

Attribute Token Arithmetic: Disentangled and Continuous Semantic Control for Visual Autoregressive Models

Xindi Yang , Yicheng Wu ^{*}, Cheng Zhang ,
Jianfei Cai , and Tien-Tsin Wong 

Faculty of IT, Monash University
{xindi.yang, yicheng.wu1, cheng.zhang, jianfei.cai, tt.wong}@monash.edu

Abstract. Autoregressive text-to-image generation has recently achieved remarkable progress, offering high-fidelity synthesis via a unified generative framework. However, fine-grained semantic control remains challenging due to the attribute entanglement and the misalignment between textual and fine-grained visual representations. In this paper, we introduce Attribute Token Arithmetic (ATA), a method that enables disentangled and continuous attribute control in visual autoregressive modelling. Inspired by the vector arithmetic property observed in word embeddings, ATA identifies semantic directions corresponding to visual attributes (*e.g.*, aging, fatness, emotion) directly within the pretrained autoregressive latent space. These directions are learned from a single reference image, without model retraining or large-scale supervision. During generation, attributes can be continuously adjusted and compositionally combined through simple arithmetic operations with other attribute tokens. Extensive experiments demonstrate that ATA achieves identity-preserving, fine-grained, and multi-attribute adjustment, outperforming existing autoregressive editing baselines in controllability, generality, and computational efficiency. Our code will be released here.

Keywords: Autoregressive Modeling · Arithmetic · Image Editing

1 Introduction

Photorealistic image synthesis and editing are fundamental tasks in visual content creation, underpinning applications in movie production, gaming, and virtual reality. Traditional computer graphics achieve these goals through geometric modeling and physics-based rendering, but such pipelines are labor-intensive and require substantial expertise. Recent years have witnessed a paradigm shift toward generative AI, which produces realistic images directly from data, significantly lowering the barrier to visual content creation [30, 40, 46–49, 54].

Among generative paradigms, diffusion models [30, 32, 33, 45] have achieved remarkable success in text-to-image generation, enabling high-fidelity and semantically aligned visual synthesis. For example, TokenVerse [11] and XVerse [4]

^{*} Corresponding author




Fig. 1: Exemplar edited images with disentangled and continuous semantic control. Given a reference and paired text prompts, our ATA extracts semantic attribute directions in the VAR-aligned text space. These directions are derived from the image concepts that correspond to the semantic differences between paired prompts. Specifically, ATA modifies source images to gradually increase the strength of two attributes. The first two results increase with the first attribute, while in the last two, the second attribute is introduced. (Left) ATA enables disentangled and continuous semantic control of “**angry**” and “**old**”, while preserving the identity. (Right) Moreover, ATA can transfer the general *disentangled* attribute concepts across object categories. Although “**fat**”-ness is not a native attribute of chairs, ATA can still transfer the “**fat**” & “**furry**” concepts and enable continuous control to yield creative chair designs, while both VAREdit and Qwen-Image-Edit fail to provide continuous control.

attempt to achieve concept-level control by associating text tokens with visual semantics, allowing multi-concept manipulation. However, diffusion models require iterative denoising steps, which are computationally expensive and limit their efficiency for real-time or high-resolution editing [20, 32]. This limitation has renewed interest in autoregressive (AR) models, which generate images token by token in a discrete latent space, offering faster sampling, superior scalability, and a unified formulation with language generation [5, 34].

Building upon this idea, Visual Autoregressive Models (VAR) [13, 37] extend AR generation to high-resolution imagery through multi-scale tokenization and causal transformers. By formulating image synthesis as a sequence prediction task, VAR naturally bridges language and vision within a discrete latent space. Despite their success in image quality and efficiency, precise and flexible editing remains a major challenge, as fine-grained semantic control remains challenging to realize within discrete token spaces. Such fine-grained control is essential for many graphics applications to allow users to fine-tune the desired visual effect. Recent works attempt to address this limitation. AREdit [41] manipulates internal representations of pretrained VARs without retraining but supports only single-attribute edits. VAREdit [23] fine-tunes VARs on paired examples for instruction-based editing but suffers from limited generalization. Overall, existing approaches struggle to achieve generic, continuous, and compositional control within the VAR framework.

Therefore, our work starts from a simple yet powerful question: *Can attribute-level control be achieved within VARs similar to text token arithmetic?* We discover that semantically meaningful vector operations exist in the VAR-aligned text-token space, where manipulating token vectors associated with specific at-



Fig. 2: Preliminary image editing results. By manipulating semantic directions to the corresponding text tokens, (a) we achieve the localized editing and the preservation of object identities; (b) we can also edit non-localized attributes, *e.g.*, background.

tributes (*e.g.*, “**angry**”, “**old**”) produces coherent and interpretable visual changes (Fig. 1). This insight reveals that text tokens inherently encode compositional visual semantics, suggesting a new paradigm for continuous and cross-category control, without any retraining.

Motivated by this observation, we propose ATA (Attribute Token Arithmetic), a lightweight optimization-based framework that learns multi-attribute vectors directly within a pretrained VAR. Given a single reference image and its textual description, ATA disentangles attribute-specific directions in the joint text–token space and leverages them for controllable image synthesis. Once learned, these attribute vectors can be stored in an offline library and reused or combined at arbitrary strengths to create new attributes and visual compositions (see Fig. 1). Practically, ATA enables intuitive edits such as “**making a face angrier**”, “**adding a smile and aging simultaneously**”, or even “**transferring the fat attribute of a cat to a chair**”.

Our contributions are as follows:

- We reveal that semantically meaningful vector arithmetic exists in the latent space of pretrained VARs, enabling attribute-level visual control.
- We propose ATA, a simple yet effective framework for learning generic and compositional attribute vectors without retraining or large-scale supervision.
- We validate ATA through extensive experiments and user studies, demonstrating its controllability and generalizability superior to existing VAR editing approaches.

2 Related Work

2.1 Autoregressive Image Generation

The success of LLMs, such as the GPT series [1], has demonstrated the scalability and versatility of the autoregressive (AR) paradigm. To adapt this to visual data, seminal works like VQ-VAE [39] and VQGAN [9] introduce learnable codebooks

to represent image patches as discrete tokens. Subsequent variants have improved codebook utilization and reconstruction via multi-codebook designs [21], residual quantization [16], and regularization [50]. Additionally, leveraging semantic priors from models like CLIP [55] has enhanced the alignment between quantized tokens and textual concepts.

More recently, research has shifted toward scalar-based quantization, such as Binary Scalar Quantization (BSQ) [53], which offers a lightweight alternative to traditional vector quantization. To mitigate information loss, studies have explored continuous or hybrid formulations: MAR [17] employs a diffusion process in continuous latent space, TokenBridge [42] utilizes post-training quantization, and HART [35] adopts a hybrid design. Despite these continuous alternatives, discrete representations remain favored for their ability to naturally unify linguistic and visual modalities.

Once tokenized, causal transformers are trained to capture sequential dependencies. While early models like DALL-E [29] directly modeled flattened sequences, they faced scalability and spatial consistency challenges. Later strategies introduced hierarchical and spatially aware modeling: TiTok [51] uses register tokens for semantic regions, while column-wise tokenization [8] generates tokens from left to right. Visual Autoregressive Modeling (VAR) [37] further refined this via next-scale prediction across resolutions, a paradigm recently extended by Infinity [13] through the integration of scalar quantization.

Collectively, these advances have significantly bridged the gap between continuous visual data and discrete sequence modeling, establishing a solid foundation for autoregressive image generation with enhanced fidelity, coherence, and flexibility. However, achieving fine-grained control in AR-based generation remains an open challenge. Text conditions, while semantically expressive, often provide only coarse guidance that fails to capture intricate visual details. Consequently, the generated images may lack completeness, transferability, precision, and structural consistency across regions or attributes. These limitations underscore the need for explicit, disentangled, and semantically grounded control mechanisms that are highly demanded in the movie and gaming industries. These motivate our exploration of content control within the autoregressive framework.

2.2 Image Editing

Image editing is a long-standing problem in computer vision and graphics. Recent foundation models have advanced visual content generation. This progress has shifted attention to achieve fine-grained semantic control in generative model. Diffusion-based models have explored in this area. [2, 14, 27, 38] manipulate the denoising process via feature injection to control target concepts in image. [10, 31] propose to train LoRA [15] modules for specific concepts to improve disentanglement and consistency. [4, 7, 11, 12] propose edits by locating semantic directions in a pretrained latent space. [22, 43] leverages large-scale image-text paired data to finetune the base model and enable instruction-based editing. These methods can be effective in fine-grained semantic editing. However, iterative sampling in the denoising process slows inference and makes real-time editing difficult.

Table 1: Comparison among existing image editing methods.

Model	Strength Control	Cross-Category Transfer	Object Composition
Concept Slider	✓	✓	✗
Tokenverse	✗	✗	✓
VAREdit, Qwen-Image-Edit	✓	✓	✗
Ours	✓	✓	✓

Recently, AR image generation has gained interest due to its higher efficiency and language-like modeling paradigm. AREdit [41] manipulates representations but focuses on single-concept edits, which limits disentanglement. [23, 26, 44] are trained on a large-scale dataset and offer instruction-based interfaces, yet they provide limited control. Our ATA work aims to bridge this gap, providing disentangled and continuous semantic control in the VAR models.

Comparison with Existing Methods. Image editing has progressed significantly in both the Diffusion and AR communities, with multiple approaches emerging to address similar objectives. Table 1 compares our work with these recent image editing methods.

3 VAR Basics

VAR [13, 37] formulates text-to-image generation as a sequence prediction task in the discrete token space. Given a text prompt, a frozen language encoder (*i.e.*, Flan-T5 [6]) produces text tokens. These text tokens are then mapped into a VAR-aligned text token space via an MLP network to align with image tokens. An image tokenizer encodes the image I into a feature map F and then quantizes F into K multi-scale residual maps $(R_1, R_2, \dots, R_k)$. In our experiments, there are 13 scales ranging from 1×1 to 64×64 to synthesize 1024×1024 images.

We adopt Infinity [13] as our VAR backbone. As illustrated in Fig. 3(a), each transformer block consists of self-attention, cross-attention, and feed-forward layers, where the text tokens guide the generation of image tokens through cross-attention. During training, the model predicts next-scale residual $\hat{R}_{k+1}$ conditioned on all previous residual maps $(R_1, R_2, \dots, R_k)$ and VAR-aligned text context. The causal transformer takes low-resolution residual maps as input and predicts the next-resolution residual maps through a progressive up-interpolation process. During the inference, image generation proceeds hierarchically from coarse to fine scales. Coarse-scale tokens capture the global semantic structure, while fine-scale tokens incrementally refine spatial details until the full-resolution image is reconstructed. The fine-scale components are represented as residual maps, forming a Laplacian pyramid-like hierarchy that progressively injects local details into global semantics.




Fig. 3: Pipeline of our ATA framework. (a) A pretrained VAR model processes image and text tokens with a causal transformer. At each next-scale prediction step, VAR applies Self-Attention, Cross-Attention, and Feed-Forward blocks. Text and image tokens interact through cross-attention. Our ATA adds the learned semantic directions to the corresponding VAR-aligned text tokens in the Cross-Attention Block. (b) Given a reference image and a pair of semantically different prompts, our ATA learns two semantic directions, d_{general} and d_{specific} , for each target concept, via a lightweight attention block, Delta-Mod. These directions capture the general attribute concept and image-specific personalized traits. They are learned with the original VAR training loss. (c) During inference, the learned directions are applied to the corresponding text tokens. Our ATA enables disentangled and continuous control in VAR.

4 Method

Inspired by the vector arithmetic behavior observed in word embeddings [24], we first investigate in Sec. 4.1 whether various attributes (such as fatness, aging) can be edited via controlling the semantic directions in the pretrained VAR space. We found that there indeed exist controllable semantic directions, which enable localized semantic manipulation in the VAR-aligned text token space. To achieve more fine-grained control within this space for personalized image editing, we propose to learn disentangled semantic directions from individual images in Sec. 4.2.

4.1 Semantic Control in Text Tokens

Prior work in NLP [24] shows that word embeddings support vector arithmetic consistent with human intuition. For example, **king + woman - man $\approx$ queen**, where simple addition and subtraction on token embeddings produce the desired output aligned with human intuition. Inspired by the vector arithmetic behavior observed in word embeddings, we explore whether VAR supports a similar form of vector arithmetic for image editing.

In VAR’s text-to-image generation task, a frozen text encoder processes the prompt P to produce text tokens, which are then mapped by a pretrained MLP to the VAR-aligned text-embedding space. These tokens interact with image tokens in the VAR transformer and guide the generation process, as shown in Fig. 3 (a). For simplicity, consider an image of a fat cat. We can construct a prompt pair P_{src} : “A cat.” and P_{tar} : “A fat cat.” We can regard “A” and “A

fat” as the attribute tokens $T_{\text{attr}}^{\text{src}}$ and $T_{\text{attr}}^{\text{tar}}$, respectively, while regarding **cat**” as the identity token T_{id} in VAR-aligned text-token space. We treat the attribute difference between $T_{\text{attr}}^{\text{src}}$ and $T_{\text{attr}}^{\text{tar}}$ as the semantic direction:

$$\Delta T_{\text{attr}} = \text{mean}(T_{\text{attr}}^{\text{tar}}) - T_{\text{attr}}^{\text{src}}, \quad (1)$$

where $\text{mean}(\cdot)$ denotes the mean of the multiple text tokens associated with $T_{\text{attr}}^{\text{tar}}$. With ΔT_{attr} , we can manipulate the semantic direction in token space as follows:

$$\hat{T}_{\text{attr}}^{\text{src}} = T_{\text{attr}}^{\text{src}} + \gamma \times \Delta T_{\text{attr}}, \quad (2)$$

where γ is a hyperparameter that controls the strength of the semantic direction. By adding a specific semantic direction to the attribute token, we can edit the corresponding attribute in the image. As shown in Fig. 2, the edits are highly localized and do not affect unrelated regions. The effect is not limited to attributes and objects; non-object concepts (*e.g.*, background) can also be edited in the same way.

4.2 Disentangled Semantic Direction Learning

Using the text token embedding difference in Eq. (1) directly as the semantic direction is insufficient to represent the general and specific concepts embodied in an image. Here, we aim to learn semantic directions from individual images, which can later be used in fine-grained personalized image editing.

Disentangled Semantic Direction Representation

Consider a text-to-image generation process, which maps the text prompt P to the distribution $\mathcal{S}_P$ in the VAR-aligned text token space (corresponding to a set of images that satisfy the text prompt P). The model sampling process picks a sample, ϕ_I , from the distribution $\mathcal{S}_P$, where I is the specific image generated. To represent the semantic shift in the latent space, we establish a direction $\mathbf{d}$ that originates from the mean of the distribution $\mathcal{S}_P$ and points toward the target sample point




Fig. 4: Concept illustration of ATA. Given a reference image I and a pair of semantically different prompts (*e.g.*, “A fat cat.” and “A cat.”), our method learns two semantic directions with respect to I : Fat_{specific} , the semantic direction from the mean of the distribution “fat cat” to I , and Fat_{general} , the semantic direction from the mean of the distribution “cat” to I , both in the VAR-aligned text space. We obtain the general attribute direction $Fat_{\text{attribute}}$ by simple vector arithmetic. The learned attribute transfers across object categories (*e.g.* transfer “fat” to chair).

ϕ_I . This orientation is formulated as:

$$\mathbf{d}_P = \phi_I - \mu(\mathcal{S}_P), \quad (3)$$

where $\mu(\mathcal{S}_P)$ denotes the mean of distribution $\mathcal{S}_P$ in the same latent space.

Let us use Fig. 4 as an example to illustrate our key idea. Specifically, we use the attribute (*i.e.*, **fat**) to describe a subject (*i.e.*, **a cat**) in the text prompt, which effectively specifies a sub-distribution (*i.e.*, **a fat cat**) out of the original distribution (*i.e.*, **a cat**). For a specific sample ϕ_I within the sub-distribution, we can form two vectors: one from the mean of the original distribution towards ϕ_I , denoted as $\mathbf{d}_{\text{general}}$ (*i.e.*, $\text{Fat}_{\text{general}}$), and the other from the mean of the sub-distribution towards ϕ_I , denoted as $\mathbf{d}_{\text{specific}}$ (*i.e.*, $\text{Fat}_{\text{specific}}$). Note that $\mathbf{d}_{\text{general}}$ and $\mathbf{d}_{\text{specific}}$ can serve as semantic directions that correspond to the identity concepts in the image I . By subtracting these two vectors, we effectively disentangle the *attribute direction* from the specific image:

$$\mathbf{d}_{\text{attribute}} = \mathbf{d}_{\text{general}} - \mathbf{d}_{\text{specific}}, \quad (4)$$

where $\mathbf{d}_{\text{attribute}}$ (*i.e.*, $\text{Fat}_{\text{attribute}}$) identifies a generic attribute direction, from the original distribution towards the sub-distribution. Amazingly, such general attributes can generalize across different object categories (here, transferring **fat**, learned from **fat cat**, to **chair**).

Learning Disentangled Semantic Directions We now present how to learn the disentangled semantic directions $\mathbf{d}_{\text{general}}$ and $\mathbf{d}_{\text{specific}}$ from a given image. Specifically, we learn them via constructing semantically different prompt pairs. Given a reference image I , we define a general prompt P_{general} and a specific prompt P_{specific} (*e.g.*, “**A cat on the ground.**” and “**A fat cat on the ground.**”), and $S(P_{\text{specific}})$ is a sub-distribution of $S(P_{\text{general}})$. We learn the semantic directions via a lightweight attention block, **Delta-Mod**, as shown in Fig. 3(b) top. Particularly, we extract the predicted residual map $\{\hat{R}_1, \hat{R}_2, \dots, \hat{R}_k\}$ from the reference image I using the visual encoder in the VAR model, which is fed into the attention block Delta-Mod as K and V . Together with two learnable input vectors as Q , Delta-Mod outputs the corresponding semantic direction $\{\mathbf{d}_{\text{general}}, \mathbf{d}_{\text{specific}}\}$. The process can be formulated as

$$\{\mathbf{d}_{\text{general}}, \mathbf{d}_{\text{specific}}\} = \text{Delta-Mod}(R_1, \dots, R_n). \quad (5)$$

Note that we do not have ground truth for the semantic directions. To supervise them, as shown in Fig. 3(b) bottom, we add $\mathbf{d}_{\text{general}}$ and $\mathbf{d}_{\text{specific}}$ to the corresponding text token positions of P_{general} and P_{specific} , respectively, both of which are required to reconstruct the given image I via the VAR model. With the VAR model frozen and the same VAR reconstruction loss, we only optimize the lightweight attention block Delta-Mod. Multiple semantic directions corresponding to different concepts or attributes of an image can be learned together via the same Delta-Mod.




Fig. 5: General attribute transfer and cross-category semantic editing. We present more examples of general attribute transfer. (Left) We apply the attribute “**Apple-shaped**” to the concept “**cake**” and perform continuous control. Our method gradually transforms the cake into an apple-shaped cake, whereas baseline methods turn the cake into an apple. (Right) We transfer the attribute concept “**Sphynx**”, originally defined for Sphynx cats, to sheep. Our method achieves a gradual transition to hairlessness with pronounced skin folds, whereas the baselines fail to do so.

Figure 3(c) illustrates our inference process. With different disentangled semantic directions extracted from different images, we can combine them for fine-grained personalized image editing by selecting either specific semantic directions $\mathbf{d}_{\text{specific}}$ (e.g., $\text{Cat}_{\text{specific}}$ and $\text{Glasses}_{\text{specific}}$) or the general attribute direction Eq. (4) (e.g., $\text{Fat}_{\text{attribute}}$), and applying them on the corresponding text tokens as in Eq. (2). Note that due to the next-scale prediction nature in VAR, we have another hyperparameter that selects the scale at which we apply the semantic direction manipulation; see the appendix for details. In addition, we can construct a semantic attribute library offline, from which semantic directions can be retrieved and combined on demand.

5 Experiments

We evaluate our method against state-of-the-art baselines and demonstrate its advantages in continuous, disentangled, and generalizable semantic control. Our experiments are organized into five parts: (1) comparison with existing editing models, (2) user studies, (3) quantitative evaluation, (4) ablation studies and (5) additional qualitative demonstrations of advanced editing capabilities.

5.1 Benchmark

Comparisons. Instruction-driven editing models follow textual commands but lack explicit mechanisms for continuous semantic control. We approximate continuous editing by prompting these methods with three attribute strengths: *slight*, *moderate*, and *extreme*. We include: **Qwen-Image-Edit (20B)** [43], a large-scale DiT-based instruction editor, and **VAREdit (Infinity-8B)** [13], a VAR-based instruction-following editor. Furthermore, we compare with **Concept Sliders** [10], which learn continuous attribute directions by optimizing LoRA modules for each attribute concept. Here, we use their SDXL-trained [28] LoRAs in our experiments.

Data Details. To assess continuous multi-attribute manipulation, we construct a controlled benchmark consisting of:

- two object categories: humans and common objects,
- 15 prompts per category, each with two seeds,
- six attribute sets, spanning common and general concepts (*e.g.*, *apple-shaped*).

Each source image is edited with two attributes at three intensity levels. We then report VQA-Score [18] for semantic alignment and LPIPS (VGG) [52] for content preservation.

Moreover, we leverage the GEdit [19] benchmark for evaluations. We randomly select 52 different attribute-editing samples. Then, Gemini [36] is used to generate attribute prompt pairs in the form of “a [attribute] [object] [context]” vs. “a [object] [context]”, sharing the same context, along with the corresponding reference image. ATA is used to obtain general semantic attribute directions. In this way, we built an attribute library in GEdit. Beyond semantic consistency and perceptual quality in GEdit, we introduce a disentanglement score to monitor unintended changes in [attribute2] when [attribute1] is modified. For the evaluation of GEdit, we employ the Qwen2.5-VL-72B [3] model as the GEdit evaluation backbone.

For comparison, we construct prompts in the form “make [object] *slightly/moderately/extremely* [attribute1]” and “*slightly/moderately/extremely* [attribute2]” to evaluate multi-attribute and continuous control in VAREdit and Qwen-Image. Concept-slider follows the same evaluation protocol as ATA.

Implementation. All experiments were conducted in an identical environment for a fair comparison (Hardware: a single NVIDIA H100 GPU; Software: PyTorch 2.5.1, CUDA 12.2). We optimize Delta-Mod for 500 steps over 5 epochs with a batch size of 1 and use the original VAR objective to optimize Delta-Mod, *i.e.*, the cross-entropy loss.

5.2 Qualitative Evaluation

Figure 5 presents side-by-side comparisons with Qwen-Image-Edit, VAREdit and Concept-slider on two challenging attribute transfer tasks: applying the *apple-shaped* attribute to cakes and transferring the *Sphynx* concept from sphynx cats to sheep.

Table 2: Quantitative comparison of semantic alignment and content preservation. We use Δ VQA and Semantic Consistency (S.C.) to measure the semantic shift toward the target attribute. Image preservation is evaluated via the inverted LPIPS metric (I-LPIPS) and Perceptual Quality (P.Q.). Moreover, we leverage Disen. for the disentanglement performance in multi-attribute editing.

Model	Intensity	Control Dataset		GEdit		
		Δ VQA($\uparrow$)	I-LPIPS($\uparrow$)	S.C.($\uparrow$)	P.Q.($\uparrow$)	Disen.($\uparrow$)
Concept-slider	Slight	0.421	0.891	4.992	7.612	5.172
	Moderate	0.513	0.847	4.826	7.261	5.018
	Extreme	0.572	0.812	4.622	7.027	4.920
Qwen-Image	Slight	+3.82%	+4.12%	+7.13%	+4.27%	+5.02%
	Moderate	+5.31%	+5.44%	+6.20%	+5.11%	+3.86%
	Extreme	+7.01%	+6.89%	+6.34%	+6.18%	+3.42%
VAREdit	Slight	+5.71%	+3.16%	+6.68%	+4.56%	+4.88%
	Moderate	+6.30%	+4.75%	+5.72%	+5.75%	+4.52%
	Extreme	+6.95%	+6.37%	+5.51%	+5.78%	+4.35%
Ours	Slight	+11.44%	+3.68%	+10.28%	+4.82%	+8.42%
	Moderate	+12.72%	+5.17%	+16.33%	+3.16%	+9.68%
	Extreme	+14.58%	+7.29%	+8.24%	+3.64%	+6.78%

General attribute transfer. Our method generates smooth, continuous transformations while preserving object identity. Baselines either perform object substitution (*e.g.*, cake $\rightarrow$ apple) or fail to express the fine-grained attribute (*e.g.*, hairlessness and wrinkles in the Sphynx concept). These results demonstrate that our semantic directions capture **class-agnostic, high-level semantic properties** that existing methods cannot reliably transfer.

5.3 User Study

To complement the above evaluation, we conduct a user study to assess human perception of the generated images. We follow the standard experimental paradigm from psychophysics, namely the two-alternative forced choice (2AFC) setting [25]. In our study, participants were given a questionnaire in which they viewed pairs of images showing continuous semantic control of a reference image. For each trial, we show a reference image and two edited results produced by different methods under the same attribute prompt. Participants are asked to choose

Table 3: User study. Preference rates for Image Preservation (I.P.) and Prompt Adherence (P.A.).

Model	I.P.($\uparrow$)	P.A.($\uparrow$)
Concept-slider	10%	6%
Qwen-Image	30%	23%
VAREdit	23%	27%
Ours	37%	43%

Table 4: Comparison with a training-free text-token baseline. We report mean scores under the same GEdit protocol. The baseline directly uses the token difference in Eq. (1) as the semantic direction.

Method	Semantic Consistency ($\uparrow$)	Perceptual Quality ($\uparrow$)	Disentanglement ($\uparrow$)
$\text{mean}(T_{\text{attr}}^{\text{tar}}) - T_{\text{attr}}^{\text{src}}$	5.048	7.447	5.296
ATA (Infinity-2B)	5.374	7.584	5.455

the result that better matches the target attribute while preserving the reference content. The responses from 30 participants are summarized in Table 3. The results show a clear preference for the images generated by our method over those produced by the baselines.

5.4 Quantitative Evaluation

Table 2 reports quantitative results. Our method achieves the largest semantic shift towards the target attribute (+11.44% / +12.72% / +14.58% ΔVQA) while maintaining competitive or superior image preservation (up to +7.29% I-LPIPS at extreme intensity) on controllable dataset. Moreover, our approach outperforms existing models in semantic consistency and disentanglement performance in GEdit, particularly in complex multi-attribute editing scenarios (*e.g.*, +16.33% Semantic Consistency and +9.68% Disentanglement). Notably, our method requires neither LoRA optimization nor fine-tuning of VAR, yet outperforms instruction-based editors across all three control settings.

5.5 Ablation Studies

We analyze the key components of our design:

Injecting attribute directions to identity. Fig. 6(a) shows that adding attribute directions to identity tokens results in semantic conflicts and degraded edits. This validates our identity–attribute separation strategy.

Token-position dependence. In Fig. 6(b), applying a direction (*e.g.*, *fat*) to an unrelated token yields no meaningful change, showing that directions must be injected at their corresponding semantic positions.

Full design. Fig. 6(c) illustrates that our full model produces clean, attribute-specific edits with no identity leakage.

Training-free text-token baseline. To examine whether attribute directions can be obtained by direct text-token arithmetic, we compare ATA with a training-free baseline that uses the raw token difference in Eq. (1), *i.e.*, $\text{mean}(T_{\text{attr}}^{\text{tar}}) - T_{\text{attr}}^{\text{src}}$, as the semantic direction. Both methods are evaluated under the same GEdit protocol, and we report mean scores in Table 4. ATA consistently outperforms this baseline across all three metrics, suggesting that learned directions capture attribute factors more effectively than naive prompt-level token differences.

Comparison with Infinity-2B using Concept Slider. We further compare with an Infinity-2B baseline trained using the Concept Slider strategy. Specifically, we perform LoRA tuning for attribute control on the Control Dataset while

Table 5: Ablations on the Control Dataset. ATA with original prompts serves as the shared reference for comparison with Infinity-2B using Concept Slider and for the prompt sensitivity study.

Analysis	Method / Setting	Mean Δ VQA ($\uparrow$)	Mean I-LPIPS ($\uparrow$)
Same-backbone comparison	Infinity-2B with Concept Slider	0.532	0.879
	ATA (Infinity-2B, original prompts)	0.567	0.895
Prompt sensitivity	ATA (Infinity-2B, perturbed prompts)	0.541	0.864




Fig. 6: Ablation study. (a) Optimizing the attribute direction on the identity token ‘cat’ instead of the attribute token ‘A’ changes the subject identity in the generated image. (b) Adding the fat direction into a non-attribute token ‘ground’ does not produce a fatter cat. (c) Our method can make a fatter cat while preserving identity.

keeping the same Infinity-2B backbone as ATA. As shown in Table 5, ATA with original prompts achieves higher mean Δ VQA and mean I-LPIPS than Infinity-2B with Concept Slider. This result suggests that the advantage of ATA is not solely due to the backbone and that its token-direction formulation provides more effective attribute control.

Prompt sensitivity. ATA is not designed around a fixed prompt template. To evaluate its sensitivity to prompt wording, we perturb prompts on the Control Dataset using natural variations, including template changes (e.g., “a cat” $\rightarrow$ “a photo of a cat”), lexical paraphrases (e.g., “a cat” $\rightarrow$ “a kitty”), and context deletion (e.g., “a cat on the floor” $\rightarrow$ “a cat”). As shown in Table 5, these perturbations lead to modest drops from the original-prompt ATA setting, suggesting that ATA remains effective under common prompt variations while leaving further robustness improvements for future work.

5.6 Advanced Editing Capabilities

We further demonstrate several advanced semantic control scenarios.

Identity–attribute compositionality. Fig. 7 demonstrates that our method learns image-specific identity directions, which can be recombined with arbitrary attribute directions at inference time. Given three reference identities (a dog, a woman, and a mosaic cat), our model extracts separate identity directions and applies them to new scenes while controlling attributes such as *glasses*, *hat*, *necklace*, or *fog*. Across all cases, the synthesized images preserve the target identity



Fig. 7: Identity-attribute compositionality. We extract identity directions from three reference images and recombine them with various attribute directions. Our method preserves the identity of the chosen reference while expressing attributes such as *glasses*, *hat*, *necklace*, or *fog*, illustrating clean semantic disentanglement.



Fig. 8: Advanced Editing Capabilities. (a) Our model can both intensify (+ fat) and diminish (− fat) the same attribute, while preserving the subject’s identity and overall visual coherence. (b) Beyond simple word-level attributes, ATA enables the transfer of complex, composite attributes across distinct object categories. For example, the attribute “**wearing a shirt with the text AI**” can be consistently applied to a human, a teddy bear, and a bronze statue.

while faithfully expressing the desired attributes, confirming that identity and attribute semantics are disentangled and compositional in our representation.

Bidirectional and continuous control. Fig. 8 (a) demonstrates that our semantic directions support positive and negative traversal (*e.g.*, −fat → original → +fat), along the semantic attribute axis.

Control of complex attributes. For simple attributes (*e.g.*, fatness, emotion), a single word can describe a specific aspect of an object. However, real-world attributes are often more complex and cannot be captured by one or two words. Interestingly, our method can also transfer such complex semantic concepts, which may involve multiple attributes or even full descriptive phrases. As shown in Fig. 8 (b), beyond simple attributes, we can leverage complex descriptions as abstract attributes for image editing, further improving our flexibility.

5.7 Limitation and Discussion

Although our model can transfer general attributes and support continuous control, the editing capacity is ultimately constrained by the generative prior of




Fig. 9: Failure cases. (a) There is a bound on attribute strength; beyond this bound, artifacts appear in the generated images; (b) The method is less effective for certain particular attributes (*e.g.*, color); (c) ATA is limited by the pre-trained generative prior of the backbone.

VAR. Very strong attribute strengths may introduce artifacts or affect unrelated regions (see Fig. 9 (a)). Additionally, global color transformations remain challenging (see Fig. 9 (b)). We hypothesize that this behavior is related to VAR’s hierarchical generation mechanism. VAR predicts tokens from coarse to fine, and its localized attributes may be spatially or hierarchically entangled across neighboring tokens/scales. Such entanglement makes it difficult to disentangle these attributes under our ATA design. Moreover, due to the limited world knowledge of the base model, ATA cannot transfer concepts that the base model itself does not understand. For example, attempting to transfer the concept of a “cat” onto a woman does not result in a woman with cat-like styling, as this type of cross-domain semantic mapping is not well supported by the backbone, as shown in Fig. 9 (c). Future work will explore more autoregressive frameworks for a comprehensive validation.

6 Conclusion

Autoregressive text-to-image models have recently demonstrated strong performance as a unified framework for image generation. However, achieving fine-grained semantic control remains challenging due to the attribute entanglement and the misalignment between textual and visual representations. In this work, we have presented Attribute Token Arithmetic (ATA), a token-level semantic direction manipulation method built upon the VAR backbone. ATA enables disentangled and continuous attribute control by performing arithmetic operations directly in the VAR token space after a few optimization epochs. Experiments demonstrate the effectiveness and generalizability of ATA, showing consistent improvements over existing approaches for controllable generation.

Acknowledgements

This research was supported in part by the Australian Government through the Australian Research Council (ARC) Discovery Project DP260100218.

References

1. Achiam, J., Adler, S., Agarwal, S., Ahmad, L., Akkaya, I., Aleman, F.L., Almeida, D., Altenschmidt, J., Altman, S., Anadkat, S., et al.: Gpt-4 technical report. arXiv preprint arXiv:2303.08774 (2023)
2. Alaluf, Y., Garibi, D., Patashnik, O., Averbuch-Elor, H., Cohen-Or, D.: Cross-image attention for zero-shot appearance transfer. In: ACM SIGGRAPH. pp. 1–12 (2024)
3. Bai, S., Chen, K., Liu, X., Wang, J., Ge, W., Song, S., Dang, K., Wang, P., Wang, S., Tang, J., et al.: Qwen2. 5-vl technical report. arXiv preprint arXiv:2502.13923 (2025)
4. Chen, B., Zhao, M., Sun, H., Chen, L., Wang, X., Du, K., Wu, X.: Xverse: Consistent multi-subject control of identity and semantic attributes via dit modulation. arXiv preprint arXiv:2506.21416 (2025)
5. Chen, X., Wu, Z., Liu, X., Pan, Z., Liu, W., Xie, Z., Yu, X., Ruan, C.: Janus-pro: Unified multimodal understanding and generation with data and model scaling. arXiv preprint arXiv:2501.17811 (2025)
6. Chung, H.W., Hou, L., Longpre, S., Zoph, B., Tay, Y., Fedus, W., Li, Y., Wang, X., Dehghani, M., Brahma, S., et al.: Scaling instruction-finetuned language models. *Journal of Machine Learning Research* **25**(70), 1–53 (2024)
7. Dalva, Y., Venkatesh, K., Yanardag, P.: Fluxspace: Disentangled semantic editing in rectified flow transformers. arXiv preprint arXiv:2412.09611 (2024)
8. Dong, R., Xu, M., Geng, Z., Li, L., Hu, H., Gu, S.: Equivariant image modeling. arXiv preprint arXiv:2503.18948 (2025)
9. Esser, P., Rombach, R., Ommer, B.: Taming transformers for high-resolution image synthesis. In: CVPR. pp. 12873–12883 (2021)
10. Gandikota, R., Materzyńska, J., Zhou, T., Torralba, A., Bau, D.: Concept sliders: Lora adaptors for precise control in diffusion models. In: ECCV. pp. 172–188. Springer (2024)
11. Garibi, D., Yadin, S., Paiss, R., Tov, O., Zada, S., Ephrat, A., Michaeli, T., Mosseri, I., Dekel, T.: Tokenverse: Versatile multi-concept personalization in token modulation space. *ACM Transactions On Graphics (TOG)* **44**(4), 1–11 (2025)
12. Guerrero-Viu, J., Hasan, M., Roullier, A., Harikumar, M., Hu, Y., Guerrero, P., Gutierrez, D., Masia, B., Deschaintre, V.: Texsliders: Diffusion-based texture editing in clip space. In: ACM SIGGRAPH. pp. 1–11 (2024)
13. Han, J., Liu, J., Jiang, Y., Yan, B., Zhang, Y., Yuan, Z., Peng, B., Liu, X.: Infinity: Scaling bitwise autoregressive modeling for high-resolution image synthesis. In: CVPR. pp. 15733–15744 (2025)
14. Hertz, A., Mokady, R., Tenenbaum, J., Aberman, K., Pritch, Y., Cohen-Or, D.: Prompt-to-prompt image editing with cross attention control. arXiv preprint arXiv:2208.01626 (2022)
15. Hu, E.J., Shen, Y., Wallis, P., Allen-Zhu, Z., Li, Y., Wang, S., Wang, L., Chen, W., et al.: Lora: Low-rank adaptation of large language models. *ICLR* **1**(2), 3 (2022)
16. Lee, D., Kim, C., Kim, S., Cho, M., Han, W.S.: Autoregressive image generation using residual quantization. In: CVPR. pp. 11523–11532 (2022)
17. Li, T., Tian, Y., Li, H., Deng, M., He, K.: Autoregressive image generation without vector quantization. In: NeurIPS. vol. 37, pp. 56424–56445 (2024)
18. Lin, Z., Pathak, D., Li, B., Li, J., Xia, X., Neubig, G., Zhang, P., Ramanan, D.: Evaluating text-to-visual generation with image-to-text generation. In: ECCV. pp. 366–384. Springer (2024)

19. Liu, S., Han, Y., Xing, P., Yin, F., Wang, R., Cheng, W., Liao, J., Wang, Y., Fu, H., Han, C., et al.: Step1x-edit: A practical framework for general image editing. arXiv preprint arXiv:2504.17761 (2025)
20. Lu, C., Zhou, Y., Bao, F., Chen, J., Li, C., Zhu, J.: Dpm-solver: A fast ode solver for diffusion probabilistic model sampling in around 10 steps. In: NeurIPS. vol. 35, pp. 5775–5787 (2022)
21. Ma, C., Jiang, Y., Wu, J., Yang, J., Yu, X., Yuan, Z., Peng, B., Qi, X.: Uniook: A unified tokenizer for visual generation and understanding. arXiv preprint arXiv:2502.20321 (2025)
22. Mao, C., Zhang, J., Pan, Y., Jiang, Z., Han, Z., Liu, Y., Zhou, J.: Ace++: Instruction-based image creation and editing via context-aware content filling. arXiv preprint arXiv:2501.02487 (2025)
23. Mao, Q., Cai, Q., Li, Y., Pan, Y., Cheng, M., Yao, T., Liu, Q., Mei, T.: Visual autoregressive modeling for instruction-guided image editing. arXiv preprint arXiv:2508.15772 (2025)
24. Mikolov, T., Yih, W.t., Zweig, G.: Linguistic regularities in continuous space word representations. In: NAACL. pp. 746–751 (2013)
25. Motamed, S., Culp, L., Swersky, K., Jaini, P., Geirhos, R.: Do generative video models learn physical principles from watching videos? arXiv preprint arXiv:2501.09038 (2025)
26. Mu, J., Vasconcelos, N., Wang, X.: Editar: Unified conditional generation with autoregressive models. In: CVPR. pp. 7899–7909 (2025)
27. Patashnik, O., Garibi, D., Azuri, I., Averbuch-Elor, H., Cohen-Or, D.: Localizing object-level shape variations with text-to-image diffusion models. In: ICCV. pp. 23051–23061 (2023)
28. Podell, D., English, Z., Lacey, K., Blattmann, A., Dockhorn, T., Müller, J., Penna, J., Rombach, R.: Sdxl: Improving latent diffusion models for high-resolution image synthesis. arXiv preprint arXiv:2307.01952 (2023)
29. Ramesh, A., Dhariwal, P., Nichol, A., Chu, C., Chen, M.: Hierarchical text-conditional image generation with clip latents. arXiv preprint arXiv:2204.06125 1(2), 3 (2022)
30. Rombach, R., Blattmann, A., Lorenz, D., Esser, P., Ommer, B.: High-resolution image synthesis with latent diffusion models. In: CVPR. pp. 10684–10695 (2022)
31. Ruiz, N., Li, Y., Jampani, V., Pritch, Y., Rubinstein, M., Aberman, K.: Dreambooth: Fine tuning text-to-image diffusion models for subject-driven generation. In: CVPR. pp. 22500–22510 (2023)
32. Song, J., Meng, C., Ermon, S.: Denoising diffusion implicit models. arXiv preprint arXiv:2010.02502 (2020)
33. Song, Y., Sohl-Dickstein, J., Kingma, D.P., Kumar, A., Ermon, S., Poole, B.: Score-based generative modeling through stochastic differential equations. arXiv preprint arXiv:2011.13456 (2020)
34. Sun, P., Jiang, Y., Chen, S., Zhang, S., Peng, B., Luo, P., Yuan, Z.: Autoregressive model beats diffusion: Llama for scalable image generation. arXiv preprint arXiv:2406.06525 (2024)
35. Tang, H., Wu, Y., Yang, S., Xie, E., Chen, J., Chen, J., Zhang, Z., Cai, H., Lu, Y., Han, S.: Hart: Efficient visual generation with hybrid autoregressive transformer. In: ICLR (2025)
36. Team, G., Anil, R., Borgeaud, S., Alayrac, J.B., Yu, J., Soricut, R., Schalkwyk, J., Dai, A.M., Hauth, A., Millican, K., et al.: Gemini: a family of highly capable multimodal models. arXiv preprint arXiv:2312.11805 (2023)

37. Tian, K., Jiang, Y., Yuan, Z., Peng, B., Wang, L.: Visual autoregressive modeling: Scalable image generation via next-scale prediction. In: *NeurIPS*. vol. 37, pp. 84839–84865 (2024)
38. Tumanyan, N., Geyer, M., Bagon, S., Dekel, T.: Plug-and-play diffusion features for text-driven image-to-image translation. In: *CVPR*. pp. 1921–1930 (2023)
39. Van Den Oord, A., Vinyals, O., et al.: Neural discrete representation learning. In: *NeurIPS*. vol. 30 (2017)
40. Wan, T., Wang, A., Ai, B., Wen, B., Mao, C., Xie, C.W., Chen, D., Yu, F., Zhao, H., Yang, J., et al.: Wan: Open and advanced large-scale video generative models. *arXiv preprint arXiv:2503.20314* (2025)
41. Wang, Y., Guo, L., Li, Z., Huang, J., Wang, P., Wen, B., Wang, J.: Training-free text-guided image editing with visual autoregressive model. *arXiv preprint arXiv:2503.23897* (2025)
42. Wang, Y., Lin, Z., Teng, Y., Zhu, Y., Ren, S., Feng, J., Liu, X.: Bridging continuous and discrete tokens for autoregressive visual generation. *arXiv preprint arXiv:2503.16430* (2025)
43. Wu, C., Li, J., Zhou, J., Lin, J., Gao, K., Yan, K., Yin, S.m., Bai, S., Xu, X., Chen, Y., et al.: Qwen-image technical report. *arXiv preprint arXiv:2508.02324* (2025)
44. Wu, C., Zheng, P., Yan, R., Xiao, S., Luo, X., Wang, Y., Li, W., Jiang, X., Liu, Y., Zhou, J., et al.: Omnigen2: Exploration to advanced multimodal generation. *arXiv preprint arXiv:2506.18871* (2025)
45. Wu, Y., Luo, X., Xu, Z., Guo, X., Ju, L., Ge, Z., Liao, W., Cai, J.: Diversified and personalized multi-rater medical image segmentation. In: *CVPR*. pp. 11470–11479 (June 2024)
46. Wu, Y., Song, T., Wu, Z., Ye, J., Ge, Z., Bai, W., Chen, Z., Cai, J.: Virtual full-stack scanning of brain mri via imputing any quantised code. In: *CVPR*. pp. 21026–21035 (June 2026)
47. Xing, J., Xia, M., Zhang, Y., Chen, H., Yu, W., Liu, H., Liu, G., Wang, X., Shan, Y., Wong, T.T.: Dynamicrafter: Animating open-domain images with video diffusion priors. In: *ECCV*. pp. 399–417. Springer (2024)
48. Yang, X., Li, B., Zhang, Y., Yin, Z., Bai, L., Ma, L., Wang, Z., Cai, J., Wong, T.T., Lu, H., Jia, X.: Vlipp: Towards physically plausible video generation with vision and language informed physical prior. In: *ICCV*. pp. 12360–12370 (October 2025)
49. Yang, Z., Teng, J., Zheng, W., Ding, M., Huang, S., Xu, J., Yang, Y., Hong, W., Zhang, X., Feng, G., et al.: Cogvideox: Text-to-video diffusion models with an expert transformer. *arXiv preprint arXiv:2408.06072* (2024)
50. Yu, J., Li, X., Koh, J.Y., Zhang, H., Pang, R., Qin, J., Ku, A., Xu, Y., Baldridge, J., Wu, Y.: Vector-quantized image modeling with improved vqgan. In: *ICLR* (2021)
51. Yu, Q., Weber, M., Deng, X., Shen, X., Cremers, D., Chen, L.C.: An image is worth 32 tokens for reconstruction and generation. In: *NeurIPS*. vol. 37, pp. 128940–128966 (2024)
52. Zhang, R., Isola, P., Efros, A.A., Shechtman, E., Wang, O.: The unreasonable effectiveness of deep features as a perceptual metric. In: *CVPR*. pp. 586–595 (2018)
53. Zhao, Y., Xiong, Y., Kraehenbuehl, P.: Image and video tokenization with binary spherical quantization. In: *ICLR* (2024)
54. Zheng, Z., Peng, X., Yang, T., Shen, C., Li, S., Liu, H., Zhou, Y., Li, T., You, Y.: Open-sora: Democratizing efficient video production for all. *arXiv preprint arXiv:2412.20404* (2024)
55. Zhu, L., Wei, F., Lu, Y., Chen, D.: Scaling the codebook size of vq-gan to 100,000 with a utilization rate of 99%. In: *NeurIPS*. vol. 37, pp. 12612–12635 (2024)