

CustomContrast: A Multilevel Contrastive Perspective For Subject-Driven Text-to-Image Customization

Nan Chen, Mengqi Huang, Zhuowei Chen, Yang Zheng, Lei Zhang, Zhendong Mao*

University of Science and Technology of China

{chen_nan, huangmq, chenzw01, zy849900389}@mail.ustc.edu.cn, {leizh23, zdmao}@ustc.edu.cn

Abstract

Subject-driven text-to-image (T2I) customization has drawn significant interest in academia and industry. This task enables pre-trained models to generate novel images based on unique subjects. Existing studies adopt a self-reconstructive perspective, focusing on capturing all details of a single image, which will misconstrue the specific image’s irrelevant attributes (*e.g.*, view, pose, and background) as the subject intrinsic attributes. This misconstruction leads to both overfitting or underfitting of irrelevant and intrinsic attributes of the subject, *i.e.*, these attributes are over-represented or under-represented simultaneously, causing a trade-off between similarity and controllability. In this study, we argue an ideal subject representation can be achieved by a cross-differential perspective, *i.e.*, decoupling subject intrinsic attributes from irrelevant attributes via contrastive learning, which allows the model to focus more on intrinsic attributes through intra-consistency (features of the same subject are spatially closer) and inter-distinctiveness (features of different subjects have distinguished differences). Specifically, we propose **Custom-Contrast**, a novel framework, which includes a Multilevel Contrastive Learning (MCL) paradigm and a Multimodal Feature Injection (MFI) Encoder. The MCL paradigm is used to extract intrinsic features of subjects from high-level semantics to low-level appearance through crossmodal semantic contrastive learning and multiscale appearance contrastive learning. To facilitate contrastive learning, we introduce the MFI encoder to capture cross-modal representations. Extensive experiments show the effectiveness of *CustomContrast* in subject similarity and text controllability.

code — <https://cn-makers.github.io/CustomContrast/>

1 Introduction

Subject-driven text-to-image (T2I) customization aims to empower pre-trained text-to-image models to generate images of unique subjects (*e.g.*, objects, animals) specified by users. Given an image of a specific subject, users can generate novel depictions of the subject in various scenes, poses, and motions guided by target prompts. This task has drawn rapidly growing research interests in both academia and industry because of its broad applications. The key challenge

*Corresponding author
Copyright © 2025, Association for the Advancement of Artificial Intelligence (www.aaai.org). All rights reserved.

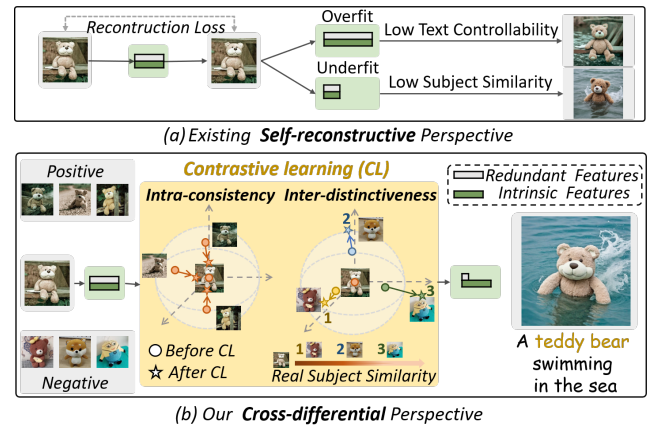


Figure 1: Comparison with existing perspective. (a) Existing studies learn each subject feature with entangled redundant features (*e.g.*, view, pose), suffering a trade-off between similarity and controllability (redundant and intrinsic features simultaneously overfit or underfit since they are coupled together). (b) In contrast, we rethink it from a cross-differential perspective. By using contrastive learning to ensure intra-consistency (features of the same subject are spatially closer) and inter-distinctiveness (features of different subjects have distinguished differences), our model disentangles the subject intrinsic features from irrelevant features for dual optimization of controllability and similarity.

of customized T2I generation lies in how to precisely generate the target subject while maximally preserving the text controllability of the pre-trained diffusion models.

Existing customized T2I methods mainly involve two streams: the finetune-based stream and the finetune-free stream. The finetune-based stream (Gal et al. 2022; Ruiz et al. 2023; Kumari et al. 2023) typically fine-tunes each new concept. In this paper, we focus on the finetune-free stream, which achieves real-time generation by training adapters based on the pre-trained encoders. Early work (Wei et al. 2023; Ye et al. 2023) aims to enhance subject similarity by injecting image embeddings into U-Net using additional visual attention. Recent studies further improve text controllability while preserving subject similarity by selectively in-

jecting image features. Some work (Li, Li, and Hoi 2024; Zhang et al. 2024) reduces background interferences by using different reference and reconstruction images, while other work (Song et al. 2024) controls the generated regions through additional inference operations. Additionally, some work (Sun et al. 2024) introduces multimodal large language models to further enhance text controllability. To summarize, existing methods adopt a self-reconstructive perspective, that is, taking target subject images as inputs, aiming to fully mimic all the subject details via the diffusion reconstruction loss (Ho, Jain, and Abbeel 2020).

However, the self-reconstructive perspective has inherent defects, *i.e.*, misconstruing the attributes in a specific image as the intrinsic attributes of the subject. The image’s attributes are entangled, including irrelevant attributes (*e.g.*, orientation, pose, and background), and intrinsic attributes of the subject that remain unchanged despite external interferences. Thus, reconstruction from a single image of the subject will inevitably result in the capture of both irrelevant and subject-intrinsic attributes. As a result, Insufficient representations result in inaccurate extraction of both intrinsic and irrelevant attributes, reducing similarity (*e.g.*, underfitting in Fig. 1(a) with changes in the toy’s color). Conversely, over-representations cause overfitting of both intrinsic and irrelevant attributes, reducing controllability (*e.g.*, overfitting in Fig. 1(a) with the toy’s pose and scene hard to alter). Expanding the dataset or changing the image background still follows this perspective, therefore these cannot significantly improve the misconstruction of redundant attributes.

Therefore, we argue an ideal subject representation can be achieved by a **cross-differential** perspective, *i.e.*, comparing differences between target samples via contrastive learning, which aims to capture each subject’s accurate representation. As shown in Fig. 1(b), this perspective achieves intra-consistency and inter-distinctiveness. Firstly, Intra-consistency is achieved by pulling images of the same subject under different contexts closer, decoupling irrelevant attributes. Secondly, Inter-distinctiveness is ensured by comparing the specific subject with others, thereby further learning the fine-grained intrinsic features. These allow the model to focus more on the intrinsic attributes than redundant attributes. By reducing redundant attributes’ interference, the cross-differential perspective can achieve more accurate extraction of subject intrinsic features while enabling flexible interaction with text (*e.g.*, the swimming toy in Fig. 1(b)).

In this paper, we propose a novel customization framework, *CustomContrast*, from the cross-differential perspective, which uses contrastive learning to extract intrinsic attributes for dual optimization of controllability and similarity. *CustomContrast* includes two key components: the Multimodal Feature Injection (MFI) Encoder, which aims to generate consistent multimodal representations to support the implementation of contrastive learning, and a Multilevel Contrastive Learning (MCL) paradigm that focuses on contrasting subject representations from low-level appearance to high-level semantics. Specifically, the MFI-Encoder includes Visual Qformer, Textual Qformer, and TextualVisual (TV) Fusion module. The first two modules are used to extract visual and textual embeddings. The TV Fusion module

uses textual embeddings as queries to capture text-related visual embeddings. The novel MCL paradigm includes Cross-modal Semantic Contrastive Learning (CSCL) and Multi-scale Appearance Contrastive Learning (MACL). CSCL enhances the semantic consistency between learned textual and visual embeddings, while MACL preserves the relative distances of samples between the learned embedding space and the real subject space.

Our main contributions are summarized as follows:

- **Concepts.** We propose *CustomContrast*, a novel paradigm from a cross-differential perspective. By ensuring intra-consistency and inter-distinctiveness, it gradually extracts the intrinsic subject representations.
- **Technology.** We propose a novel MCL paradigm that extracts intrinsic representations of subjects from high-level semantics to low-level appearance through CSCL and MACL. To support the MCL paradigm, we introduce the MFI-encoder to capture cross-modal representations.
- **Experiments.** Our model, trained on SD-V1.5 and SDXL, outperforms corresponding advanced methods. Experiments show our model improves text controllability by 3.8% and 5.4% respectively, and subject similarity (E-DI) by 5.9% and 2.4%, while easily extending to multi-subject and human domain generation.

2 Related Work

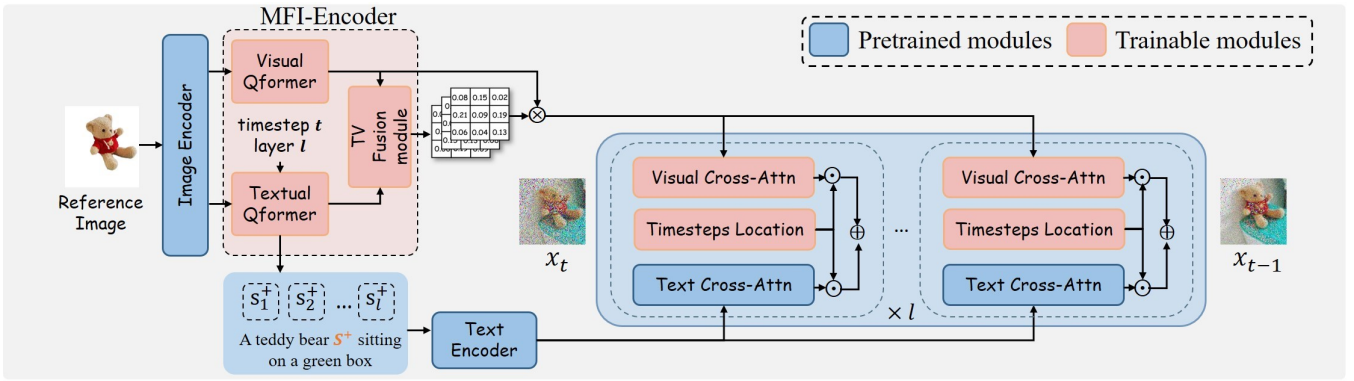
2.1 Subject-Driven Text-to-image Customization

Existing customized text-to-image work includes finetune-based streams and finetune-free streams. Early work (Gal et al. 2022; Ruiz et al. 2023; Kumari et al. 2023; Han et al. 2023; Voynov et al. 2023; Alaluf et al. 2023; Zhang et al. 2023; Nam et al. 2024; Huang et al. 2022) focuses on the finetune-based stream, which requires learning pseudo-words as the subject representations and fine-tuning each new subject. Custom Diffusion (Kumari et al. 2023) and SVDiff (Han et al. 2023) reduce computational overhead by minimizing fine-tuning parameters. Recent work (Voynov et al. 2023; Zhang et al. 2023) in the \mathcal{P}^+ space uses different pseudo-words for different cross-attention and timesteps.

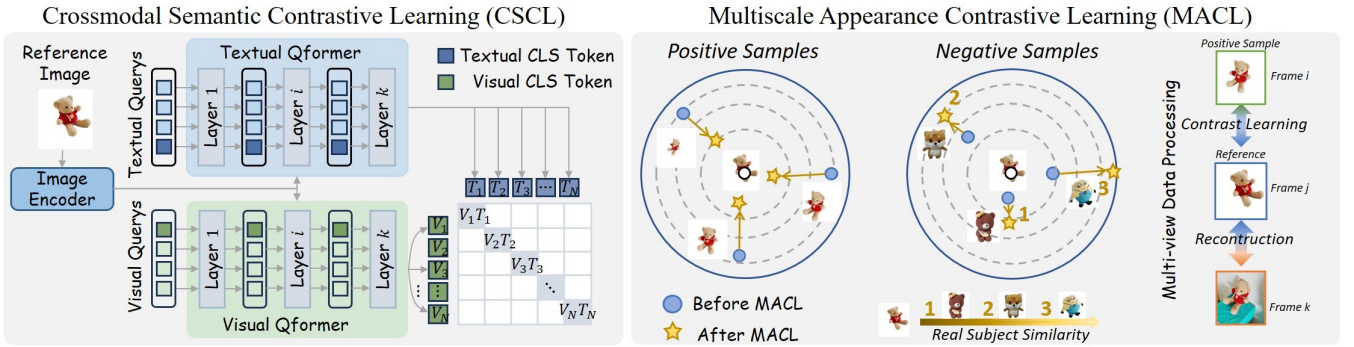
The finetune-free stream aims to eliminate optimization for each subject, achieving real-time generation. Recent work (Wei et al. 2023; Li, Li, and Hoi 2024; Huang et al. 2024; Zhang et al. 2024; Song et al. 2024; Pan et al. 2023; Wang et al. 2024; Patel et al. 2024; Sun et al. 2024; Chen et al. 2023; Huang et al. 2023b,a) trains an additional adapter to map reference images to image prompts. Elite (Wei et al. 2023) uses local mapping to enhance subject features. IP-adapter (Ye et al. 2023) injects visual features into additional cross-attention, struggling with attribute changes (*e.g.*, shape and orientation). RealCustom (Huang et al. 2024; Mao et al. 2024) and SSR-Encoder (Zhang et al. 2024) selectively inject subject-related features to reduce background features. Other studies (Song et al. 2024; Pan et al. 2023) use multimodal large language models for image-text alignment.

2.2 Contrastive Learning of Representations

Early zero-shot classification studies (Khattak et al. 2023; Roy and Etemad 2024) based on CLIP (Radford et al. 2021)



(a) Our training pipeline



(b) Our Multilevel Contrastive Learning (MCL) paradigm

Figure 2: Overview of the proposed *CustomContrast*. (a) Training pipeline. The consistency between textual and visual features is accurately learned by the MFI-Encoder, which includes a Textual-Visual (TV) Fusion module to enhance feature consistency from visual and textual Qformers. (b) The MCL paradigm includes CSCL, aligning high-level semantics by contrasting visual and textual embeddings via CLS tokens, and MACL, which is applied to text embeddings from different cross-attention layers. MACL decouples redundant subject features by aligning positive samples (segmented images of the same subject from various views, positions, and sizes), while preserving relative distances by contrasting with other subjects.

have validated the effectiveness of contrastive learning in aligning visual features with textual features. By dynamically adjusting these representation spaces in downstream tasks, the consistency between visual and textual representations can be further improved. In customized tasks, past work based on the Multimodal Large Language Model (MLLM) (Patel et al. 2024; Pan et al. 2023; Sun et al. 2024) has explored using contrastive learning to reduce the gap between image and text spaces. Summarily, similar to traditional contrastive learning, previous studies in customized task primarily focus on aligning image embeddings with text embeddings, usually at the expense of subject similarity. In contrast, our work uses contrastive learning to extract intrinsic features from redundant features. It decouples redundant features through intra-consistency and further learns intrinsic features through inter-distinctiveness, achieving dual optimization of text controllability and subject similarity.

3 Methodology

The pipeline of *CustomContrast* is depicted in Fig. 2, which mainly contains MFI-Encoder and Multilevel Contrastive Learning (MCL) paradigm. Given reference images dur-

ing training, *CustomContrast* extracts consistent multimodal representations through the MFI-Encoder, which is designed to support the implementation of contrastive learning. Then, the MCL paradigm contrasts these representations from low-level appearance to high-level semantics, extracting intrinsic subject representations. In this section, we will first introduce the preliminaries of the diffusion model. Then, we will describe the MFI-Encoder and MCL paradigm in detail.

3.1 Preliminaries

Our work is based on Stable Diffusion (SD) (Rombach et al. 2022), which includes an autoencoder and a UNet denoiser. The encoder $\mathcal{E}(\cdot)$ of the autoencoder maps a given image $\mathbf{x} \in \mathbb{R}^{H \times W \times 3}$ to a low-dimensional space $\mathbf{z} = \mathcal{E}(\mathbf{x}) \in \mathbb{R}^{h \times w \times d}$, where d denotes the hidden space dimension. The corresponding decoder $\mathcal{D}(\cdot)$ maps \mathbf{z} back to pixel space. In customized task, the text \mathbf{c}_t and the reference image \mathbf{c}_i are used as conditional inputs to the denoiser $\epsilon_\theta(\cdot)$. Ultimately, the denoiser $\epsilon_\theta(\cdot)$ is trained by mean-squared loss:

$$\mathcal{L}_{\text{LDM}} = \mathbb{E}_{\mathbf{z}, \epsilon, t} \|\epsilon - \epsilon_\theta(\mathbf{z}_t, t, \mathbf{c}_t, \mathbf{c}_i)\|_2^2, \quad (1)$$

where ϵ refers to unscaled noise and t means denoising timestep. z_t is the hidden space tensor at t -th timestep.

3.2 Multimodal Feature Injection Encoder

The Multimodal Feature Injection (MFI) Encoder is designed to extract multimodal features (textual and image features), which consists of three main components: Visual-Qformer, Textual-Qformer, and the TV Fusion Module.

The CLIP image encoder is used to extract image features. Previous studies (Wei et al. 2023; Zhang et al. 2024) show that using only the last layer’s feature of CLIP loses fine-grained details. We propose to concatenate features from different CLIP layers to get coarse-to-fine subject features $\mathbf{f}_I = \{\mathbf{f}_k\}_{k=0}^K$, where K refers to layer numbers, defaulting to 3. Then \mathbf{f}_I is fed into textual and visual Qformer.

Here, Qformer denotes the Transformer module where the queries are learnable queries, and image features are input as part of keys and values. PerceiverAttention (PA) (Awadalla et al. 2023) is used in each Qformer to aggregate image features with learnable queries. $PA(\mathbf{f}_I, \mathbf{f}_q)$ aggregates image features \mathbf{f}_I and learnable queries \mathbf{f}_q as follows:

$$\text{Softmax} \left(\frac{Q(\mathbf{f}_q)K([\mathbf{f}_I, \mathbf{f}_q])}{\sqrt{d}} \right) V([\mathbf{f}_I, \mathbf{f}_q]), \quad (2)$$

where $[\mathbf{f}_I, \mathbf{f}_q]$ denotes concatenation of features \mathbf{f}_I and \mathbf{f}_q .

To enable MFI Encoder to dynamically adapt to variations in timesteps and UNet layers l , we use Fourier mapping and learnable embeddings for UNet layer and timestep information to generate spatiotemporal queries $\mathbf{f}_q^{st} \in \mathbb{R}^{(l+1) \times d}$. These queries \mathbf{f}_q^{st} are used by Textual Qformer to query the image features \mathbf{f}_I , mapping subjects to predefined text space as $\mathbf{f}_t \in \mathbb{R}^{(l+1) \times d}$. The first l features of \mathbf{f}_t are respectively fed into the l text cross-attention layers after passing through the text encoder. Differently, vanilla learnable visual queries \mathbf{f}_q^v are used by the visual Qformer to extract visual features $\mathbf{f}_v \in \mathbb{R}^{(m+1) \times d}$, where $m+1$ denotes visual query numbers. The first m features of \mathbf{f}_v are finally injected into each visual cross-attention layer. The remaining single features of \mathbf{f}_t and \mathbf{f}_v are used for CSCL in Section 3.3.

TV Fusion Module To further align visual features \mathbf{f}_v with text space, we use the $PA(\mathbf{f}_t, \mathbf{f}_v)$ module in Eq. 2 to extract refined visual features $\hat{\mathbf{f}}_v$ related to text features \mathbf{f}_t . Finally, $\hat{\mathbf{f}}_v$ and \mathbf{f}_t are injected into parallel visual and text cross-attention. The attention integration is as follows:

$$\text{Attention}(Q, K_t, V_t, K_v, V_v) = \text{Softmax} \left(\frac{QK_t^\top}{\sqrt{d}} \right) V_t + \gamma \cdot \text{Softmax} \left(\frac{QK_v^\top}{\sqrt{d}} \right) V_v, \quad (3)$$

where $Q = W_q z_t$, $K_v = W_k^v \hat{\mathbf{f}}_v$, $K_t = W_k^t \hat{\mathbf{f}}_t$, $V_v = W_v^v \hat{\mathbf{f}}_v$, and $V_t = W_v^t \hat{\mathbf{f}}_t$. W_q , W_k^t , W_v^t are frozen, while W_k^v and W_v^v are trainable projection layers for visual cross attention. z_t denotes image latents and $\hat{\mathbf{f}}_t$ denotes text feature generated by feeding \mathbf{f}_t and the corresponding caption into the text encoder. γ is a weight factor, defaulting to 1.

3.3 Multilevel Contrastive Learning Paradigm

Cross-differential perspective aims to capture subjects’ distinguished intrinsic features by contrasting multilevel differences between subjects. To achieve this, we propose a Multilevel Contrastive Learning paradigm: Crossmodal Semantic Contrastive Learning aligns the textual and visual features for semantic intrinsic consistency, and Multiscale Appearance Contrastive Learning ensures feature distances are consistent with those of real subjects across different scales.

Crossmodal Semantic Contrastive Learning Features of subjects in the CLIP text and image space are extracted through Visual and Textual Qformer, respectively. However, a gap often exists between these features, making it challenging for CLIP image features to respond sensitively to complex transformations (e.g., shape transformations).

To address this, we propose Crossmodal Semantic Contrastive Learning (CSCL) to align the features generated by the Visual Qformer and Textual Qformer, ensuring they are in the same space, as shown in Fig. 2(b)(left). We add a visual CLS token v^{cls} and a textual CLS token t^{cls} to the visual queries and textual queries, applying CSCL loss \mathcal{L}_c on them, which is as follows:

$$\mathcal{L}_c = - \sum_{i \in N} \log \left[\frac{\exp(a_{ii})}{\sum_{j \neq i} \exp(a_{ji})} \right] - \sum_{i \in N} \log \left[\frac{\exp(a_{ii})}{\sum_{j \neq i} \exp(a_{ij})} \right], \quad (4)$$

where i and j donate two different samples within a batch, N donates batch size. a_{ij} represents $\cos(\mathbf{f}_{v_i}^{cls}, \mathbf{f}_{t_j}^{cls})$, measuring consistency between $\mathbf{f}_{v_i}^{cls}$ and $\mathbf{f}_{t_j}^{cls}$ by cosine similarity. $\mathbf{f}_{v_i}^{cls}$ refers to the remaining CLS features of i -th sample’s \mathbf{f}_v , and $\mathbf{f}_{t_j}^{cls}$ refers to the remaining CLS feature of j -th sample’s \mathbf{f}_t .

Multiscale Appearance Contrastive Learning The key idea of Multi-scale Appearance Similarity Contrastive Learning (MACL) is to ensure that the distance relationships between multiscale features are consistent with those of real subjects. This means the features of the same subject with different situations should be as close as possible (intra-consistency), while the distances between different samples’ features should match those between real subjects (inter-distinctiveness). As shown in Fig. 2(b)(right), we achieve intra-consistency by pulling positive samples of the reference subject closer, and inter-distinctiveness by introducing scaling factors to align the feature distances with negative samples to real subject distances. In this section, we will introduce the S^+ Space and MACL in the S^+ Space.

A. S^+ Space

Injecting token $\mathbf{s}^* \in \mathbb{R}^{k \times d}$ into textual space is used for finetune-based studies (Ruiz et al. 2023) as shown in Fig. 3(a). Here, k and d refer to token numbers and dimensions. Define S as the space spanned by \mathbf{s}^* . We then define extended visual concept space S^+ as follows:

$$S^+ := [\mathbf{s}_1^*, \mathbf{s}_2^*, \dots, \mathbf{s}_l^*], \quad (5)$$




Figure 3: (a) In S space, a token s^* influences all cross-attention layers. (b) In S^+ space, different s_i^* control cross-attention layers. (c) MACL is applied separately to each s_i^* .

where s_i^* denotes tokens for i -th cross-attention layer and l denotes cross-attention layer numbers. $s^+ \in \mathbb{R}^{l \times k \times d}$ is a tensor in the S^+ space. As depicted in Fig. 3(b), l features of s^+ are injected into l cross-attention. MACL is applied to s_i^* for different cross-attention as shown in Fig. 3(c).

B. MACL in S^+ Space

MACL in the S^+ space preserves the multi-scale similarity structure, ensuring the similarities of learned features are positively correlated with those of real subjects. Thus, the sample similarities must satisfy the following conditions:

$$\mathcal{D}_r(\mathbf{x}_i, \mathbf{x}_j) < \mathcal{D}_r(\mathbf{x}_i, \mathbf{x}_k) \Leftrightarrow \mathcal{D}_g(\mathbf{z}_i, \mathbf{z}_j) < \mathcal{D}_g(\mathbf{z}_i, \mathbf{z}_k), \quad (6)$$

where $\mathcal{D}_r(\cdot)$ denotes the ideal distance metric for real subjects \mathbf{x} , and $\mathcal{D}_g(\cdot)$ is relative distance metric for learned features \mathbf{z} . i, j and k denote the i -th, j -th, and k -th samples. We use the CLIP image encoder $\psi_C(\cdot)$ to extract appearance features of real subjects, approximating relative distance \mathcal{D}_r between real subjects. Note that $\psi_C(\cdot)$ is a specific case to approximate \mathcal{D}_r and can be replaced by others (e.g., DINO-V2 (Sinhamahapatra et al. 2024), SigLIP (Zhai et al. 2023)).

We design MACL scaling factors to implement the aforementioned constraints. These factors scale the generated samples' similarity based on the real subjects' similarity. We use cosine similarity, denoted as $\cos(\cdot)$, to measure the similarity between the s^+ components of different samples at all l cross-attention layers. The appearance representation of the segmented subject images $\hat{\mathbf{x}}$, obtained by CLIP image encoder, is utilized to compute the appearance similarity matrix $R_{ap} = \{r_{ij}\}_{i,j=1}^{N \times N}$, where $r_{ij} = \cos(\psi_C(\hat{\mathbf{x}}_i), \psi_C(\hat{\mathbf{x}}_j))$. Here, segmented subject images $\hat{\mathbf{x}}$ approximate the real subjects \mathbf{x} , allowing MACL to focus on the subjects themselves rather than the background. The appearance scaling factor is $\alpha_{ij} = 1/r_{ij}$. The MACL loss \mathcal{L}_m^{ij} between s_i^+ and s_j^+ extracted by Textual Qformer is:

$$\mathcal{L}_m^{ij} = - \sum_{a=0}^l \log \left[\frac{\exp(\alpha_{ij} \cdot \cos_{ij}^a / \tau)}{\sum_{k \neq i}^N \exp(\alpha_{ik} \cdot \cos_{ik}^a / \tau)} \right], \quad (7)$$

where \cos_{ij}^a and \cos_{ik}^a are abbreviation for $\cos(s_{ia}^+, s_{ja}^+)$ and $\cos(s_{ia}^+, s_{ka}^+)$. s_{ia}^+ , s_{ja}^+ and s_{ka}^+ respectively represent components of the a -th cross attention layer for i -th sample's s_i^+ , j -th sample's s_j^+ , and k -th sample's s_k^+ , τ denotes temperature and N dnotes batch size. The scaling factor α_{ij} aligns

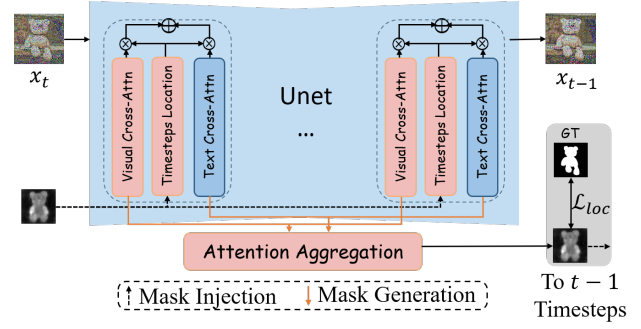


Figure 4: Illustration of timesteps-specific subject location.

relative distance relationships between learned features and real subjects. The detailed derivation is in Appendix A.1.

As shown in Fig. 2(b) (right), We select frames different from the reference images as MACL positive samples. By aligning images of the same subject, *CustomContrast* effectively decouples irrelevant features of the subject. The processing details of positive samples are in Appendix B.

3.4 Timesteps-specific Subject Location

Similar to MOMA(Song et al. 2024), we adopt timestep injection inference operations. The main difference is that we use location loss (Xiao et al. 2023) to further reduce redundant features' interference, as shown in Fig. 4. During training, we extract all visual cross-attention maps A_v and p -th text cross-attention maps A_t^p related to S^+ space (where p is the placeholder index). The global attention map \bar{A}_g is then aggregated by averaging these attention maps. \bar{A}_g is constrained with the actual mask M_g by loss function \mathcal{L}_{loc} (Xiao et al. 2023):

$$\mathcal{L}_{loc} = \text{mean}((1 - M_g) \odot (1 - \bar{A}_g)) - \text{mean}(M_g \odot \bar{A}_g), \quad (8)$$

where \odot denotes the element-wise multiplication. During inference, the extracted mask at timestep $t+1$ is injected into the visual attention maps and the p -th text attention maps at timestep t , as shown in Fig. 4. Then, the location mask generated at timestep t is injected in the next timestep $t-1$, and so on. The visualization details are in Appendix C.7.

The total loss is as follows:

$$\mathcal{L}_{total} = \mathcal{L}_{LDM} + \lambda_1(\mathcal{L}_c + \mathcal{L}_m) + \lambda_2(\mathcal{L}_{loc}), \quad (9)$$

where λ_1 and λ_2 are loss hyper-parameters.

4 Experiments

4.1 Experimental Setups

Implementation. Our model is implemented on both SD V1-5 and SDXL. The filtered subset of MVImageNet (Yu et al. 2023) and OpenImage (Kuznetsova et al. 2020) are used as training sets. The model is trained on 6 A100 GPUs for 200k iterations with learning rate $3e-5$ and $\lambda_1=1e-2$ and $\lambda_2=1e-3$. The layer numbers of Textual and Visual Qformer are set to 4. More details are in Appendix C.1.

| Methods | <i>controllability</i> | | <i>similarity</i> | | | |
|-----------------------|------------------------|---------------|-------------------|-----------------|-----------------|-----------------|
| | CT \uparrow | IR \uparrow | R-CI \uparrow | E-CI \uparrow | R-DI \uparrow | E-DI \uparrow |
| Custom Diffusion | 0.258 | -1.41 | 0.863 | 0.756 | 0.661 | 0.517 |
| DreamMatcher | 0.298 | 0.16 | 0.875 | 0.781 | 0.670 | 0.540 |
| BLIP-Diffusion | 0.278 | -0.80 | 0.865 | 0.768 | 0.629 | 0.535 |
| ELITE | 0.291 | -0.66 | 0.856 | 0.762 | 0.661 | 0.533 |
| IP-Adapter-plus | 0.295 | -0.61 | 0.878 | 0.766 | 0.702 | 0.558 |
| SSR-Encoder | 0.303 | -0.89 | 0.857 | 0.767 | 0.623 | 0.524 |
| KOSMOS-G | 0.268 | -1.45 | 0.864 | 0.753 | 0.677 | 0.531 |
| MOMA | 0.313 | -0.24 | 0.879 | 0.772 | 0.679 | 0.550 |
| ours (SD-v1.5) | 0.325 | 0.45 | 0.919 | 0.788 | 0.737 | 0.591 |

Table 1: Quantitative results with SD-v1.5 based methods. Bolded numbers indicate the best performance.

| Methods | <i>controllability</i> | | <i>similarity</i> | | | |
|--------------------|------------------------|---------------|-------------------|-----------------|-----------------|-----------------|
| | CT \uparrow | IR \uparrow | R-CI \uparrow | E-CI \uparrow | R-DI \uparrow | E-DI \uparrow |
| λ -ECLIPSE | 0.279 | -1.03 | 0.882 | 0.773 | 0.677 | 0.511 |
| IP-Adapter-XL | 0.289 | -0.71 | 0.918 | 0.790 | 0.723 | 0.570 |
| EMU2 | 0.298 | -0.203 | 0.91 | 0.795 | 0.719 | 0.583 |
| MS-Diffusion | 0.312 | -0.24 | 0.901 | 0.772 | 0.682 | 0.540 |
| ours (SDXL) | 0.329 | 0.45 | 0.927 | 0.796 | 0.742 | 0.597 |

Table 2: Quantitative results with more advanced base models, i.e., SDXL and kv2.2 (Razzhigaev et al. 2023).

Test Prompt. To evaluate complex scenarios, we create a prompt set with four categories: style, shape, scene, and accessorization modification. See Appendix C.2 for details.

Evaluation metrics. We evaluate our model based on DINO-I (DI) (Oquab et al. 2023), CLIP-I (CI), and CLIP-T (CT) scores on DreamBench (Ruiz et al. 2023). **Similarity.** To exclude background interference, we calculate the subject similarity after segmenting the reference subjects and generated subjects with Grounded SAM (Ren et al. 2024). Similarity metrics are refined by measuring subject similarity in reconstruction (R-CI and R-DI) and editing (E-CI and E-DI). **Controllability.** CLIP-T metric is calculated using cosine similarity in the CLIP text-image embedding space. Additionally, ImageReward (IR) (Xu et al. 2024) is also used to evaluate controllability.

4.2 Main Results

In this section, We will compare *CustomContrast* with previous SOTAs through qualitative and quantitative analyses.

Quantitative results. *CustomContrast* is compared with previous SOTAs on both SD-v1.5 and SDXL. Details of SOTAs are in Appendix C.1. As shown in Tab. 1, our model demonstrates significant improvements in both text controllability and image similarity compared to previous SD-V1.5 methods, with a notable increase in the ImageReward metric from 0.157 to 0.454. The quantitative comparison for SDXL, as shown in Tab. 2, indicates that our method also improves both controllability and similarity, with the CLIP-T increasing from 0.312 to 0.329.

Qualitative results. Fig. 5 illustrates the qualitative results compared with existing methods. *CustomContrast* (SDXL) demonstrates high fidelity and controllability,

| ID | Visual | Textual | TV Fusion | E-CI \uparrow | E-DI \uparrow | CT \uparrow |
|----|--------------|--------------|--------------|-----------------|-----------------|---------------|
| 0 | \checkmark | | | 0.731 | 0.554 | 0.302 |
| 1 | \checkmark | \checkmark | | 0.765 | 0.568 | 0.282 |
| 2 | \checkmark | \checkmark | \checkmark | 0.754 | 0.563 | 0.311 |

Table 3: Ablation studies on MFI-Encoder. "Visual" and "Textual" refer to Qformer types. Compared with ID-0, the MFI-encoder (ID-2) improves similarity and controllability.

| ID | MFI-encoder | CSCL | MACL | E-CI \uparrow | E-DI \uparrow | CT \uparrow |
|----|--------------|--------------|--------------|-----------------|-----------------|---------------|
| 2 | \checkmark | | | 0.754 | 0.563 | 0.311 |
| 3 | \checkmark | \checkmark | | 0.765 | 0.575 | 0.323 |
| 4 | \checkmark | | \checkmark | 0.780 | 0.584 | 0.318 |
| 5 | \checkmark | \checkmark | \checkmark | 0.788 | 0.591 | 0.325 |

Table 4: Ablation studies on the effectiveness of each module based on SD-v1.5. With both MACL and CSCL losses combined (ID-5), *CustomContrast* achieves dual optimization in text controllability and image similarity.

achieving superior zero-shot customization capabilities even for complex accessorization, shape, and style modification, where other models underperform (e.g., a cat toy playing the guitar in the first row, a driver turning into a cat in the fourth row, and watercolor style in the fifth row). Qualitative results of *CustomContrast* (SD-v1.5) are in Appendix C.3.

To further demonstrate *CustomContrast*'s capabilities in complex editing, we present qualitative results in multi-object and human domain generation. The qualitative results for the multi-subject task are shown in Fig. 6. Our model can preserve the multiple subject similarities while achieving superior editing capabilities (e.g., only the cat toy wears a hat in the first row; only the cat wears glasses in the third row). Human domain qualitative results are in Appendix C.5.

4.3 Ablation Studies

To demonstrate the effectiveness of essential components of *CustomContrast*, we conduct extensive ablation experiments. All ablation experiments are trained on SD-v1.5. We present the key ablation study in the main text, with the remaining ablation experiments shown in Appendix C.6.

Effectiveness of Each Module for MFI-Encoder. As indicated in Tab. 3, compared to using only Visual-Qformer (ID-0), introducing Textual-Qformer (ID-1) further map the subject features to the S^+ space, enhancing subject similarity. However, simply adding Textual-Qformer may reduce text controllability, due to the lack of unified constraints between the visual and textual spaces. To address this issue, we introduce the TV-Fusion module (ID-2), which extracts text-related visual features, simultaneously improving both text controllability and image similarity compared to ID-0.

Effectiveness of Contrastive Learning Paradigm. As indicated in Tab. 4, the contrastive learning paradigm achieves dual optimization in text controllability and image similarity. Compared to the baseline without contrastive learning constraints (ID-2), applying only CSCL (ID-3) significantly improves text controllability. MACL (ID-4) makes

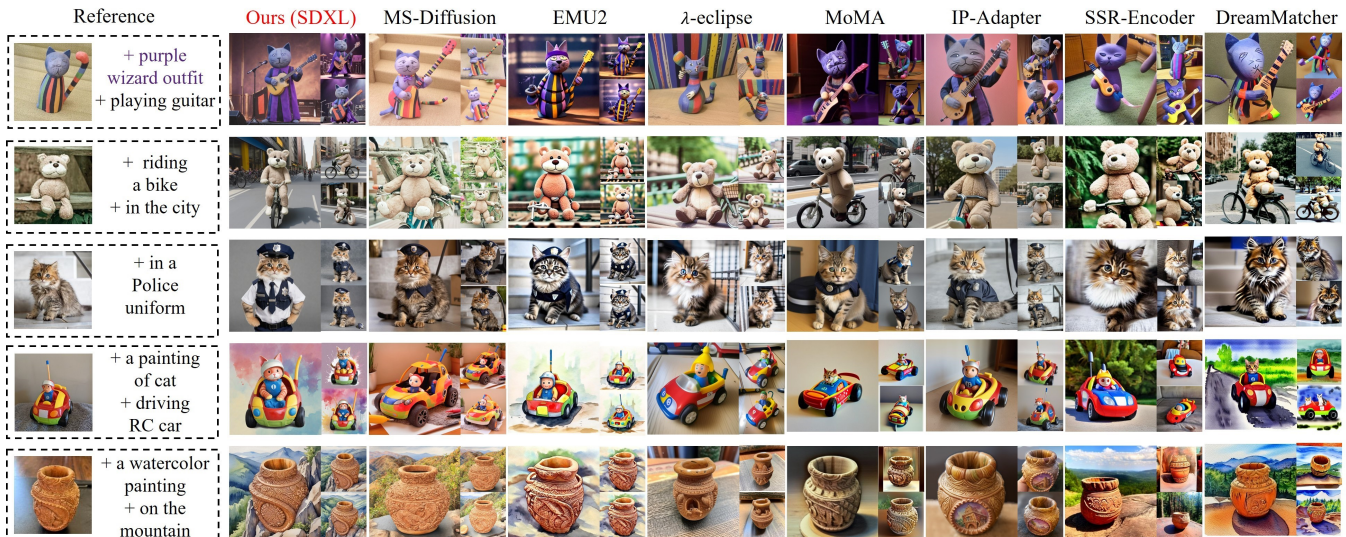


Figure 5: **Qualitative comparison** with existing methods. *CustomContrast* decouples intrinsic features from redundant features, enabling flexible text control over complex pose (e.g., the cat toy in the first row) and shape (e.g., cat driving car in the fourth row) transformations. In contrast, other methods underperform due to the influence of coupled redundant features.

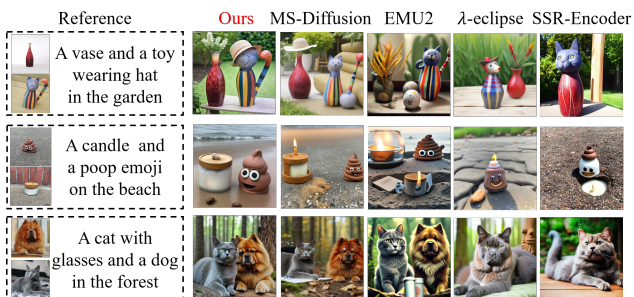


Figure 6: **Qualitative results** of multi-subject generation.

the model focus on the subject essential features, thereby achieving better similarity. Meanwhile, due to the decoupling of irrelevant features, the model can better respond to the text. Ultimately, with both losses (ID-5), *CustomContrast* achieves dual optimization in controllability and similarity, increasing text controllability from 0.311 to 0.325.

4.4 Visualization

To further illustrate our model’s superiority, we use t-SNE to analyze the subject representations of ELITE (Wei et al. 2023) and our model in Fig. 7. We use dogs with front and side views from DogFaceNet (Mougeot, Li, and Jia 2019) as samples. Different shapes represent different dogs, and different colors indicate different views. In Fig. 7(b), the ELITE representations are more entangled by view interference. In contrast, s^+ learned by *CustomContrast* better preserves relative distances between samples in Fig. 7(c), with different views of the same dog being closer. These analyses demonstrate our model decouples view redundant features from intrinsic subject features.

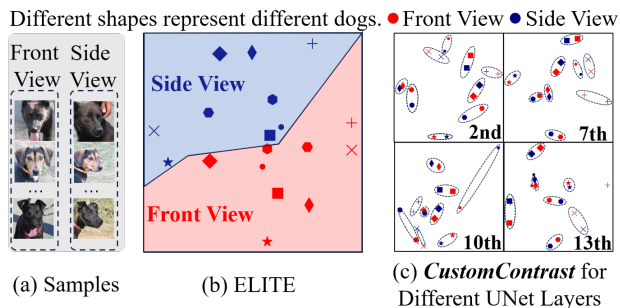


Figure 7: (a) Dog’s view examples. (b) T-SNE of ELITE representations. (c) T-SNE of our representations, where 2nd represents the second cross attention, and so on. Ideal representations should achieve intra-consistency (i.e., points with the same shape should be spatially closer in the figure). ELITE is influenced by different views, with view features coupled into the representations. However, *CustomContrast* decouples view features from the intrinsic features, preserving relative distances between samples.

5 Conclusion

In this paper, we propose *CustomContrast*, a novel paradigm for customized T2I from a cross-differential perspective. *CustomContrast* introduces a novel Multilevel Contrastive Learning paradigm that captures intrinsic subject representations from high-level semantics to low-level appearance, supported by the MFI-encoder for consistent cross-modal representations. Extensive experiments demonstrate the superiority of our model, effectively achieving high subject similarity and textual controllability, while easily extending to multi-subject and the human domain generation.

Acknowledgements

This research is supported by Artificial Intelligence-National Science and Technology Major Project 2023ZD0121200, National Natural Science Foundation of China under Grant 62222212 and National Natural Science Foundation of China under Grant 623B2094.

References

- Alaluf, Y.; Richardson, E.; Metzger, G.; and Cohen-Or, D. 2023. A neural space-time representation for text-to-image personalization. *ACM Transactions on Graphics (TOG)*, 42(6): 1–10.
- Awadalla, A.; Gao, I.; Gardner, J.; Hessel, J.; Hanafy, Y.; Zhu, W.; Marathe, K.; Bitton, Y.; Gadre, S.; Sagawa, S.; Jitsev, J.; Kornblith, S.; Koh, P. W.; Ilharco, G.; Wortsman, M.; and Schmidt, L. 2023. OpenFlamingo: An Open-Source Framework for Training Large Autoregressive Vision-Language Models. *arXiv preprint arXiv:2308.01390*.
- Chen, Z.; Fang, S.; Liu, W.; He, Q.; Huang, M.; Zhang, Y.; and Mao, Z. 2023. Dreamidentity: Improved editability for efficient face-identity preserved image generation. *arXiv preprint arXiv:2307.00300*.
- Gal, R.; Alaluf, Y.; Atzmon, Y.; Patashnik, O.; Bermano, A. H.; Chechik, G.; and Cohen-Or, D. 2022. An image is worth one word: Personalizing text-to-image generation using textual inversion. *arXiv preprint arXiv:2208.01618*.
- Han, L.; Li, Y.; Zhang, H.; Milanfar, P.; Metaxas, D.; and Yang, F. 2023. Svdiff: Compact parameter space for diffusion fine-tuning. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 7323–7334.
- Ho, J.; Jain, A.; and Abbeel, P. 2020. Denoising diffusion probabilistic models. *Advances in neural information processing systems*, 33: 6840–6851.
- Huang, M.; Mao, Z.; Chen, Z.; and Zhang, Y. 2023a. Towards accurate image coding: Improved autoregressive image generation with dynamic vector quantization. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 22596–22605.
- Huang, M.; Mao, Z.; Liu, M.; He, Q.; and Zhang, Y. 2024. RealCustom: Narrowing Real Text Word for Real-Time Open-Domain Text-to-Image Customization. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 7476–7485.
- Huang, M.; Mao, Z.; Wang, P.; Wang, Q.; and Zhang, Y. 2022. Dse-gan: Dynamic semantic evolution generative adversarial network for text-to-image generation. In *Proceedings of the 30th ACM International Conference on Multimedia*, 4345–4354.
- Huang, M.; Mao, Z.; Wang, Q.; and Zhang, Y. 2023b. Not all image regions matter: Masked vector quantization for autoregressive image generation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2002–2011.
- Khattak, M. U.; Rasheed, H.; Maaz, M.; Khan, S.; and Khan, F. S. 2023. Maple: Multi-modal prompt learning. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 19113–19122.
- Kumari, N.; Zhang, B.; Zhang, R.; Shechtman, E.; and Zhu, J.-Y. 2023. Multi-Concept Customization of Text-to-Image Diffusion. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*.
- Kuznetsova, A.; Rom, H.; Alldrin, N.; Uijlings, J.; Krasin, I.; Pont-Tuset, J.; Kamali, S.; Popov, S.; Mallocci, M.; Kolesnikov, A.; et al. 2020. The open images dataset v4: Unified image classification, object detection, and visual relationship detection at scale. *International Journal of Computer Vision*, 128(7): 1956–1981.
- Li, D.; Li, J.; and Hoi, S. 2024. Blip-diffusion: Pre-trained subject representation for controllable text-to-image generation and editing. *Advances in Neural Information Processing Systems*, 36.
- Mao, Z.; Huang, M.; Ding, F.; Liu, M.; He, Q.; Chang, X.; and Zhang, Y. 2024. RealCustom++: Representing Images as Real-Word for Real-Time Customization. *arXiv preprint arXiv:2408.09744*.
- Mougeot, G.; Li, D.; and Jia, S. 2019. A Deep Learning Approach for Dog Face Verification and Recognition. In Nayak, A. C.; and Sharma, A., eds., *PRICAI 2019: Trends in Artificial Intelligence*, 418–430. Cham: Springer International Publishing. ISBN 978-3-030-29894-4.
- Nam, J.; Kim, H.; Lee, D.; Jin, S.; Kim, S.; and Chang, S. 2024. DreamMatcher: Appearance Matching Self-Attention for Semantically-Consistent Text-to-Image Personalization. *arXiv:2402.09812*.
- Oquab, M.; Darcet, T.; Moutakanni, T.; Vo, H.; Szafraniec, M.; Khalidov, V.; Fernandez, P.; Haziza, D.; Massa, F.; El-Nouby, A.; et al. 2023. Dinov2: Learning robust visual features without supervision. *arXiv preprint arXiv:2304.07193*.
- Pan, X.; Dong, L.; Huang, S.; Peng, Z.; Chen, W.; and Wei, F. 2023. Kosmos-G: Generating Images in Context with Multimodal Large Language Models. *ArXiv*, abs/2310.02992.
- Patel, M.; Jung, S.; Baral, C.; and Yang, Y. 2024. λ -ECLIPSE: Multi-Concept Personalized Text-to-Image Diffusion Models by Leveraging CLIP Latent Space. *arXiv preprint arXiv:2402.05195*.
- Radford, A.; Kim, J. W.; Hallacy, C.; Ramesh, A.; Goh, G.; Agarwal, S.; Sastry, G.; Askell, A.; Mishkin, P.; Clark, J.; et al. 2021. Learning transferable visual models from natural language supervision. In *International conference on machine learning*, 8748–8763. PMLR.
- Razhigayev, A.; Shakhmatov, A.; Maltseva, A.; Arkhipkin, V.; Pavlov, I.; Ryabov, I.; Kuts, A.; Panchenko, A.; Kuznetsov, A.; and Dimitrov, D. 2023. Kandinsky: An improved text-to-image synthesis with image prior and latent diffusion. In *Proceedings of the 2023 Conference on Empirical Methods in Natural Language Processing: System Demonstrations*, 286–295.
- Ren, T.; Liu, S.; Zeng, A.; Lin, J.; Li, K.; Cao, H.; Chen, J.; Huang, X.; Chen, Y.; Yan, F.; et al. 2024. Grounded sam: Assembling open-world models for diverse visual tasks. *arXiv preprint arXiv:2401.14159*.
- Rombach, R.; Blattmann, A.; Lorenz, D.; Esser, P.; and Ommer, B. 2022. High-resolution image synthesis with latent

- diffusion models. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 10684–10695.
- Roy, S.; and Etemad, A. 2024. Consistency-guided Prompt Learning for Vision-Language Models. In *Proceedings of the International Conference on Learning Representations (ICLR)*.
- Ruiz, N.; Li, Y.; Jampani, V.; Pritch, Y.; Rubinstein, M.; and Aberman, K. 2023. Dreambooth: Fine tuning text-to-image diffusion models for subject-driven generation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 22500–22510.
- Sinhamahapatra, P.; Schwaiger, F.; Bose, S.; Wang, H.; Roscher, K.; and Guennemann, S. 2024. Finding Dino: A plug-and-play framework for unsupervised detection of out-of-distribution objects using prototypes. *arXiv preprint arXiv:2404.07664*.
- Song, K.; Zhu, Y.; Liu, B.; Yan, Q.; Elgammal, A.; and Yang, X. 2024. MoMA: Multimodal LLM Adapter for Fast Personalized Image Generation. *arXiv preprint arXiv:2404.05674*.
- Sun, Q.; Cui, Y.; Zhang, X.; Zhang, F.; Yu, Q.; Wang, Y.; Rao, Y.; Liu, J.; Huang, T.; and Wang, X. 2024. Generative multimodal models are in-context learners. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 14398–14409.
- Voynov, A.; Chu, Q.; Cohen-Or, D.; and Aberman, K. 2023. *P+*: Extended Textual Conditioning in Text-to-Image Generation. *arXiv preprint arXiv:2303.09522*.
- Wang, X.; Fu, S.; Huang, Q.; He, W.; and Jiang, H. 2024. MS-Diffusion: Multi-subject Zero-shot Image Personalization with Layout Guidance. *arXiv preprint arXiv:2406.07209*.
- Wei, Y.; Zhang, Y.; Ji, Z.; Bai, J.; Zhang, L.; and Zuo, W. 2023. Elite: Encoding visual concepts into textual embeddings for customized text-to-image generation. *arXiv preprint arXiv:2302.13848*.
- Xiao, G.; Yin, T.; Freeman, W. T.; Durand, F.; and Han, S. 2023. Fastcomposer: Tuning-free multi-subject image generation with localized attention. *arXiv preprint arXiv:2305.10431*.
- Xu, J.; Liu, X.; Wu, Y.; Tong, Y.; Li, Q.; Ding, M.; Tang, J.; and Dong, Y. 2024. Imagereward: Learning and evaluating human preferences for text-to-image generation. *Advances in Neural Information Processing Systems*, 36.
- Ye, H.; Zhang, J.; Liu, S.; Han, X.; and Yang, W. 2023. Ip-adapter: Text compatible image prompt adapter for text-to-image diffusion models. *arXiv preprint arXiv:2308.06721*.
- Yu, X.; Xu, M.; Zhang, Y.; Liu, H.; Ye, C.; Wu, Y.; Yan, Z.; Zhu, C.; Xiong, Z.; Liang, T.; et al. 2023. Mvimnet: A large-scale dataset of multi-view images. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 9150–9161.
- Zhai, X.; Mustafa, B.; Kolesnikov, A.; and Beyer, L. 2023. Sigmoid loss for language image pre-training. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 11975–11986.
- Zhang, Y.; Dong, W.; Tang, F.; Huang, N.; Huang, H.; Ma, C.; Lee, T.-Y.; Deussen, O.; and Xu, C. 2023. Prospect: Prompt spectrum for attribute-aware personalization of diffusion models. *ACM Transactions on Graphics (TOG)*, 42(6): 1–14.
- Zhang, Y.; Song, Y.; Liu, J.; Wang, R.; Yu, J.; Tang, H.; Li, H.; Tang, X.; Hu, Y.; Pan, H.; et al. 2024. Ssr-encoder: Encoding selective subject representation for subject-driven generation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 8069–8078.