.

Guidelines:

- The answer NA means that the paper poses no such risks.
- Released models that have a high risk for misuse or dual-use should be released with necessary safeguards to allow for controlled use of the model, for example by requiring that users adhere to usage guidelines or restrictions to access the model or implementing safety filters.
- Datasets that have been scraped from the Internet could pose safety risks. The authors should describe how they avoided releasing unsafe images.
- We recognize that providing effective safeguards is challenging, and many papers do not require this, but we encourage authors to take this into account and make a best faith effort.

12. Licenses for existing assets

Question: Are the creators or original owners of assets (e.g., code, data, models), used in the paper, properly credited and are the license and terms of use explicitly mentioned and properly respected?

Answer: [Yes]

Justification: This paper cites the original paper that produced dataset and models.

Guidelines:

- The answer NA means that the paper does not use existing assets.
- The authors should cite the original paper that produced the code package or dataset.
- The authors should state which version of the asset is used and, if possible, include a URL.
- The name of the license (e.g., CC-BY 4.0) should be included for each asset.
- For scraped data from a particular source (e.g., website), the copyright and terms of service of that source should be provided.
- If assets are released, the license, copyright information, and terms of use in the package should be provided. For popular datasets, paperswithcode.com/datasets has curated licenses for some datasets. Their licensing guide can help determine the license of a dataset.
- For existing datasets that are re-packaged, both the original license and the license of the derived asset (if it has changed) should be provided.

- If this information is not available online, the authors are encouraged to reach out to the asset's creators.

13. **New assets**

Question: Are new assets introduced in the paper well documented and is the documentation provided alongside the assets?

Answer: [NA]

Justification: This paper does not release new assets.

Guidelines:

- The answer NA means that the paper does not release new assets.
- Researchers should communicate the details of the dataset/code/model as part of their submissions via structured templates. This includes details about training, license, limitations, etc.
- The paper should discuss whether and how consent was obtained from people whose asset is used.
- At submission time, remember to anonymize your assets (if applicable). You can either create an anonymized URL or include an anonymized zip file.

14. **Crowdsourcing and research with human subjects**

Question: For crowdsourcing experiments and research with human subjects, does the paper include the full text of instructions given to participants and screenshots, if applicable, as well as details about compensation (if any)?

Answer: [NA]

Justification: This paper does not involve crowdsourcing nor research with human subject.

Guidelines:

- The answer NA means that the paper does not involve crowdsourcing nor research with human subjects.
- Including this information in the supplemental material is fine, but if the main contribution of the paper involves human subjects, then as much detail as possible should be included in the main paper.
- According to the NeurIPS Code of Ethics, workers involved in data collection, curation, or other labor should be paid at least the minimum wage in the country of the data collector.

15. **Institutional review board (IRB) approvals or equivalent for research with human subjects**

Question: Does the paper describe potential risks incurred by study participants, whether such risks were disclosed to the subjects, and whether Institutional Review Board (IRB) approvals (or an equivalent approval/review based on the requirements of your country or institution) were obtained?

Answer: [NA]

Justification: This paper does not involve crowdsourcing nor research with human subjects.

Guidelines:

- The answer NA means that the paper does not involve crowdsourcing nor research with human subjects.
- Depending on the country in which research is conducted, IRB approval (or equivalent) may be required for any human subjects research. If you obtained IRB approval, you should clearly state this in the paper.
- We recognize that the procedures for this may vary significantly between institutions and locations, and we expect authors to adhere to the NeurIPS Code of Ethics and the guidelines for their institution.
- For initial submissions, do not include any information that would break anonymity (if applicable), such as the institution conducting the review.

16. **Declaration of LLM usage**

Question: Does the paper describe the usage of LLMs if it is an important, original, or non-standard component of the core methods in this research? Note that if the LLM is used only for writing, editing, or formatting purposes and does not impact the core methodology, scientific rigorousness, or originality of the research, declaration is not required.

Answer: [NA]

Justification: The core method development in this research does not involve LLMs.

Guidelines:

- The answer NA means that the core method development in this research does not involve LLMs as any important, original, or non-standard components.
- Please refer to our LLM policy (<https://neurips.cc/Conferences/2025/LLM>) for what should or should not be described.

SoftREPA for better Text-to-Image Alignment

Supplementary Materials

A Information Theoretical Analysis of Diffusion Models

Several studies have explored diffusion models from an information-theoretic viewpoint. According to [18], the mutual information between random variables X and Y can be expressed as the expectation of pointwise mutual information over their joint distribution:

$$I(X, Y) = \mathbb{E}_{p(\mathbf{x}, \mathbf{y})}[i(\mathbf{x}, \mathbf{y})], \quad (19)$$

where the pointwise mutual information (PMI) is defined as the difference between the conditional and marginal log-likelihoods of a given sample $\mathbf{x}$:

$$i(\mathbf{x}, \mathbf{y}) = \log p(\mathbf{x}|\mathbf{y}) - \log p(\mathbf{x}). \quad (20)$$

This quantity measures how much the presence of condition $\mathbf{y}$ affects the probability distribution near $\mathbf{x}$.

Song et al. [36] and Kong et al. [18] showed that, under the assumption of an optimal diffusion model trained on given data, the log-likelihood of an image $\mathbf{x}$ can be formulated as:

$$-\log p(\mathbf{x}) = \frac{1}{2} \int_0^T \lambda(t) \mathbb{E}_{\epsilon} [\|\epsilon_{\theta}(\mathbf{x}_t, t) - \epsilon\|^2] dt + C \quad (21)$$

This result also extends to conditional diffusion models, where the likelihood of $\mathbf{x}$ given condition $\mathbf{y}$ (e.g., text embeddings) is given by:

$$-\log p(\mathbf{x}|\mathbf{y}) = \frac{1}{2} \int_0^T \lambda(t) \mathbb{E}_{\epsilon} [\|\epsilon_{\theta}(\mathbf{x}_t, t, \mathbf{y}) - \epsilon\|^2] dt + C \quad (22)$$

B Implementation Details

Architecture Details For Stable Diffusion 3, the main experiments use soft tokens of length 4, applied across 5 layers. In Stable Diffusion 1.5, soft tokens of length 4 are applied only to the Down block layers of the UNet’s conditional text embeddings. For Stable Diffusion XL, soft tokens of length 8 are used, applied to both the Down and Middle block layers of the UNet’s conditional text embeddings, without incorporating timestep dependency.

Training Details The hyperparameters used during training are listed in table 5. For Stable Diffusion 3, the soft tokens are optimized solely using the contrastive score matching loss (L_{SoftREPA}). In contrast, for Stable Diffusion 1.5 and Stable Diffusion XL, optimization combines the contrastive score matching loss (L_{SoftREPA}) with a small weighting of the denoising score matching loss (L_{DSM}). table 6 shows that the model consistently achieves improved performance regardless of whether the denoising score matching loss is applied. We observed that initializing the soft tokens with a random distribution led to performance degradation in Stable Diffusion 1.5. To address this, we initialized the soft tokens using unconditional text embeddings. The code is publicly available at <https://github.com/softrepa/SoftREPA>.

Models	lr	wd	batch size (positive, negative)	iterations	token init	optimizer	lr scheduler
SD1.5	1e-3	1e-4	32(4, 28)	26,000	$\emptyset$	AdamW	CosineAnnealingWarmRestarts
SDXL	1e-3	1e-4	16(1, 15)	30,000	$N(0, 0.02)$	AdamW	CosineAnnealingWarmRestarts
SD3	1e-3	1e-4	16(4, 12)	30,000	$N(0, 0.02)$	AdamW	CosineAnnealingWarmRestarts

Table 5: The implementation details for training.

Model	loss	Human Preference		Text Alignment		Image Quality	
		ImageReward $\uparrow$	PickScore $\uparrow$	CLIP $\uparrow$	HPS $\uparrow$	FID $\downarrow$	LPIPS $\downarrow$
SD1.5	-	17.72	21.47	26.45	25.08	24.59	43.80
	$L_{\text{SoftREPA}} + L_{\text{DSM}}$	32.89	21.50	27.33	25.18	23.43	43.38
	L_{SoftREPA}	32.63	21.61	27.10	25.94	29.25	44.09
SDXL	-	75.06	22.38	26.76	27.35	24.69	42.05
	$L_{\text{SoftREPA}} + L_{\text{DSM}}$	85.29	22.62	26.80	28.30	26.04	42.39
	L_{SoftREPA}	81.09	22.59	26.87	28.32	26.42	42.69

Table 6: Quantitative comparison of T2I generation between contrastive score matching loss (L_{SoftREPA}) with and without denoising score matching loss (L_{DSM}). Generation quality is evaluated on the COCO-val 1K [23]. ImageReward, CLIP, HPS, and LPIPS are scaled by $\times 10^2$.

Algorithm 1 Image Generation with Soft Tokens in MM-DiT

Require: Gaussian noise $\mathbf{z} \sim \mathcal{N}(\mathbf{0}, \mathbf{I})$
Require: Text $\mathbf{Y} \sim p_{\text{data}}$
Require: Soft token $\mathbf{s} \sim \text{Embedding}(k, t)$
Require: Number of layers N , Threshold layer L
Require: Time steps $\{t_T, t_{T-1}, \dots, t_0\}$

- 1: Initialize $\mathbf{H}_{\text{img}}^{(0,T)} \leftarrow \mathbf{z}$
- 2: Initialize $\mathbf{H}_{\text{text}}^{(0,T)} \leftarrow \text{TextEncoder}(\mathbf{Y})$
- 3: $n = |\mathbf{H}_{\text{text}}^{(0,T)}|$
- 4: **for** t in $\{t_T, t_{T-1}, \dots, t_0\}$ **do**
- 5: **for** $l = 1$ to N **do**
- 6: **if** $k \leq L$ **then**
- 7: $\mathbf{s}^{(k,t)} \leftarrow \text{Embedding}(k, t)$
- 8: $\hat{\mathbf{H}}_{\text{text}}^{(k-1,t)} \leftarrow [\mathbf{s}^{(k,t)}; \mathbf{H}_{\text{text}}^{(k-1,t)}]$
- 9: **else**
- 10: $\hat{\mathbf{H}}_{\text{text}}^{(k-1,t)} \leftarrow \mathbf{H}_{\text{text}}^{(k-1,t)}$
- 11: **end if**
- 12: $\mathbf{H}_{\text{img}}^{(k,t)}, \hat{\mathbf{H}}_{\text{text}}^{(k,t)} \leftarrow \text{Layer}_l(\mathbf{H}_{\text{img}}^{(k-1,t)}, \hat{\mathbf{H}}_{\text{text}}^{(k-1,t)})$
- 13: $\mathbf{H}_{\text{text}}^{(k,t)} \leftarrow \hat{\mathbf{H}}_{\text{text}}^{(k,t)}[-n :, :]$ $\triangleright$ Drop soft tokens
- 14: **end for**
- 15: **end for**
- 16: **return** $\hat{\mathbf{X}} = \text{Decoder}(\mathbf{H}_{\text{img}}^{(N,t_0)})$

Inference Details A detailed description of the image generation algorithm on MM-DiT is provided in algorithm 1. At each layer, a distinct set of soft tokens is prepended to the text features and used exclusively within that layer. These soft tokens are not carried over to subsequent layers; instead, new soft tokens are introduced or omitted as appropriate. Their primary role is to guide the text tokens toward better text-image alignment, particularly during the early stages of the model.

Memory Efficiency For the memory efficiency metrics reported in table 1, values were measured by averaging results over 50 runs with an A100 GPU. Image resolution was set to 512×512 for SD1.5 and 1024×1024 for SDXL and SD3.

C Prompt for Editing dataset

To generate source/target text prompt from images, we leverage LLaVA [25]¹ and Llama [10]² sequentially. Specifically, we extract source text description from 800 training images of DIV2K dataset by giving each image and following text prompt to LLaVA.

"Describe the object and background in the image"

¹We use checkpoint from 4bit/llava-v1.5-13b-3GB

²We use Llama-3.1-8B-Instruct model.

Then, we generate the target text prompt that has only a single different concept compared with the source text prompt using Llama with the following instruction.

"You are an AI assistant for generating paired text prompts for real image editing tasks. Your goal is to modify a given text description by replacing an object with other while strictly following these rules:

- 1) Modify only one object (i.e., a single meaningful concept such as an object). It could be small object.
- 2) The replacement must be significantly different from the original concept but contextually appropriate. Avoid unrealistic substitutions (e.g., changing "rabbit on grass" to "rocket on grass").
- 3) Ensure diversity in word choices across different modifications.
- 4) Preserve all other words exactly as they are. Do not change sentence structure, introduce new elements, or modify additional details.
- 5) Do not provide any additional words—output only the modified text description.
- 6) Do not change or add colors. Specifically, when modifying a building, change only the appearance, not the type of building (e.g., do not change "building" to "church" or "lighthouse").
- 7) Modify only one feature at a time. If changing an object (e.g., "starfish" to "sea turtle"), do not alter its color, shape, or other attributes.

Example:

Input: The image features a close-up of a brown dog with a blue nose. The dog is standing in a grassy field, and the background is blurred, creating a focus on the dog's face. The dog's ears are perked up, and its eyes are open, giving it a curious and attentive expression. The dog's fur is brown, and the grass in the background is green, creating a natural and vibrant scene.

Output: The image features a close-up of a brown fox with a blue nose. The fox is standing in a grassy field, and the background is blurred, creating a focus on the fox's face. The fox's ears are perked up, and its eyes are open, giving it a curious and attentive expression. The fox's fur is brown, and the grass in the background is green, creating a natural and vibrant scene."

D Discussion on counting metric

Adopting Soft Tokens Only A Few Layers As shown in section 4, our model demonstrates notable improvements in generating single and multiple objects, as well as in color and position-aware synthesis. However, it exhibits a relative decline in accurately generating the specified number of objects described in the text prompts. We hypothesize that this drop in the counting metric stems from the soft tokens excessively emphasizing textual cues, which can lead to overgeneration of object instances. To address this, we limited the use of soft tokens to the early layers of the model. As presented in table 7, applying soft tokens only to layers 1 ~ 2 preserves strong overall performance while mitigating the degradation in counting accuracy observed when soft tokens are applied across layers 1 ~ 5.

Incorporating Counting Loss during training To further enhance counting fidelity, we trained SoftREPA with a lightweight object-counting loss. Specifically, we employ a mean squared error (MSE) loss between the predicted object count, which is derived from denoised images, and ground-truth labels using the lightweight object detection module, YOLOv8 [40]. This variant, shown in the last row of table 7, shows improved counting accuracy without sacrificing overall generation quality, suggesting that it effectively reduces the tendency to replicate objects due to textual overemphasis.

E Additional Results on T2I Generation

Additional qualitative results for SD1.5 and SDXL on the COCO-val5K dataset are presented in fig. 7 and fig. 8, respectively. Furthermore, qualitative examples on the GenEval benchmark are provided in fig. 9.

Model	Layers	GenEval						
		Mean↑	Single↑	Two↑	Counting↑	Colors↑	Position↑	Color Attribution↑
SD3		0.68	<u>0.99</u>	0.86	<u>0.56</u>	0.85	0.27	0.55
SD3 (Ours)	1	0.70	<u>0.99</u>	<u>0.91</u>	<u>0.51</u>	0.89	<u>0.31</u>	<u>0.58</u>
SD3 (Ours)	1~2	<u>0.69</u>	1.00	0.89	0.50	0.88	0.34	0.56
SD3 (Ours)	1~5	0.70	1.00	0.95	0.29	0.92	0.34	0.68
SD3 (Ours) + count	1~5	<u>0.69</u>	<u>0.99</u>	0.88	0.59	0.86	0.25	0.55

Table 7: Additional quantitative comparison on the GenEval [9] benchmark. **Bold** indicates the best performance, and underline denotes the second best. "Layers" refers to the transformer layers where soft tokens are applied.




Figure 7: The qualitative results of text-to-image generation comparing SD1.5 and SD1.5 with proposed method. The given text is from COCO dataset.




Figure 8: The qualitative results of text-to-image generation comparing SDXL and SDXL with proposed method. The given text is from COCO dataset.




Figure 9: The qualitative results of text-to-image generation comparing SD3 and SD3 with proposed method. The given text is from GenEval dataset.

F Additional Results on Image Editing

Qualitative comparisons of SoftREPA across various editing methods, including PnP [38], MasaCtrl [4] with DDIM [35], Direct Inversion [14], and FlowEdit [19], on the PIEBench dataset are presented in fig. 11, corresponding to the quantitative results in table 2.

Quantitative results on the DIV2K and Cat2Dog datasets under varying target CFG scales are reported in table 8 and table 9, respectively. The corresponding qualitative results for different CFG scales are also shown in table 8. Additionally, qualitative examples across different numbers of editing steps are provided in fig. 12 and fig. 10.

Model	NMAX	target CFG	Human Preference		Text Alignment		Image Quality	
			ImageReward↑	PickScore↑	CLIP↑	HPS↑	FID↓	LPIPS↓
FlowEdit	33	13.5	0.380	21.749	0.261	0.255	52.038	0.154
FlowEdit	33	16.5	0.466	21.824	0.263	0.259	55.371	0.180
FlowEdit	33	19.5	0.510	21.869	0.265	0.261	58.069	0.200
FlowEdit +Ours	30	9	0.466	21.884	0.263	0.260	52.633	0.149
FlowEdit +Ours	30	11	0.564	21.985	0.268	0.265	56.900	0.178
FlowEdit +Ours	30	13	0.627	22.050	0.272	0.270	59.710	0.201

Table 8: Quantitative results of image editing regarding image quality and target text alignment of generated images with various target CFG scales on **DIV2K** dataset.

Model	NMAX	target CFG	Human Preference		Text Alignment		Image Quality	
			ImageReward↑	PickScore↑	CLIP↑	HPS↑	FID↓	LPIPS↓
FlowEdit	33	13.5	0.937	20.726	0.225	0.266	197.1	0.199
FlowEdit	33	16.5	0.908	20.641	0.229	0.272	201.05	0.217
FlowEdit	33	19.5	0.882	20.572	0.232	0.276	202.92	0.234
FlowEdit +Ours	30	7	1.111	21.143	0.226	0.269	198.71	0.173
FlowEdit +Ours	30	9	1.144	21.221	0.234	0.283	208.18	0.204
FlowEdit +Ours	30	11	1.144	21.224	0.239	0.290	214.72	0.230
FlowEdit +Ours	30	13	1.135	21.200	0.241	0.293	219.35	0.252

Table 9: Quantitative results of image editing regarding image quality and target text alignment of generated images with various target CFG scales on **Cat2Dog** dataset.




Figure 10: Additional qualitative results of text-guided image editing on Cat2Dog dataset.




Figure 11: The qualitative results of various editing methods including PnP, MasaCtrl, FlowEdit with DDIM and Direct inversion on PIEBench.

G Additional Results on Ablation Study

Quantitative results for varying the number of layers and soft tokens are shown in table 10, with corresponding qualitative examples in fig. 13 and fig. 14. As seen in fig. 13, text prompts are increasingly emphasized as more layers adopt soft tokens. The same input image is used in fig. 14 to illustrate the effect of different token counts.




Figure 12: The additional qualitative results on editing using FlowEdit. The edited1-3 using soft tokens are vary on the number of editing steps, 27, 29, 30 out of 50 total timesteps, respectively.

H Additional Results on comparing with LoRA

In table 11, we conducted various conditions on LoRA finetuning, including top-5-layer and full-layer LoRA, as well as the corresponding number of trainable parameters and the chosen LoRA rank. Specifically, LoRA with top-5 layers uses 0.4M parameters, LoRA applied to all layers uses 2.3M parameters, and our soft token approach uses 0.9M parameters for 5 layers. We adopt a LoRA rank of 4, which is commonly used for diffusion model finetuning.

To enable a direct comparison with a similar number of parameters, we additionally report results for LoRA applied to the top-10 layers (0.9M parameters) and top-5 layers with rank 8 (0.9M parameters). Among these settings, LoRA on the top-5 layers with rank 4 yields the best results within LoRA baselines. Importantly, the combination of Soft Tokens and contrastive Loss consistently achieves the strongest performance across key metrics, supporting our claim that both techniques independently and meaningfully improve text-image alignment.



Figure 13: Qualitative results of the ablation study on the **number of layers** adopting soft tokens.

# of tokens	# of layers	Human Preference		Text Alignment		Image Quality	
		ImageReward↑	PickScore↑	CLIP↑	HPS↑	FID↓	LPIPS↓
8	2	0.993	22.556	0.263	0.286	72.253	0.428
8	4	1.009	22.464	0.266	0.285	72.308	0.424
8	6	0.984	22.403	0.267	0.286	73.840	0.427
8	7	0.944	22.197	0.269	0.280	74.546	0.425
1	5	1.054	22.492	0.272	0.283	58.430	0.426
4	5	1.063	22.493	0.269	0.287	60.766	0.429
8	5	1.056	22.516	0.268	0.288	62.644	0.431
16	5	0.801	22.095	0.265	0.278	60.743	0.431
32	5	0.675	21.876	0.257	0.275	61.623	0.426

Table 10: Quantitative results of image generation ablation study on the number of soft tokens and the number of layers. The evaluation is conducted on COCO val 1K dataset using SD3 backbone.



Figure 14: Qualitative results of the ablation study on the **number of tokens** adopting soft tokens.

Model	Rank	Layers	# of Params	Human Preference		Text Alignment		Image Quality	
				ImageReward $\uparrow$	PickScore $\uparrow$	CLIP $\uparrow$	HPS $\uparrow$	FID $\downarrow$	LPIPS $\downarrow$
SD3				90.89	22.50	26.26	27.97	56.87	42.47
SD3 + LoRA + Contrastive loss	4	1-5	0.4M	97.72	22.57	26.42	28.44	70.33	42.71
SD3 + LoRA + Contrastive loss	8	1-5	0.9M	95.92	22.55	26.28	28.28	70.30	42.55
SD3 + LoRA + Contrastive loss	4	1-10	0.9M	92.72	22.54	26.23	28.10	70.53	42.59
SD3 + LoRA + Contrastive loss	4	all layers	2.3M	80.89	22.20	26.70	25.97	71.57	41.54
SD3 + Soft tokens + Contrastive loss (Ours)	-	1-5	0.9M	106.32	22.49	26.92	28.70	60.76	42.99

Table 11: Quantitative comparison of T2I generation between soft tokens and various configurations of LoRA on SD3. Generation quality is evaluated on the COCO-val 1K [23]. ImageReward, CLIP, HPS, and LPIPS are scaled by $\times 10^2$.