

Learning Disentangled Identifiers for Action-Customized Text-to-Image Generation

Siteng Huang^{1,4*}, Biao Gong², Yutong Feng², Xi Chen², Yuqian Fu³, Yu Liu², Donglin Wang^{4†}

¹Zhejiang University ²Alibaba Group ³ETH Zürich

⁴Machine Intelligence Lab (MiLAB), AI Division, School of Engineering, Westlake University

{siteng.huang, a.biao.gong, yuqianfu0207}@gmail.com

{fengyutong.fyt, xizhi.cx, ly103369}@alibaba-inc.com wangdonglin@westlake.edu.cn



Figure 1. **Action customization results of our ADI method.** By inverting representative action-related features, the learned identifiers “<A>” can be paired with a variety of characters and animals to contribute to the generation of accurate, diverse and high-quality images.

Abstract

This study focuses on a novel task in text-to-image (T2I) generation, namely **action customization**. The objective of this task is to learn the co-existing action from limited data and generalize it to unseen humans or even animals. Experimental results show that existing subject-driven customization methods fail to learn the representative characteristics of actions and struggle in decoupling actions from context features, including appearance. To overcome the preference for low-level features and the entanglement of high-level features, we propose an inversion-based method **Action-Disentangled Identifier (ADI)** to learn action-specific iden-

tifiers from the exemplar images. ADI first expands the semantic conditioning space by introducing layer-wise identifier tokens, thereby increasing the representational richness while distributing the inversion across different features. Then, to block the inversion of action-agnostic features, ADI extracts the gradient invariance from the constructed sample triples and masks the updates of irrelevant channels. To comprehensively evaluate the task, we present an **Action-Bench** that includes a variety of actions, each accompanied by meticulously selected samples. Both quantitative and qualitative results show that our ADI outperforms existing baselines in action-customized T2I generation. Our project page is at <https://adi-t2i.github.io/ADI>.

*Work done during internship at Alibaba Group.

†Corresponding author.




Figure 2. **Action customization results of existing subject-driven customization methods.** Due to the preference to search for low-level invariants when asked to learn high-level action features, some methods fail to generate the specified actions, while others confuse the animals with human appearances.

1. Introduction

Thanks to the remarkable advances in text-to-image generation models [3, 10, 18, 20], in particular the recent diffusion model [22, 26], high-quality and diverse images can be synthesized under the control of text descriptions. However, it is difficult to provide precise descriptions of the desired actions, which are highly abstracted and summarized concepts. Therefore, relying solely on textual descriptions to generate actions tends to reduce fidelity to user requirements. Additionally, controllable generation methods [14, 35] that rely on the conditioning of a skeleton or sketch image suffer from limited diversity and freedom, and they show difficulty generalizing to unseen subjects without retraining. In this paper, we study the **action customization** task, capturing the common action in the given images to generate new images with various new subjects.

To better understand the challenge of action customization, we start by examining existing subject-driven customization methods. Observations shown in Fig. 2 can be divided into two categories. Several methods including DreamBooth [25], Textual Inversion [5], and ReVersion [7], generate images that are unrelated to specific actions, suggesting that they fail to capture the representative characteristics of the actions. Since most of them are designed to invert appearance features with a pixel-level reconstruction loss, low-level details are emphasized during optimization while high-level action features are neglected. Benefiting from fine-tuning cross-attention or utilizing per-layer tokens, Custom Diffusion [9] and P+ [30] offer a larger semantic conditioning space for learning new concepts. Consequently, they are capable of encoding action-related knowledge such as “*raises one finger*” or “*raises both arms for cheering*” from exemplar images. However, they fail to decouple the focus from action-agnostic features, such as the appearance of the human body. These pieces of information are also encoded into the learned identifiers and “contaminate” the generation of animals during inference. As a result, the intended gorilla is replaced by a

woman, and the tigers generated by the two methods exhibit human arms instead.

To avoid the appearance leakage while accurately modeling the target action, we propose **Action-Disentangled Identifier (ADI)** to learn the optimal action-specific identifiers. Firstly, we expand the semantic conditioning space by applying layer-wise identifier tokens. Since existing works have analyzed that different layers have varying degrees of control over low-level and high-level features [30], such an expansion increases the accommodation of various features, making it easier to invert action-related features. Furthermore, we would like to decouple the action-agnostic features from the learning of action identifiers. To achieve this, we discover invariant mechanisms in the data that are difficult to vary across examples. Specifically, given an exemplar image with the specific action, another same-action image can be randomly sampled from the training data, forming a context-different pair. Meanwhile, leveraging mature subject-driven customization techniques, an image that shares the similar context can be quickly synthesized to form an action-different pair. To decouple the highly-coupled features, we disentangle action-agnostic features at the gradient level, and construct two context gradient masks by comparing the difference on the gradients over the input pairs. By overwriting the merged gradient mask to the gradient of the anchor image, the update of action-agnostic channels on the identifiers is discarded.

Moreover, as a pioneering effort in this direction, we also contribute to a new benchmark named **ActionBench**, which provides a testbed of unique actions with diverse images for the under-explored task. We conduct extensive experiments on ActionBench, and a quick glance at the performance of ADI is illustrated in Fig. 1, where users can freely combine the designated action identifiers with various unseen humans and even animals. In summary, the main contributions of our work are three-fold:

- We propose a novel action customization task, which requires learning the desired action from limited data for future generation. While existing customization fo-

cuses on reprinting appearances, we highlight this understudied but important problem.

- We contribute the ActionBench, where a variety of unique actions with manually filtered images provide the evaluation conditions for the task.
- We devise the Action-Disentangled Identifier (ADI) method, which successfully inverts action-related features into the learned identifiers that can be freely combined with various characters and animals to generate high-quality images.

2. Related Work

Text-to-Image (T2I) Generation. Generating high-quality and diverse images from textual conditions has received considerable attention from both the research community and the general public. The previous dominant generative adversarial networks (GANs) [20, 28, 31, 38], consisting of a generator and a discriminator, suffer from unstable optimization and less diverse generations due to the adversarial training [2]. And variational autoencoders (VAEs) [8], which apply a probabilistic encoder-decoder architecture, are also prone to posterior collapse and over-smoothed generations [37]. Text-conditional auto-regressive models [3, 4, 18, 34] have shown more impressive results, but require time-consuming iterative processes to achieve high-quality image sampling. More recently, diffusion models have emerged as a promising alternative, achieving impressive results with open-vocabulary text descriptions through their natural fitting to inductive biases of image data [15, 19, 22, 26]. GLIDE [15] introduces text conditions into the diffusion process through the use of an unclassified guide. DALL-E 2 [19] employs a diffusion prior module and cascading diffusion decoders to generate high-resolution images based on the CLIP [17] text encoder. Imagen [26] focuses on language understanding by using a large T5 language model to better represent semantics. The latent diffusion model [22] improves computational efficiency by performing the diffusion process in low-dimension latent space with an autoencoder. Finally, Stable Diffusion (SD) [22] employs a cross-attention mechanism to inject textual conditions into the diffusion generation process, aligning with the provided textual input. However, it is difficult to provide precise action descriptions in text, since user intent and machine understanding are not aligned. Furthermore, experimental results in Fig. 4 show that some actions are difficult to generate correctly without re-training, e.g., “performs a handstand”.

Controllable Action Generation. The paper focuses on transferring the desired action from exemplar images to unseen people, characters, and even animals for photorealistic image generation. Existing efforts take source images and pose information (e.g., skeletal images or body parsing) as conditions to control the generation. Previous controllable

solutions based on GANs [12, 13, 36] and VAEs [21, 33] suffer from training difficulties and poor generation results. Some subsequent works [24, 32] introduce text conditions to guide the action generation, yet fail with open vocabulary due to the small size of the vocabulary pools. Thanks to the significant advances of T2I diffusion models, recent methods [10, 14, 29], in particular the popular ControlNet [35], add arbitrary conditions to improve the versatility and controllability. While gaining a tremendous amount of traction from the community, ControlNet refers to the provided skeleton image to generate the action, which reduces flexibility and diversity. In addition, the objective of designing a general framework with additional trainable modules makes it not well-targeted to animals. In this work, we investigate customization solutions for action generation.

Subject-Driven Customization. Due to the demand for generating images with user-specified subjects, customization methods [5, 9, 25, 30] tailored to the appearance have been studied in the context of T2I generation. Specifically, DreamBooth [25] binds rare new words with specific subjects through fine-tuning the whole T2I generator. Textual Inversion [5] learns an extra identifier to represent the subject and adds the identifier as a new word to the dictionary of the text encoder. Custom Diffusion [9] only fine-tunes the key and value matrices of the cross-attention to represent new concepts. P+ [30] extends the textual-conditioning space with per-layer tokens to allow for greater disentangling and control. Despite the success achieved, the experimental results in Fig. 2 show their failure in action customization. A recent work ReVersion [7] makes progress in learning specific relations including some interactions from exemplar images. However, the design of the method, which specializes in learning spatial relations, makes it difficult to invert action information.

3. Action Customization Benchmark

Given a set of exemplar images $\mathcal{X} = \{\mathbf{x}_1, \mathbf{x}_2, \dots, \mathbf{x}_N\}$, we assume that all images contain the same action performed by different people. The action-agnostic descriptions associated to the exemplar images are also provided, which can be used as prompt templates during training. The objective of the action customization task is to extract the co-existing action and transfer it to the synthesis of action-specific images with different new subjects. In order to provide suitable conditions for systematic comparisons on this task, we present a new **ActionBench**, which consists of diverse actions accompanied by meticulously selected sample images. The benchmark can be used for both quantitative and qualitative comparisons.

Action Categories. To determine the involved actions, we first request GPT-4 [16] to provide 50 candidate action categories, and then attempt to collect images for these candidates. Only actions that can collect sufficient high-quality

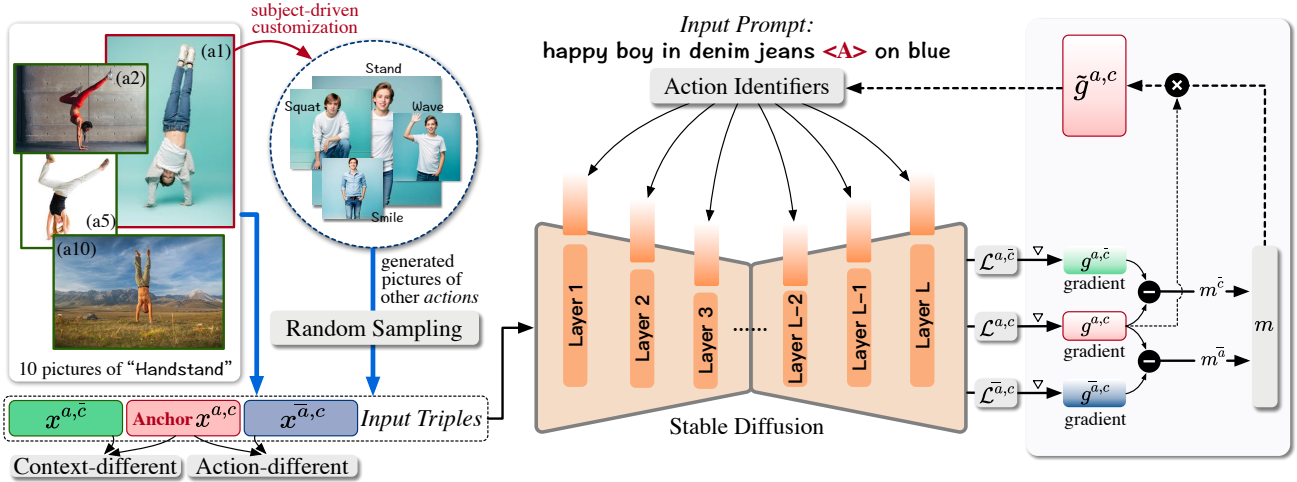


Figure 3. **Overview of our ADI method.** ADI learns more efficient action identifiers by extending the semantic conditioning space and masking gradient updates to action-agnostic channels.

images are preserved. We finally define eight unique actions, ranging from single-handed (e.g., “raises one finger”) to full-body movements (e.g., “performs a handstand”).

Exemplar Images and Prompts. For each action, we collect ten example images with corresponding textual descriptions, featuring different people. We manually remove action-related descriptions from the textual content to make them suitable as prompt templates.

Evaluation Subjects. We provide a list containing 23 subjects, including generic humans (e.g., “An old man”), well-known personalities (e.g., “David Beckham”), and animals (e.g., “A panda”). The latter two are guaranteed to be completely unseen, which tests the generalization of methods.

4. Methodology

We start with the technical background in Sec. 4.1. Then, we provide a comprehensive description of our proposed ADI in Sec. 4.2.

4.1. Preliminaries

Our study is based on the Stable Diffusion (SD) [22] model, which is considered to be the public state-of-the-art text-to-image generator. Specifically, to operate the diffusion process [6] in a low-dimensional latent space, SD employs a hierarchical VAE that consists of an encoder \mathcal{E} and a decoder \mathcal{D} . The encoder \mathcal{E} is tasked with encoding the given image \mathbf{x} into latent features \mathbf{z} , and the decoder \mathcal{D} reconstructs the image $\hat{\mathbf{x}}$ from the latent, i.e., $\hat{\mathbf{x}} = \mathcal{D}(\mathbf{z}) = \mathcal{D}(\mathcal{E}(\mathbf{x}))$. To control the generation with the textual conditions, given the noisy latent \mathbf{z}_t , current time step t and text tokens \mathbf{y} , a conditional U-Net [23] denoiser is trained to predict the noise ϵ added to the latent \mathbf{z} :

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

where \mathbf{y} is obtained by feeding the prompt into a CLIP [17] text encoder. During inference, the pre-trained SD first samples a latent \mathbf{z}_T from the standard normal distribution $\mathcal{N}(0, 1)$. Iteratively, \mathbf{z}_{t-1} can be obtained by removing noise from \mathbf{z}_t conditioned on \mathbf{y} . After the final denoising step, the latent \mathbf{z}_0 is mapped to generate an image $\hat{\mathbf{x}}$ with the decoder \mathcal{D} .

4.2. Action-Disentangled Identifier (ADI)

Given exemplar images that all contain a specific entity, existing subject-driven inversion methods [5, 9] learn to represent the entity as an identifier token $v \in \mathbb{R}^d$. And the learned v can then be employed in text prompts to produce diverse and novel images, where the entity can be generated with different contexts. In this paper, we continue the vein of capturing the common action in exemplar images by finding the optimal identifiers. An overview of our proposed ADI is illustrated in Fig. 3.

Expanding Semantic Inversion. To overcome the preference to low-level appearance features, we apply layer-wise identifier tokens to increase the accommodation of various features. Specifically, for the l -th layer where $l \in [1, L]$ and L is the number of cross-attention layers in the T2I model, a new identifier token $v_l \in \mathbb{R}^d$ is initialized. Feeding the prompt with v_l into the text encoder, the output tokens \mathbf{y}_l control the update of the latents in the l -th layer, thus influencing the generation of the visual content. And the learned tokens from all layers can form a token set \mathcal{V} , which can then be paired with different subjects for generation. Rather than having a single identifier token take on the responsibility of reconstruction, having separate identifiers at different layers effectively ensures that more features are converted, including the action-related features we care about.

Learning Gradient Mask with Context-Different Pair. The next step is to prevent the identifiers from inverting fea-

tures that are not relevant to the action and thus contaminating the subsequent image generation. Given $\mathbf{x}^{(a,c)} \in \mathcal{X}$ as an *anchor* sample, where a denotes the specific action, and c denotes the action-agnostic context contained in the image including human appearance and background, we can randomly sample another image $\mathbf{x}^{(a,\bar{c})}$ from \mathcal{X} , where \bar{c} represents that the context is different from c . Taking the context-different pair $\mathbf{x}^{(a,c)}$ and $\mathbf{x}^{(a,\bar{c})}$ as the input, we can calculate two gradients of the denoising loss \mathcal{L} with respect to the identifier token v :

$$g^{(a,c)} = \frac{\partial \mathcal{L}^{(a,c)}}{\partial v}, \quad (2)$$

$$g^{(a,\bar{c})} = \frac{\partial \mathcal{L}^{(a,\bar{c})}}{\partial v}. \quad (3)$$

Note that the subscript l is omitted for the sake of uniformity and clarity. Each identifier token contains multiple channels, each carrying semantically distinct and independent information. And the gradient consistency of a channel indicates that the channel is likely to carry information about the specific action. Therefore, we calculate the absolute value of the difference between the two gradients:

$$\Delta g^{\bar{c}} = |g^{(a,c)} - g^{(a,\bar{c})}|, \quad (4)$$

where the semantic channels with a small difference can be regarded as **action-related** channels of the action a , which are expected to be **preserved**. Specifically, we sort the difference from the largest to the smallest, taking the value at β percent γ^β as a threshold. In other words, $\beta\%$ of the channels are masked. Then, the mask that shares the same dimension as v can be calculated. For the k -th channel,

$$m_k^{\bar{c}} = \begin{cases} 0, & \Delta g_k^{\bar{c}} \geq \gamma^\beta \\ 1, & \Delta g_k^{\bar{c}} < \gamma^\beta \end{cases}. \quad (5)$$

By overwriting the mask to the gradient of the anchor sample, the action-related knowledge is preserved and incorporated into the update of v , while the updates on action-agnostic channels are ignored. Note that since the specific visual invariance about the action changes slightly depending on the sample pair, the masked channels may not be exactly the same each time. Furthermore, both samples use the prompt of the anchor sample $\mathbf{x}^{(a,c)}$ when calculating the gradients. Since the visual context of $\mathbf{x}^{(a,\bar{c})}$ is inconsistent with the description in the prompt, the reconstruction loss favours larger gradients in the context-related channels. In this way, the action-related channels found through the threshold will be more accurate.

Learning Gradient Mask with Action-Different Pair. Although the context-different pairs have the same action semantics, there may be differences in the visualization of the actions, and therefore the channels associated with the most representative action features do not necessarily have a smaller gradient difference. Since learning the gradient

mask with the context-different only pair is not stable and effective enough, we also construct action-different pairs to generate the gradient mask from another perspective. For each sample $\mathbf{x}^{(a,c)} \in \mathcal{X}$, we can use it to quickly train a subject-driven customization model (e.g., DreamBooth) that effectively inverts the most of the low-level context information. Therefore, by filling the prompt template of $\mathbf{x}^{(a,c)}$ with action descriptions that are different from a , the trained customization model can generate various action images as $\mathcal{X}^{(\bar{a},c)}$. Note that this step is not necessary if the users can compile a dataset of varied actions by the same individual using pre-captured images. However, the data collection is usually arduous and lengthy, making fast training of a subject-driven customization a more convenient solution. Due to the one-shot training and the concise text, the generated images $\mathcal{X}^{(\bar{a},c)}$ may not be consistent with the action descriptions, or the context may differ from the original $\mathbf{x}^{(a,c)}$, but in practice, we have found that the quality is sufficient to diversify the action variation. In this way, when $\mathbf{x}^{(a,c)}$ is sampled during training, we can randomly sample a image $\mathbf{x}^{(\bar{a},c)}$ from $\mathcal{X}^{(\bar{a},c)}$ to construct the action-different pair. And the gradient of $\mathbf{x}^{(\bar{a},c)}$ with respect to the token v can be calculated as

$$g^{(\bar{a},c)} = \frac{\partial \mathcal{L}^{(\bar{a},c)}}{\partial v}. \quad (6)$$

Similarly, both samples use the prompt of the anchor sample $\mathbf{x}^{(a,c)}$. We can also calculate the absolute value of the gradient difference:

$$\Delta g^{\bar{a}} = |g^{(a,c)} - g^{(\bar{a},c)}|, \quad (7)$$

where the semantic channels with small difference can be regarded as **context-related** channels of the action a , which are expected to be **masked**. Therefore, we have

$$m_k^{\bar{a}} = \begin{cases} 0, & \Delta g_k^{\bar{a}} < \lambda^\beta \\ 1, & \Delta g_k^{\bar{a}} \geq \lambda^\beta \end{cases}, \quad (8)$$

where λ^β is the threshold here to mask $\beta\%$ of the channels.

Merging Gradient Masks for Context. Due to the noise introduced by context variations, identifying action-relevant channels using only context-different or action-different pairs would be difficult and unreliable. As an evidence, the average overlap rate of channels preserved by both masks at each training step is 30.26%. Therefore, given the input triple $\mathcal{I} = \{\mathbf{x}^{(a,\bar{c})}, \mathbf{x}^{(a,c)}, \mathbf{x}^{(\bar{a},c)}\}$, we can merge the two obtained masks $m^{\bar{a}}$ and $m^{\bar{c}}$ to get the final context mask m . In practice, we keep only the intersection of the unmasked channels as unmasked, as we find this merging strategy performs better. Formally, we have

$$m = m^{\bar{c}} \cap m^{\bar{a}}, \quad (9)$$

which is overwritten to the gradient of the anchor sample:

$$\tilde{g}^{(a,c)} = m \odot g^{(a,c)}. \quad (10)$$




Figure 4. **Visual comparisons of all methods.** For each action, we present the generated results showcasing its pairing with a human character and an animal.

Note that the masked gradient $\tilde{g}^{(a,c)}$, where action-agnostic channels are considered to be masked, is the only gradient used to update v . Therefore, our identifiers can adequately invert action-related features.

5. Experiments

5.1. Experiment Setup

Baselines. For the baselines included in the comparison, we select Stable Diffusion [22], ControlNet [35], DreamBooth [25], Textual Inversion [5], ReVersion [7], Custom Diffusion [9] and P+ [30].

Implementation Details. For ADI, we set the masking ratio β to 0.6 and use the AdamW [11] optimizer with a learning rate of $2e-4$, while the training takes 3000 steps. For efficiency, DreamBooth for action-different pairs does not generate class-preservation images. While only one image is used for training, the initial learning rate is $1e-6$, and the training takes 2000 steps. For a fair comparison, we use 50 steps of the DDIM [27] sampler with a scale of 7.5 for all methods. Unless otherwise specified, Stable Diffusion

v2-1-base is selected as the default pre-trained model, and images are generated at a resolution of 512×512 . All experiments are conducted on a NVIDIA A100 GPU.

5.2. Quantitative Comparison

We perform the quantitative comparison with human evaluators to assess the quantitative performance. For each subject-action pair, four images are randomly sampled from the images generated by different methods. Given (1) the exemplar images of a specific action, and (2) the textual name of the subjects, human evaluators are asked to determine whether (1) the generated action is consistent with those in the exemplar images, and (2) the generated character corresponds with the name without obvious deformations, defects, or abnormalities. A generated image will only be considered totally correct if both the action and the character are correctly generated.

Tab. 1 reports the action, subject and total accuracy for all methods. Some observations are worth highlighting: (1) Given the textual descriptions of the actions, Stable Diffusion yields the highest total accuracy of all baseline meth-



Figure 5. **Ablation study.** We remove or revise one implementation at a time to demonstrate the effects of the identifier extension and the gradient masking.

Table 1. **Quantitative comparisons with competing methods.** Action, subject and total accuracies (%) are reported.

Methods	Action	Subject	Total
Stable Diffusion [22]	30.71	84.51	27.17
ControlNet [35]	41.30	42.66	19.29
DreamBooth [25]	2.45	95.65	2.45
Textual Inversion [5]	2.17	86.14	1.90
ReVersion [7]	1.63	84.51	1.63
Custom Diffusion [9]	29.62	53.53	7.07
P+ [30]	26.90	80.16	20.92
ADI (Ours)	60.33	85.87	51.09

ods. This suggests that the existing baselines do not take full advantage of the exemplar images. (2) Despite relying on the skeleton as the condition to improve the action generation, ControlNet fail to maintain the performance of subject generation, resulting in an unsatisfactory total accuracy. (3) The action accuracy of DreamBooth, Textual Inversion, and ReVersion is incredibly low, reflecting their complete failure to invert the action-related features. (4) Custom Diffusion and P+ improve action accuracy at more or less the expense of subject accuracy. (5) Attribute to the extended semantic conditioning space and the gradient masking strategy, our ADI dramatically improves the accuracy of action generation while maintaining excellent subject accuracy. As a result, ADI achieves the best total accuracy, outperforming the baselines by 23.92%.

5.3. Qualitative Comparison

Fig. 4 illustrates the qualitative comparison of all methods involved. It can be observed that although text descriptions of the actions are provided, the actions generated by Stable Diffusion still differ from the examples. ControlNet can only maintain a rough consistency in posture and struggles to match the generated subjects to the desired requirements, resulting in incomplete or distorted body structures, while

sacrificing diversity. And the subject-driven customization methods, as discussed earlier, fail to generate the actions or exhibit appearance characteristics that differ from the specified subjects. This suggests that they are unable to convert only the features associated with the actions. Giving the credit to the design from a perspective of gradient, our ADI decouples action-related features from action-agnostic information and blocks the inversion of the latter. This allows ADI to effectively model the invariance of the action and transfer it to different characters and animals without sacrificing image quality and variety.

5.4. Ablation Study

We conduct ablation experiments on ActionBench to verify the individual effects of the proposed contributions. From the generation results in Fig. 5, it can be observed that (1) The removal of the extension to the semantic conditioning space diminishes the inversion ability of ADI. (2) Both the gradient masks learned from the context-different and the action-different pairs are essential. Removing either one can lead to inadequate learning of action knowledge or a degradation in the quality of the subject’s appearance. We attribute this to the fact that learning from a single pair is inherently noisy due to varied action visuals and the interference of action-irrelevant information. (3) We also attempt to reverse the gradient masks, *i.e.*, updates to channels that should have been masked are preserved, and updates to other channels are cancelled. Obviously, this will result in action-related features not being inverted.

5.5. Further Analysis

Impact of Masking Strategy. To validate the masking strategy in our ADI, we compare it with four other strategies in Fig. 6. Specifically, on the gradients for each update: (1) Uniform: we uniformly mask β percent of channels. (2)



Figure 6. **Visual comparison of masking strategies.** The four compared strategies fail to mask the updates from agnostic channels and invert the action-related features.



Figure 7. **Effect of gradient mask merging strategy.** Preserving the intersection of gradient masks can better invert the representative action features.

Random: we randomly mask β percent of channels. (3) Min: we mask β percent of channels with the lowest value. (4) Max: we mask β percent of channels with the highest value. We observe that none of these four strategies successfully captures high-level features related to actions, since the images they generate are independent of the specified action. And the comparison also shows that the effectiveness of our ADI not only depends on the masking itself, but also requires learning action-agnostic channels by modeling the invariance of action and context.

Impact of Gradient Mask Merging Strategy. As shown in Eq. (9), ADI takes the intersection of the two gradient masks as the default merging strategy. We compare this with selecting the union of the two masks, and illustrate the generation results in Fig. 7. Since only channels that are preserved on both masks are updated, taking the intersection can effectively filter out action-agnostic features, leading to better customization of the actions. In contrast, taking the intersection may dilute the most representative action features due to the preserved context information.

Impact of Masking Ratio β . In Fig. 8, we vary the masking ratio β from 0.2 to 0.8. When β is small, fewer dimensions of the gradient are masked, and more action-agnostic features are retained to hinder the generation of the subject’s appearance. This situation improves as β is gradually increased. However, when β is relatively large, due to the large number of masked dimensions, some of the most dis-

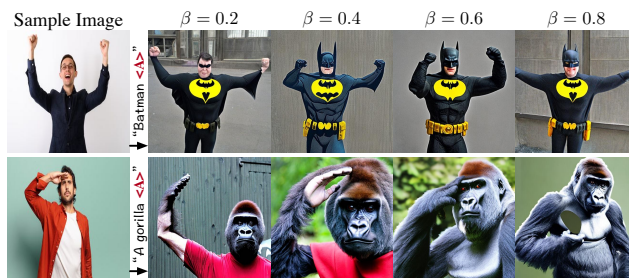


Figure 8. **Effect of the masking ratio β .** A value close to 0.5 can balance the inversion of action-related features and the removal of action-agnostic features.

criminative features of actions may not be inverted, resulting in incomplete learning of actions. Note that the optimal value of β may be different for different actions.

6. Conclusion

In this paper, we investigate an under-explored text-to-image generation task, namely action customization. To understand the challenge of the task, we first visualize the inadequacy of existing subject-driven methods in extracting action-related features from the entanglement of action-agnostic context features. Then, we propose a novel method named ADI to learn action-specific identifiers from the given images. To increase the accommodation of knowledge relevant to the action, ADI extends the inversion process with layer-wise identifier tokens. Furthermore, ADI generates gradient masks to block the contamination of action-agnostic features at the gradient level. We also contribute the ActionBench for evaluating performance on the task. Since there is a growing need to synthesize action-specific images with various new subjects, we hope that our work can highlight this important direction.

Acknowledgement This work was supported by STI 2030—Major Projects (2022ZD0208800), NSFC General Program (Grant No. 62176215). This work was supported by Alibaba Group through Alibaba Research Intern Program.

References

- [1] Zhe Cao, Gines Hidalgo, Tomas Simon, Shih-En Wei, and Yaser Sheikh. OpenPose: Realtime multi-person 2d pose estimation using part affinity fields. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, pages 172–186, 2021. [1](#)
- [2] Prafulla Dhariwal and Alexander Quinn Nichol. Diffusion models beat GANs on image synthesis. In *Proceedings of the Advances in Neural Information Processing Systems*, pages 8780–8794, 2021. [3](#)
- [3] Ming Ding, Zhuoyi Yang, Wenyi Hong, Wendi Zheng, Chang Zhou, Da Yin, Junyang Lin, Xu Zou, Zhou Shao, Hongxia Yang, and Jie Tang. CogView: Mastering text-to-image generation via transformers. In *Proceedings of the Advances in Neural Information Processing Systems*, pages 19822–19835, 2021. [2](#), [3](#)
- [4] Oran Gafni, Adam Polyak, Oron Ashual, Shelly Sheynin, Devi Parikh, and Yaniv Taigman. Make-A-Scene: Scene-based text-to-image generation with human priors. In *Proceedings of the European Conference on Computer Vision*, pages 89–106, 2022. [3](#)
- [5] Rinon Gal, Yuval Alaluf, Yuval Atzmon, Or Patashnik, Amit Haim Bermano, Gal Chechik, and Daniel Cohen-Or. An image is worth one word: Personalizing text-to-image generation using textual inversion. In *Proceedings of the International Conference on Learning Representations*, 2023. [2](#), [3](#), [4](#), [6](#), [7](#), [1](#)
- [6] Jonathan Ho, Ajay Jain, and Pieter Abbeel. Denoising diffusion probabilistic models. In *Proceedings of the Advances in Neural Information Processing Systems*, 2020. [4](#)
- [7] Ziqi Huang, Tianxing Wu, Yuming Jiang, Kelvin C. K. Chan, and Ziwei Liu. ReVersion: Diffusion-based relation inversion from images. *arXiv preprint arXiv:2303.13495*, 2023. [2](#), [3](#), [6](#), [7](#), [1](#)
- [8] Diederik P. Kingma and Max Welling. Auto-encoding variational bayes. In *Proceedings of the International Conference on Learning Representations*, 2014. [3](#)
- [9] Nupur Kumari, Bingliang Zhang, Richard Zhang, Eli Shechtman, and Jun-Yan Zhu. Multi-concept customization of text-to-image diffusion. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 1931–1941, 2023. [2](#), [3](#), [4](#), [6](#), [7](#), [1](#)
- [10] Yuheng Li, Haotian Liu, Qingyang Wu, Fangzhou Mu, Jianwei Yang, Jianfeng Gao, Chunyuan Li, and Yong Jae Lee. GLIGEN: Open-set grounded text-to-image generation. *arXiv preprint arXiv:2301.07093*, 2023. [2](#), [3](#)
- [11] Ilya Loshchilov and Frank Hutter. SGDR: Stochastic gradient descent with warm restarts. In *Proceedings of the International Conference on Learning Representations*, 2017. [6](#)
- [12] Liqian Ma, Xu Jia, Qianru Sun, Bernt Schiele, Tinne Tuytelaars, and Luc Van Gool. Pose guided person image generation. In *Proceedings of the Advances in Neural Information Processing Systems*, pages 406–416, 2017. [3](#)
- [13] Yifang Men, Yiming Mao, Yuning Jiang, Wei-Ying Ma, and Zhouhui Lian. Controllable person image synthesis with attribute-decomposed GAN. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 5083–5092, 2020. [3](#)
- [14] Chong Mou, Xintao Wang, Liangbin Xie, Jian Zhang, Zhong-gang Qi, Ying Shan, and Xiaoju Qie. T2I-Adapter: Learning adapters to dig out more controllable ability for text-to-image diffusion models. *arXiv preprint arXiv:2302.08453*, 2023. [2](#), [3](#)
- [15] Alexander Quinn Nichol, Prafulla Dhariwal, Aditya Ramesh, Pranav Shyam, Pamela Mishkin, Bob McGrew, Ilya Sutskever, and Mark Chen. GLIDE: Towards photorealistic image generation and editing with text-guided diffusion models. In *Proceedings of the International Conference on Machine Learning*, pages 16784–16804, 2022. [3](#)
- [16] OpenAI. GPT-4 technical report. *arXiv preprint arXiv:2303.08774*, 2023. [3](#)
- [17] Alec Radford, Jong Wook Kim, Chris Hallacy, Aditya Ramesh, Gabriel Goh, Sandhini Agarwal, Girish Sastry, Amanda Askell, Pamela Mishkin, Jack Clark, Gretchen Krueger, and Ilya Sutskever. Learning transferable visual models from natural language supervision. In *Proceedings of the International Conference on Machine Learning*, pages 8748–8763, 2021. [3](#), [4](#)
- [18] Aditya Ramesh, Mikhail Pavlov, Gabriel Goh, Scott Gray, Chelsea Voss, Alec Radford, Mark Chen, and Ilya Sutskever. Zero-shot text-to-image generation. In *Proceedings of the International Conference on Machine Learning*, pages 8821–8831, 2021. [2](#), [3](#)
- [19] Aditya Ramesh, Prafulla Dhariwal, Alex Nichol, Casey Chu, and Mark Chen. Hierarchical text-conditional image generation with CLIP latents. *arXiv preprint arXiv:2204.06125*, 2022. [3](#)
- [20] Scott E. Reed, Zeynep Akata, Xinchun Yan, Lajanugen Logeswaran, Bernt Schiele, and Honglak Lee. Generative adversarial text to image synthesis. In *Proceedings of the International Conference on Machine Learning*, pages 1060–1069, 2016. [2](#), [3](#)
- [21] Yurui Ren, Xiaoming Yu, Junming Chen, Thomas H. Li, and Ge Li. Deep image spatial transformation for person image generation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 7687–7696, 2020. [3](#)
- [22] 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 10674–10685, 2022. [2](#), [3](#), [4](#), [6](#), [7](#)
- [23] Olaf Ronneberger, Philipp Fischer, and Thomas Brox. U-Net: Convolutional networks for biomedical image segmentation. In *Proceedings of the International Conference on Medical Image Computing and Computer Assisted Intervention Society*, pages 234–241, 2015. [4](#)
- [24] Prasun Roy, Subhankar Ghosh, Saumik Bhattacharya, Uma-pada Pal, and Michael Blumenstein. TIPS: text-induced pose synthesis. In *Proceedings of the European Conference on Computer Vision*, pages 161–178, 2022. [3](#)
- [25] Nataniel Ruiz, Yuanzhen Li, Varun Jampani, Yael Pritch, Michael Rubinstein, and Kfir Aberman. 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*, pages 22500–22510, 2023. [2](#), [3](#), [6](#), [7](#), [1](#)
- [26] Chitwan Saharia, William Chan, Saurabh Saxena, Lala Li, Jay Whang, Emily L. Denton, Seyed Kamyar Seyed Ghasemipour, Raphael Gontijo Lopes, Burcu Karagol Ayan, Tim Salimans, Jonathan Ho, David J. Fleet, and Mohammad Norouzi. Photorealistic text-to-image diffusion models with deep language understanding. In *Proceedings of the Advances in Neural Information Processing Systems*, pages 36479–36494, 2022. [2](#), [3](#)
- [27] Jiaming Song, Chenlin Meng, and Stefano Ermon. Denoising diffusion implicit models. In *Proceedings of the International Conference on Learning Representations*, 2021. [6](#)
- [28] Ming Tao, Hao Tang, Fei Wu, Xiaoyuan Jing, Bing-Kun Bao, and Changsheng Xu. DF-GAN: A simple and effective baseline for text-to-image synthesis. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 16494–16504, 2022. [3](#)
- [29] 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*, pages 1921–1930, 2023. [3](#)
- [30] Andrey Voynov, Qinghao Chu, Daniel Cohen-Or, and Kfir Aberman. P+: Extended textual conditioning in text-to-image generation. *arXiv preprint arXiv:2303.09522*, 2023. [2](#), [3](#), [6](#), [7](#), [1](#)
- [31] Tao Xu, Pengchuan Zhang, Qiuyuan Huang, Han Zhang, Zhe Gan, Xiaolei Huang, and Xiaodong He. AttnGAN: Fine-grained text to image generation with attentional generative adversarial networks. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 1316–1324, 2018. [3](#)
- [32] Xiaogang Xu, Ying-Cong Chen, Xin Tao, and Jiaya Jia. Text-guided human image manipulation via image-text shared space. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, pages 6486–6500, 2022. [3](#)
- [33] Lingbo Yang, Pan Wang, Chang Liu, Zhanning Gao, Peiran Ren, Xinfeng Zhang, Shanshe Wang, Siwei Ma, Xian-Sheng Hua, and Wen Gao. Towards fine-grained human pose transfer with detail replenishing network. *IEEE Transactions on Image Processing*, pages 2422–2435, 2021. [3](#)
- [34] Jiahui Yu, Yuanzhong Xu, Jing Yu Koh, Thang Luong, Gunjian Baid, Zirui Wang, Vijay Vasudevan, Alexander Ku, Yinfei Yang, Burcu Karagol Ayan, Ben Hutchinson, Wei Han, Zarana Parekh, Xin Li, Han Zhang, Jason Baldridge, and Yonghui Wu. Scaling autoregressive models for content-rich text-to-image generation. *Transactions on Machine Learning Research*, 2022. [3](#)
- [35] Lvmin Zhang and Maneesh Agrawala. Adding conditional control to text-to-image diffusion models. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 2023. [2](#), [3](#), [6](#), [7](#), [1](#)
- [36] Pengze Zhang, Lingxiao Yang, Jianhuang Lai, and Xiaohua Xie. Exploring dual-task correlation for pose guided person image generation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 7703–7712, 2022. [3](#)
- [37] Shengjia Zhao, Jiaming Song, and Stefano Ermon. InfoVAE: Balancing learning and inference in variational autoencoders. In *Proceedings of the AAAI Conference on Artificial Intelligence*, pages 5885–5892, 2019. [3](#)
- [38] Minfeng Zhu, Pingbo Pan, Wei Chen, and Yi Yang. DM-GAN: Dynamic memory generative adversarial networks for text-to-image synthesis. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 5802–5810, 2019. [3](#)