

# Unified Personalized Understanding, Generating and Editing

Yu Zhong<sup>1</sup>, Tianwei Lin<sup>1</sup>, Ruike Zhu<sup>1</sup>, Yuqian Yuan<sup>1</sup>, Haoyu Zheng<sup>1</sup>, Liang Liang<sup>1</sup>,  
Wenqiao Zhang<sup>1,†</sup>, Feifei Shao<sup>1</sup>, Haoyuan Li<sup>2</sup>, Wanggui He<sup>2</sup>, Hao Jiang<sup>2</sup>, Yueting Zhuang<sup>1</sup>

<sup>1</sup>Zhejiang University, <sup>2</sup>Alibaba Group

<https://zyyy9.github.io/OmniPersona/>

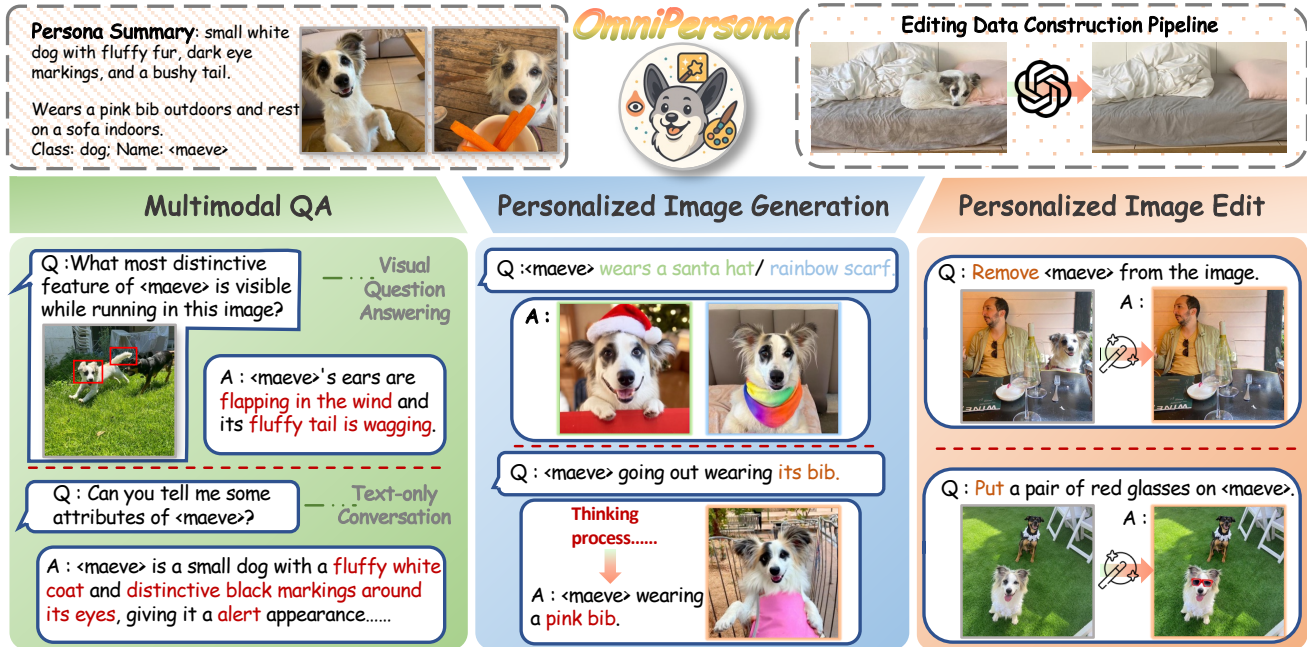


Figure 1. **Capability Overview of OmniPersona.** OmniPersona leverages decoupled learnable prompts to achieve **unified personalized understanding, generation, and editing** from only a few concept images paired with textual descriptions. Notably, OmniPersona is the first framework to enable personalized image editing, a critical capability overlooked by previous works.

## Abstract

Unified large multimodal models (LMMs) have achieved remarkable progress in general-purpose multimodal understanding and generation. However, they still operate under a “one-size-fits-all” paradigm and struggle to model user-specific concepts (e.g., generate a photo of <maeve>) in a consistent and controllable manner. Existing personalization methods typically rely on external retrieval, which is inefficient and poorly integrated into unified multimodal pipelines. Recent personalized unified models introduce learnable soft prompts to encode concept information, yet they either couple understanding and generation or depend on complex multi-stage training, leading to cross-task interference and ultimately to fuzzy or misaligned personalized knowledge. We present **OmniPersona**, an end-to-end personalization framework for unified LMMs that, for

the first time, integrates personalized understanding, generation, and image editing within a single architecture. OmniPersona introduces structurally decoupled concept tokens, allocating dedicated subspaces for different tasks to minimize interference, and incorporates an explicit knowledge replay mechanism that propagates personalized attribute knowledge across tasks, enabling consistent personalized behavior. To systematically evaluate unified personalization, we propose **OmniPBench**, extending the public UnifyBench concept set with personalized editing tasks and cross-task evaluation protocols integrating understanding, generation, and editing. Experimental results demonstrate that OmniPersona delivers competitive and robust performance across diverse personalization tasks. We hope OmniPersona will serve as a strong baseline and spur further research on controllable, unified personalization.

† denotes corresponding author

# 1. Introduction

Recently, unified large multimodal models (LMMs) for joint understanding and generation have demonstrated remarkable potential as general-purpose AI assistants [8, 13, 25, 40, 42]. Models such as Chameleon [24], Janus-Series [5, 16, 27], Blip3-o [4], Show-o2 [31] and Bagel [6] can reliably follow user instructions across diverse task formats, indicating that unified LMMs exhibit strong generalization ability over a broad spectrum of multimodal tasks in open-world settings. In realistic human-AI interaction, however, user demands are often organized around their own entities and preferences, requiring assistants to understand and persistently memorize personalized concepts and to align with user intent under brief and natural instructions. Existing unified models remain substantially constrained in such user-level scenarios. For example, when using the identifier `<maeve>` for a specific dog, even with attribute descriptions, current models still struggle to handle `<maeve>` consistently and controllably. Consequently, robustly integrating personalized concepts into unified LMMs is crucial for building truly personal AI assistants.

To enable personalization in LMMs, some efforts have been made. One line of work uses retrieval-augmented generation (RAG) [9], injecting concept attributes and descriptions as external context; another learns soft prompts to encode specific concepts in the latent space of the model [1, 17].

In parallel, several methods propose dedicated personalization mechanisms for unified LMMs and report promising results on certain benchmarks [3, 18], yet three key challenges remain for achieving comprehensive and robust personalization in unified modeling:

**(i) Coupled and Conflicting Representations.** In unified LMMs, sharing a common parameter space across tasks naturally induces representational conflict [12, 27, 40], which is further exacerbated in few-shot personalization, where a single concept representation must simultaneously support understanding, generation, and editing. Existing methods (e.g., UniToken [11]) mitigate such interference via multi-stage or alternating training, but personalized information remains compressed into a single latent space, without structurally distinguishable “slots” to host task-specific solution spaces. As a result, it is difficult to selectively leverage pretrained knowledge for understanding versus generation in a targeted and decoupled manner.

**(ii) Opaque Knowledge Latents.** Existing methods compress concept knowledge into opaque embeddings that lack explicit interpretability—we cannot verify what knowledge is encoded or how it drives the model’s understanding and generation. This opacity poses a critical challenge for *personalized attribute-reasoning generation (PARG)* [3], where models generate images that must leverage learned textual attributes (e.g., “`<wangkai>` in his home” without



Figure 2. Current state-of-the-art personalized unified models fail to stably leverage textual knowledge during generation, resulting in images misaligned with concept descriptions (top). Moreover, existing methods neglect personalized image editing, producing outputs unrelated to the input image (bottom).

specifying the home’s features). As shown in Fig. 2, even approaches like Yo’Chameleon [18] and UniCTokens [3], which attempt to enhance knowledge expression via implicit prompts and shared tokens, struggle to distinguish whether the model genuinely leverages textual attributes or merely memorizes visual patterns from training data. Developing interpretable knowledge representations that enable genuine PARG capabilities remains an open challenge.

**(iii) Personalized Editing Gap.** Prior works do not examine the feasibility of personalized image editing (see Fig. 2). However, editing is arguably the most demanding test of personalized tasks: the model must follow user instructions to perform local or structural modifications while accurately locating and preserving the target identity [19, 41], effectively coupling understanding and generation within a single operation. At the same time, current benchmarks neither systematically evaluate this crucial capability nor investigate whether incorporating editing data can in turn enhance performance across other personalization tasks.

To address these challenges, we propose **OmniPersona**, a novel end-to-end framework for unified personalized understanding, generation, and editing. The proposed framework is composed of the following key components:

**Token Representations Decoupling.** We relax the assumption that a single shared parameter space suffices for personalization by introducing structurally decoupled concept tokens and assigning them to task-specialized expert subspaces, so that the same personalized concept is captured by related but distinct tokens along different task pathways. This decoupling design mitigates cross-task interference and enables selective reuse of pretrained knowledge for understanding versus generation.

**Explicit Knowledge Replay.** OmniPersona introduces a reasoning-based knowledge replay mechanism that externalizes personalized knowledge as explicit intermediate representations. Specifically, the model self-generates descriptions of concept attributes as conditioning signals and replays the concept–text–image pathway, explicitly aligning knowledge acquired in understanding with downstream generation and enforcing semantic consistency of the concept across unified tasks.

**Understanding–Generation–Editing Synergy.** We emphasize personalized editing as a crucial component and investigate whether the fine-grained constraints imposed by editing can in turn regularize and strengthen unified personalized representations. To this end, we propose **OmniPBench**, which extends the existing unified personalization benchmark UnifyBench [3] with carefully curated personalized editing data and cross-task evaluation protocols spanning understanding, generation, and editing, enabling a systematic assessment of how editing supervision benefits overall personalization performance.

In summary, our contributions are as follows:

- To the best of our knowledge, the proposed **OmniPersona** serves as the first end-to-end unified framework that simultaneously achieves personalized image understanding, generation, and editing.
- Our framework combines decoupled token representations to reduce task conflict, and uses inference-time explicit knowledge replay to improve personalized attribute-reasoning generation stability and interpretability.
- We propose **OmniPBench** with integrated editing data to systematically assess synergy across tasks.
- Comprehensive experiments demonstrate competitive performance across personalized tasks and state-of-the-art (SOTA) results in personalized editing.

## 2. Related Work

**Unified Large Multimodal Models.** As multimodal understanding and generation advance, unified models become increasingly crucial across domains [14, 37, 39]. In recent years, unified large multimodal models have attracted significant attention [13, 15, 16, 26, 28, 30]. Early works such as Janus-Pro [5] and UniToken [11] adopt a pure autoregressive paradigm, converting multiple modalities into discrete vocabularies and performing joint modeling through a shared decoding process, thereby unifying multimodal inputs and outputs within a common semantic space. However, these approaches often suffer from low generation efficiency and limited flexibility in content synthesis. To enhance generation quality and training efficiency, several studies such as SEED-X [8], MetaQueries [20], and BLIP3-o [4] introduce diffusion-based generation heads, which enable rapid alignment between understanding and generation capabilities using limited multimodal paired data, and even

demonstrate synergistic improvements across tasks. Nevertheless, these methods remain constrained by the sequential decoding bottleneck, making real-time high-resolution generation challenging. To address this issue, subsequent works (e.g., Transfusion [42], Show-o [31] and BAGEL [6]) combine the strengths of autoregressive and diffusion models to construct hybrid architectures that maintain the unified multimodal framework while balancing generation flexibility and computational efficiency. As unified models continue to mature, an emerging research direction lies in effectively customizing and adapting them to recognize and leverage user-specific concepts, enabling personalized understanding and generation within a unified modeling paradigm.

**Personalized Multimodal Understanding and Generation.** Personalized adaptation aims to enable models to recognize and utilize specific concepts, achieving rapid user-level semantic alignment [33–36, 38]. MyVLM [1] and Yo’LLaVA [17] inject personalized concepts through learnable soft prompts, avoiding lengthy contextual descriptions while enhancing understanding capability. Retrieval-augmented approaches such as LaMP [22] and RAP [9] treat conceptual profiles as external memory, leveraging contextual retrieval to improve personalized reasoning. MC-LLaVA [2] further optimizes multi-concept discrimination and co-occurrence understanding through joint training and visual token initialization, enhancing generalization in multi-entity scenarios. On the generation side, works such as DreamBooth [21] achieve high-fidelity subject-driven image synthesis via few-shot fine-tuning of text-to-image diffusion models, laying the foundation for personalized visual generation [10, 23, 29, 32]. More recently, Yo’Chameleon [18] unifies personalized understanding and generation within a single framework through dual soft prompts and a self-prompting mechanism, while UniCTokens [3] integrates complementary semantics of both tasks via unified concept tokens and a three-stage mutual learning paradigm, achieving higher efficiency and performance in unified personalized modeling. However, existing approaches still lack an end-to-end personalized modeling mechanism that can fully exploit the potential of unified multimodal architectures.

## 3. Methodology

We propose an end-to-end personalized multimodal framework that unifies understanding, generation, and editing with token representations decoupling, explicit knowledge replay, and understanding–generation–editing synergy.

### 3.1. Token Representations Decoupling

To mitigate the representational entanglement discussed in Sec. 1, we introduce token representations decoupling. Instead of learning a unified set of prompts, we route task-

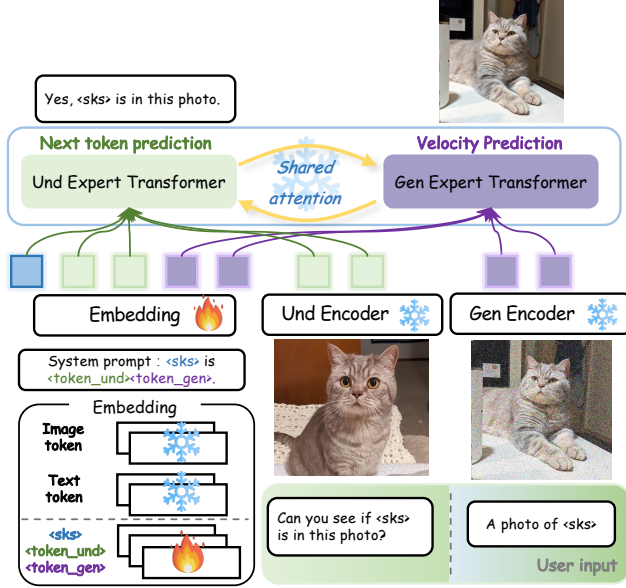


Figure 3. **Model Overview of OmniPersona.** We employ an end-to-end dual-branch routing mechanism to train two sets of parameter-decoupled prompts.

specific learnable tokens to specialized expert subspaces, enabling each branch to develop modality-aware representations while preserving cross-task synergy underpinned by the inherent modality-interaction of the unified model.

Formally, during fine-tuning, we represent each personalized concept using a special identifier (e.g.,  $\langle sks \rangle$ ) composed of multiple learnable tokens that are routed to different expert branches. The system prompt is structured as:

$$\langle sks \rangle \text{ is } \langle und_1 \rangle \dots \langle und_{N_u} \rangle \langle gen_1 \rangle \dots \langle gen_{N_g} \rangle.$$

where tokens enclosed in angle brackets are learnable concept embeddings. Specifically, as shown in Fig. 1, the concept identifier  $\langle sks \rangle$  together with  $\langle und_i \rangle$  tokens are routed to the *understanding expert*, while  $\langle gen_j \rangle$  tokens are routed to the *generation expert*. Let

$$\mathbf{P}^{(und)} = [\mathbf{p}_{sks}, \mathbf{p}_1^{(und)}, \dots, \mathbf{p}_{N_u}^{(und)}] \in \mathbb{R}^{(N_u+1) \times d},$$

$$\mathbf{P}^{(gen)} = [\mathbf{p}_1^{(gen)}, \dots, \mathbf{p}_{N_g}^{(gen)}] \in \mathbb{R}^{N_g \times d},$$

denote the embedding matrices for the two expert-specific token subsets, where  $d$  is the embedding dimension. Note that  $\mathbf{p}_{sks}$  is included in the understanding expert to enable concept identification. During the forward pass, each expert processes only its routed prompt subset and corresponding input embeddings:

$$\mathbf{H}^{(und)} = \mathcal{F}_{und}(\mathbf{P}^{(und)}, \mathbf{X}^{(und)}), \quad \mathbf{H}^{(gen)} = \mathcal{F}_{gen}(\mathbf{P}^{(gen)}, \mathbf{X}^{(gen)}), \quad (1)$$

where  $\mathcal{F}_{und}(\cdot)$  and  $\mathcal{F}_{gen}(\cdot)$  denote the expert branches of the unified model for understanding and generation, respectively, with  $\mathbf{X}^{(und)}$  and  $\mathbf{X}^{(gen)}$  representing the corresponding personalized instructions.

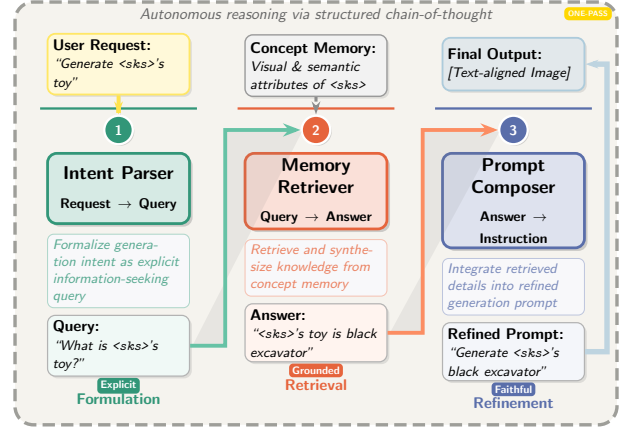


Figure 4. Explicit knowledge replay pipeline.

This routing formulation enforces expert-level parameter decoupling, enabling each branch to specialize on its designated modality while preserving synergy under a universal unified model architecture.

### 3.2. Explicit Knowledge Replay

While token decoupling mitigates cross-task interference, the learned embeddings remain opaque—it is unclear whether the model genuinely leverages semantic knowledge or merely memorizes visual patterns. To address this, we propose explicit knowledge replay, an *inference-time* mechanism that actively retrieves and externalizes concept knowledge. As shown in Fig. 4, we introduce a three-stage framework that transforms implicit generation requests into explicit, grounded, and faithful prompts. Importantly, all the aforementioned stages are recursively executed within a single unified multimodal model  $\mathcal{UMM}$ , achieving one-step generation.

**Stage 1: Intent Parser.** Given the original user request  $T$ , we reformulate it into an explicit information-seeking query using  $\mathcal{UMM}$  in text mode:

$$Q = \mathcal{UMM}^{\text{text}}(T), \quad (2)$$

which transforms generation intent into explicit queries  $Q$  such as “What is  $\langle sks \rangle$ ’s [attribute]?”.

**Stage 2: Memory Retriever.** The formulated query  $Q$  is used to retrieve concept memory from the learned token representations  $\mathbf{P}$ :

$$\mathcal{A} = \mathcal{UMM}^{\text{text}}(Q, \mathbf{P}), \quad (3)$$

where  $\mathcal{A}$  represents the grounded answer, e.g., “ $\langle sks \rangle$ ’s toy is a black excavator.”

**Stage 3: Prompt Composer.** The retrieved answer  $\mathcal{A}$  is integrated back into the generation context:

$$\hat{T} = \mathcal{UMM}^{\text{text}}(\mathcal{A}, T), \quad (4)$$

producing a refined, faithful prompt such as “Generate <sk>’ black excavator.”

**Final Generation.** The refined prompt  $\hat{T}$  is passed through the model’s generation mode for personalized synthesis:

$$I_{\text{gen}} = \mathcal{U}\mathcal{M}\mathcal{M}^{\text{gen}}(\hat{T}, \mathbf{P}). \quad (5)$$

This unified pipeline ensures outputs are (1) *explicit* through query formulation, (2) *grounded* through memory retrieval, (3) *faithful* through refinement, and (4) *personalized* through concept-aware generation, all within a single model in one pass.

### 3.3. Understanding-Generation-Editing Synergy

While existing personalized multimodal models focus on understanding and text-to-image generation, they overlook the editing dimension and lack suitable datasets for personalized editing. To enhance conceptual understanding and editing capability, we introduce a **personalized image editing dataset**.

We construct a personalized editing dataset  $\mathcal{D}_{\text{edit}}$ , where each sample is a triplet  $(I_{\text{src}}, E_{\text{edit}}, I_{\text{tgt}})$  consisting of a source image  $I_{\text{src}}$ , an editing instruction  $E_{\text{edit}}$  (e.g., “remove <sk> in the photo”), and the target image  $I_{\text{tgt}}$ . The model is trained end-to-end, jointly optimizing all soft tokens across understanding, generation, and editing tasks.

**Training Objectives.** The model is jointly optimized via two complementary losses:

(i) **Cross-Entropy Loss** for autoregressive token prediction based on the personalized understanding task:

$$\mathcal{L}_{\text{text}}^{\text{CE}} = - \sum_{i=1}^C x_i \log \hat{x}_i, \quad (6)$$

where  $x_i$  and  $\hat{x}_i$  denote the ground-truth and predicted token distributions, and  $C$  is the number of classes.

(ii) **Mean Squared Error Loss** following the Rectified Flow paradigm. Given a clean latent  $\mathbf{x}_0$  and noise  $\mathbf{x}_1$ , we obtain the interpolated latent  $\mathbf{x}_t = (1 - t)\mathbf{x}_0 + t\mathbf{x}_1$  for  $t \in [0, 1]$ . The MSE objective for image generation is:

$$\mathcal{L}_{\text{image}}^{\text{MSE}} = \mathbb{E}[\|g_{\theta}(\mathbf{x}_t | c) - (\mathbf{x}_0 - \mathbf{x}_1)\|_2^2], \quad (7)$$

where  $g_{\theta}$  predicts the velocity field conditioned on  $\mathbf{x}_t$  and context  $c$ . Similarly, the editing loss is also defined as MSE:

$$\mathcal{L}_{\text{edit}}^{\text{MSE}} = \mathbb{E}[\|g_{\theta}(\mathbf{x}_t | c_{\text{edit}}) - (\mathbf{x}_0^{\text{tgt}} - \mathbf{x}_1^{\text{tgt}})\|_2^2], \quad (8)$$

where  $c_{\text{edit}}$  includes the source image and editing instruction, and  $\mathbf{x}_0^{\text{tgt}}$  is the target latent.

**Joint Training with Editing Data.** The total objective combines understanding, generation, editing optimizations:

$$\mathcal{L}_{\text{total}} = \mathcal{L}_{\text{text}}^{\text{CE}} + \lambda_{\text{image}} \mathcal{L}_{\text{image}}^{\text{MSE}} + \lambda_{\text{edit}} \mathcal{L}_{\text{edit}}^{\text{MSE}}, \quad (9)$$

where  $\lambda_{\text{image}}$  and  $\lambda_{\text{edit}}$  control the relative contribution of image generation and editing.

## 4. Experiment

We conduct comprehensive experiments to evaluate our framework across four personalized tasks: *understanding*, *generation*, *personalized attribute-reasoning generation*, and *image editing*, where editing represents a novel evaluation dimension absent in prior benchmarks.

**Implementation Details.** We allocate  $N = 32$  learnable tokens per concept, split equally between understanding-specific tokens ( $N_u = 16$ ) and generation-specific tokens ( $N_g = 16$ ). We optimize using AdamW for 2,000 steps per concept with batch size 8. All experiments are conducted on H20 GPUs using Bagel [6] as the backbone model. Additional training details are provided in the Appendix.

**Dataset.** We introduce **OmniPBench**, a comprehensive benchmark for unified personalization spanning understanding, generation, and editing. Building upon the 20 concepts from UnifyBench [3], which include Person (10), Pets (5), and Objects (5), we contribute a novel **personalized editing dataset**. For each concept with identifier <sk> and class label  $c$ , we automatically generate removal instructions following the template “remove <sk>:  $c$  from the photo”. Target images are synthesized via inpainting models. We perform manual verification, with 90% of generated pairs passing quality control. Moreover, we design diverse edit-instruction templates to systematically probe each concept’s editing robustness under various modification scenarios. Additional details are provided in the Appendix.

**Baselines.** We compare against four categories of personalization methods: *understanding-only* approaches [1, 17], *generation-only* methods [7, 21], *unified personalization* frameworks [3, 18], and *retrieval-augmented* systems [9]. To establish context-aware performance, we additionally include *Bagel* with text-only descriptions (no learned tokens) as a zero-shot baseline.

**Metrics.** For personalized image editing, our key contribution, we employ an LLM-as-judge protocol to assess semantic alignment between edited images and textual instructions, and edit quality evaluated by visual naturalness. For understanding, generation, and personalized attribute-reasoning generation, we adopt standard UnifyBench metrics. All results are reported as mean scores across concepts.

### 4.1. Personalized Understanding and Generation

As shown in Table 1, our method achieves strong performance across understanding and generation tasks. Compared to SOTA unified models, we outperform Unictoken by 7.8% on recognition and 13.1% on average across question-answering tasks (VQA and text-based QA). For generation, we achieve the highest CLIP-I score (0.791), indicating superior alignment with personalized visual concepts, alongside substantial improvements in identity-critical metrics: face similarity (0.413 vs. 0.334).

Table 1. Quantitative comparison on **OmniPBench**. TP = Text Prompt. IP = Image Prompt. PARG = Personalized Attribute-Reasoning Generation. SEMA-C = Semantic Consistency. QUAL-I = Quality of Image. We compare open-source unified models; others are for reference. The **best** are highlighted.

Method	Model Size	Token	Training Images	Personalized Understanding				Personalized Generation				PARG		Personalized Edit				
				Rec.		VQA		QA		Pure Gen.		People Gen.		Score	CLIP-I	SEMA-C	QUAL-I	Avg.
				Weight	BLEU	GPT	BLEU	GPT	CLIP-I	CLIP-T	DINO	Face-Simi						
<b>Upper Bound</b>																		
GPT-4o+TP	200B	~100	-	0.742	0.473	0.676	0.610	0.685	0.689	0.301	0.626	0.198	0.780	0.690	0.463	0.645	0.554	
GPT-4o+IP	200B	~1,000	-	0.773	0.543	0.685	0.589	0.652	0.794	0.310	0.722	0.559	0.782	0.792	0.512	0.603	0.558	
Real Images	-	-	-	-	-	-	-	-	0.832	-	0.728	0.739	-	0.832	-	-	-	
<b>Understanding Only</b>																		
Yo'LLaVA	13B	16	~100	0.919	0.609	0.629	0.612	0.593	-	-	-	-	-	-	-	-	-	
MC-LLaVA	13B	16	~10	0.924	0.628	0.637	0.601	0.583	-	-	-	-	-	-	-	-	-	
RAP-MLLM	13B	~1,000	-	0.940	0.616	0.616	0.712	0.722	-	-	-	-	-	-	-	-	-	
Qwen2.5-VL + TP	3B	~100	-	0.660	0.407	0.727	0.574	0.774	-	-	-	-	-	-	-	-	-	
Yo'LLaVA(Phi-1.5)	1.3B	16	~100	0.765	0.488	0.497	0.510	0.494	-	-	-	-	-	-	-	-	-	
<b>Generation Only</b>																		
Text inversion	1.0B	-	~10	-	-	-	-	-	0.630	0.247	0.569	0.371	0.070	0.628	-	-	-	
DreamBooth (SD)	1.0B	-	~10	-	-	-	-	-	0.649	0.281	0.591	0.436	0.071	0.650	-	-	-	
<b>Unified Model</b>																		
Chameleon+TP	7B	~100	-	0.690	0.413	0.488	0.509	0.564	0.547	0.176	0.509	0.011	0.329	0.549	0.122	0.520	0.321	
Chameleon+IP	7B	~1,000	-	0.493	0.445	0.498	0.407	0.535	0.523	0.160	0.469	0.066	0.299	0.499	0.093	0.442	0.268	
Show-o+TP	1.3B	~100	-	0.566	0.461	0.409	0.504	0.579	0.664	0.264	0.553	0.048	0.770	0.660	0.250	0.340	0.295	
Bagel+TP	7B MoT	~100	-	0.788	0.420	0.542	0.445	0.632	0.697	0.284	0.425	0.309	0.813	0.725	0.297	0.566	0.432	
Yo'Chameleon	7B	32	~1,000	0.764	0.474	0.507	0.510	0.581	0.697	0.236	0.590	0.224	0.266	0.698	0.108	0.360	0.234	
Unictoken	1.3B	32	~10	0.790	0.503	0.523	0.544	0.603	0.750	0.282	0.646	0.334	0.359	0.749	0.155	0.212	0.184	
<b>OmniPersona (Ours)</b>	7B MoT	<b>32</b>	<b>~10</b>	<b>0.852</b>	<b>0.529</b>	<b>0.603</b>	<b>0.646</b>	<b>0.678</b>	<b>0.791</b>	<b>0.273</b>	<b>0.579</b>	<b>0.413</b>	<b>0.613</b>	<b>0.788</b>	<b>0.711</b>	<b>0.605</b>	<b>0.658</b>	

While our DINO (0.579) is lower than Unictoken (0.646) and our CLIP-T (0.273) is lower than Bagel+TP (0.284), this reflects our design priority: *preserving concept identity over generic text alignment*. The highest CLIP-I validates this trade-off, and our CLIP-T still outperforms other unified models (Show-o: 0.264, Yo'Chameleon: 0.236), confirming that identity-first optimization does not compromise overall text-image consistency. Compared to Bagel+TP (0.788), which directly injects concept descriptions as long context, we achieve superior understanding performance (0.852) with more compact token representations, demonstrating more efficient knowledge internalization.

## 4.2. Personalized Attribute-Reasoning Generation

Table 1 evaluates personalized attribute-reasoning generation (PARG), where the model must leverage concept-related textual knowledge during generation. For example, given "Generate a photo of <wangkai> in his home", the model must recall that <wangkai>'s home is by the sea, a piece of knowledge that is absent from visual priors. This tests explicit textual knowledge integration rather than image memorization.

With our explicit knowledge replay mechanism, we achieve SOTA CLIP-I (0.788), outperforming Unictoken (0.749) by 5.2%, demonstrating superior identity

preservation in knowledge-conditioned generation. While Bagel+TP achieves a higher overall PARG score (0.813 vs. 0.613) by directly injecting long-context descriptions, our approach internalizes knowledge in compact learnable tokens without increasing inference overhead. Notably, our score (0.613) substantially surpasses other learned-token methods like Unictoken (0.359) by 70.8%, validating that explicit knowledge replay effectively bridges understanding and generation.

## 4.3. Personalized Image Editing

We introduce the first systematic evaluation of personalized image editing for unified models. As shown in Table 1, we assess editing capability through LLM-as-judge metrics measuring (1) semantic alignment with instructions (SEMA-C) and (2) edited image quality (QUAL-I).

Our method achieves SOTA overall editing performance (0.658), surpassing GPT-4o+IP by 17.9% (0.558) and demonstrating the effectiveness of our unified architecture for identity-preserving editing. Notably, we excel in semantic alignment (SEMA-C: 0.711), substantially outperforming all baselines including Bagel+TP (0.297), validating that our decoupled token representations enable precise instruction following while maintaining concept fidelity. Our competitive image quality (QUAL-I: 0.605) further con-

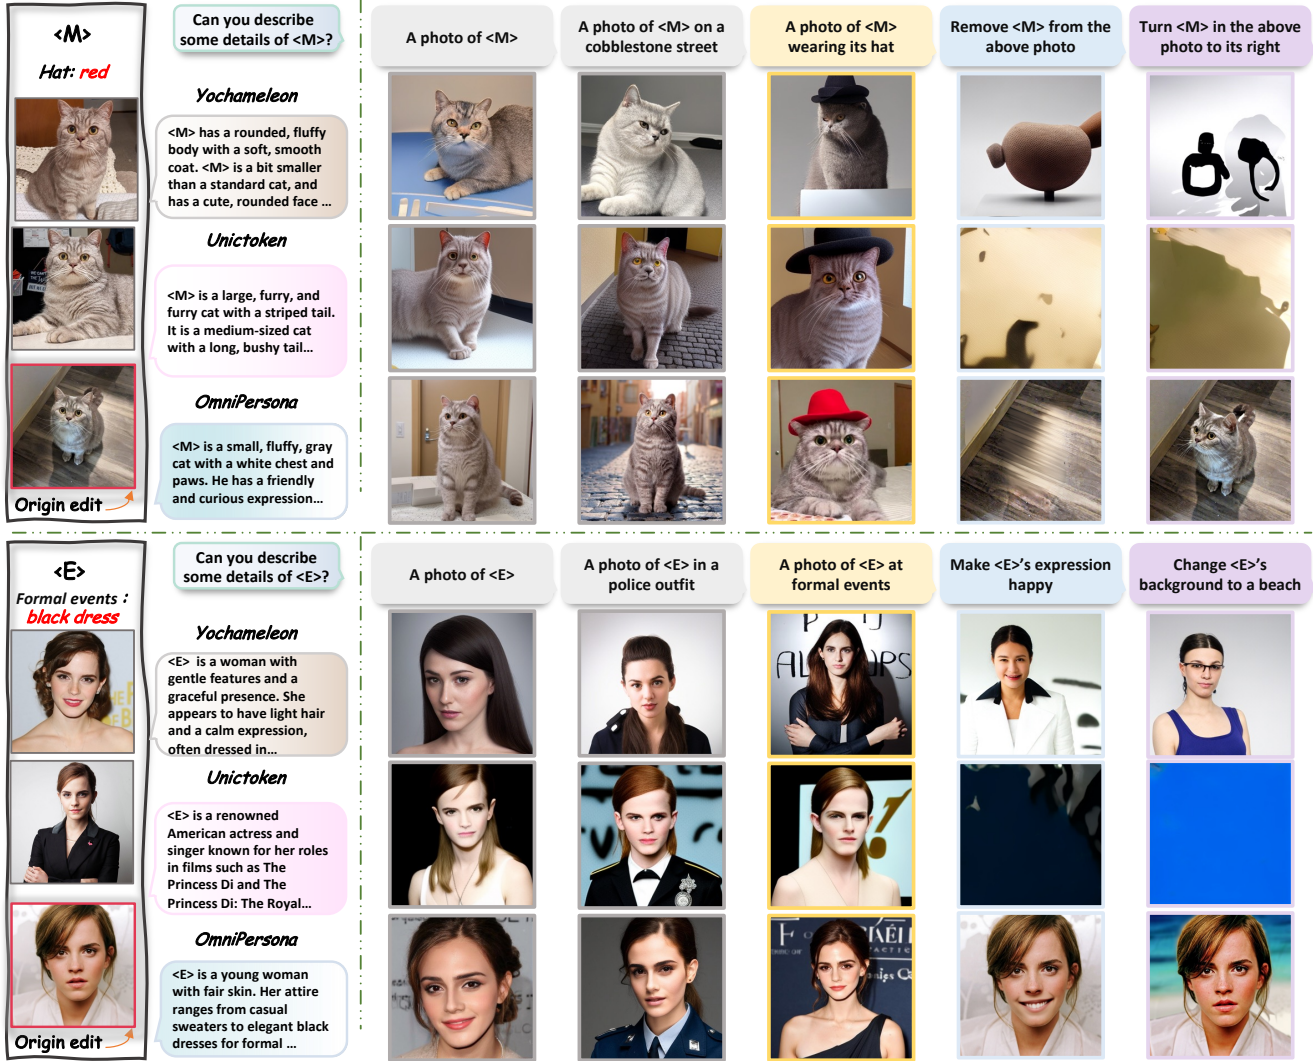


Figure 5. Qualitative comparison with Yo’Chameleon and Unictoken on **OmniPBench**. **OmniPersona (Ours)** demonstrates more precise, personalized image generation and excellent capability of personalized image editing.

firmly that explicit knowledge replay preserves visual coherence during complex editing operations.

## 5. Ablation Studies and Analysis

**Decoupling Feature Space.** Fig. 6 and Table 2 validate our token decoupling mechanism. Without disentangled tokens, all embeddings share a unified transformer, causing severe representational entanglement (Fig. 6 (a), left), where overlapping clusters indicate destructive task interference. This directly degrades understanding performance: recognition drops 5.6% (0.852→0.804), VQA BLEU decreases 9.8% (0.529→0.477), and VQA GPT falls 10.3% (0.603→0.541). Understanding tasks suffer disproportionately because concept features become contaminated by generation-specific gradients in the shared latent space,

forcing conflicting optimization objectives.

Our decoupled architecture (Fig. 6 (a), right) produces well-separated clusters, routing understanding and generation tokens to task-specialized expert subspaces. This enables selective reuse of pretrained knowledge: understanding tokens leverage semantic reasoning without interference from synthesis priors. The consistent gains across all understanding metrics validate our core insight: structurally distinguishable token “slots” mitigate gradient conflicts and unlock the model’s capacity for unified personalization.

**Knowledge Replay Promotes Personalized Knowledge Flow.** Table 2 validates the critical role of knowledge replay: without it, PARG score plummets 49.1% (0.613→0.312) while CLIP-I remains relatively stable (0.788→0.765). This reveals that implicit embeddings preserve identity but fail to leverage semantic knowledge,

Table 2. Ablation study on **OmniPBench**.

Methods	Personalized Understanding					Personalized Generation			PARG	Personalized Edit				
	Rec.	VQA		QA		Pure Gen.				People Gen.	Personalized Edit			
	Weight	BLEU	GPT	BLEU	GPT	CLIP-I	CLIP-T	DINO	Face-Simi	Score	CLIP-I	SEMA-C	Qual-I	Avg.
W/o Disentangled Tokens	0.804	0.477	0.541	0.589	0.652	0.778	0.289	0.585	0.402	0.431	0.732	0.652	0.613	0.633
W/o Knowledge Replay	-	-	-	-	-	-	-	-	-	0.312	0.765	-	-	-
W/o Edit Data	0.840	0.517	0.614	0.632	0.659	0.772	0.296	0.553	0.397	0.584	0.781	0.683	0.592	0.638
Ours (Full)	0.852	0.529	0.603	0.646	0.678	0.791	0.273	0.579	0.413	0.613	0.788	0.711	0.605	0.658

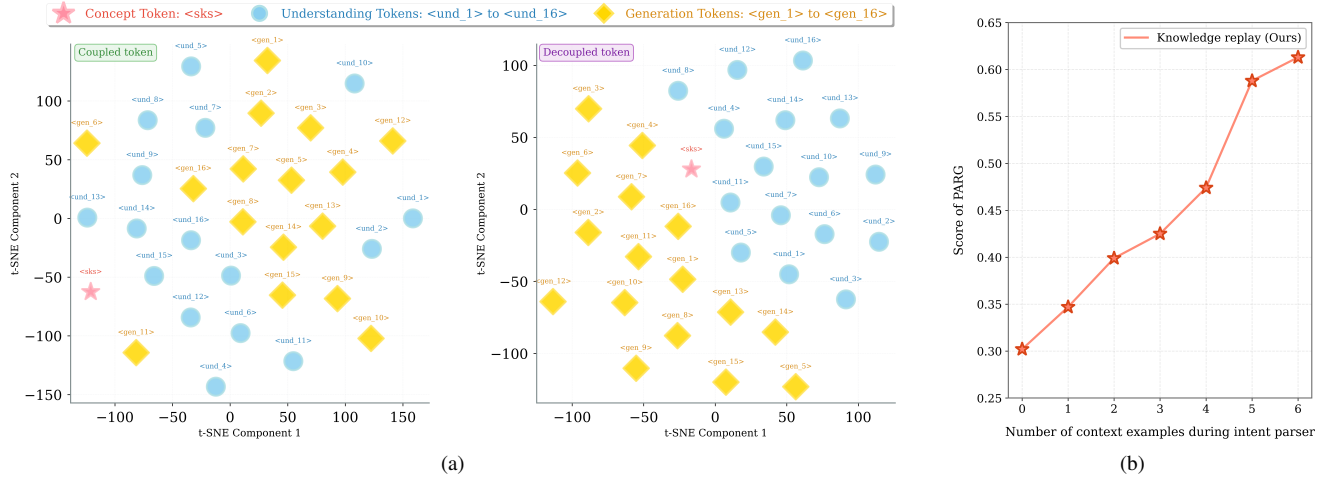


Figure 6. Token embedding visualization via t-SNE (a) and the impact of in-context exemplar quantity on PARG score (b).

demonstrating that explicit text externalization is essential for personalized attribute-reasoning generation.

Fig. 6 (b) analyzes the Intent Parser’s generalization through in-context learning. As exemplars increase from 0 to 6, generation score rises consistently from 0.30 to 0.62 (+106.7%), demonstrating that structured demonstrations enable the parser to accurately reformulate diverse user intents into explicit knowledge queries. This validates that in-context learning strengthens the concept→text→image pathway, enabling robust knowledge-grounded synthesis regardless of input phrasing.

### Multi-task Collaboration Yields Cross-task Gains.

During training, we observe that models trained solely on understanding and generation spontaneously acquire preliminary editing capabilities, which motivates us to investigate whether editing supervision could reciprocally enhance understanding and generation. Table 2 confirms this hypothesis: incorporating editing data improves recognition by 1.4% (0.840→0.852), QA GPT by 2.9% (0.659→0.678), and face similarity by 4.0% (0.397→0.413), alongside editing performance gains of 3.1% (0.638→0.658).

This bidirectional synergy stems from editing’s unique constraint structure, which requires simultaneous concept *localization* (understanding), identity *preservation* (gener-

ation), and instruction *alignment*. Unlike isolated tasks where the model can rely on visual pattern memorization, editing forces concept tokens to encode disentangled representations capturing both invariant identity attributes and modifiable contextual properties. These fine-grained constraints regularize identity preservation under distribution shift, transferring to all personalized generation scenarios. Additional details are provided in the Appendix.

## 6. Conclusion

We present **OmniPersona**, the first end-to-end framework unifying personalized understanding, generation, and image editing within unified multimodal models. Our framework introduces token representations decoupling to mitigate cross-task interference, explicit knowledge replay to enable personalized attribute-reasoning generation, and demonstrates understanding–generation–editing synergy through joint training strategies. We propose **OmniPBench**, the first benchmark for unified personalization evaluation, extending UnifyBench with personalized editing tasks and cross-task protocols. Experiments demonstrate competitive performance in personalized tasks. This work represents a significant step toward truly personal AI assistants capable of consistent, controllable concept modeling.

## Acknowledgements

This work has been supported in part by the NSFC (No. 62436007,62402426,62506333), the China Postdoctoral Science Foundation under Grant Number 2024M752794, the ZJNSF (No. LZ25F020004), the Key Research and Development Projects in Zhejiang Province (No. 2025C01128, 2025C02156), Ningbo Yongjiang Talent Introduction Programme (2023A400-G) and the Postdoctoral Fellowship Program of CPSF (GZC20251077).

## References

- [1] Yuval Alaluf, Elad Richardson, Sergey Tulyakov, Kfir Aberman, and Daniel Cohen-Or. Myvlm: Personalizing vlms for user-specific queries. In *European Conference on Computer Vision*, pages 73–91. Springer, 2024. 2, 3, 5
- [2] Ruichuan An, Sihan Yang, Ming Lu, Renrui Zhang, Kai Zeng, Yulin Luo, Jiajun Cao, Hao Liang, Ying Chen, Qi She, et al. Mc-llava: Multi-concept personalized vision-language model. *arXiv preprint arXiv:2411.11706*, 2024. 3
- [3] Ruichuan An, Sihan Yang, Renrui Zhang, Zijun Shen, Ming Lu, Gaole Dai, Hao Liang, Ziyu Guo, Shilin Yan, Yulin Luo, et al. Unictokens: Boosting personalized understanding and generation via unified concept tokens. *arXiv preprint arXiv:2505.14671*, 2025. 2, 3, 5
- [4] Jiu-hai Chen, Zhiyang Xu, Xichen Pan, Yushi Hu, Can Qin, Tom Goldstein, Lifu Huang, Tianyi Zhou, Saining Xie, Silvio Savarese, et al. Blip3-o: A family of fully open unified multimodal models-architecture, training and dataset. *arXiv preprint arXiv:2505.09568*, 2025. 2, 3
- [5] Xiaokang Chen, Zhiyu Wu, Xingchao Liu, Zizheng Pan, Wen Liu, Zhenda Xie, Xingkai Yu, and Chong Ruan. Janus-pro: Unified multimodal understanding and generation with data and model scaling. *arXiv preprint arXiv:2501.17811*, 2025. 2, 3
- [6] Chaorui Deng, Deyao Zhu, Kunchang Li, Chenhui Gou, Feng Li, Zeyu Wang, Shu Zhong, Weihao Yu, Xiaonan Nie, Ziang Song, et al. Emerging properties in unified multimodal pretraining. *arXiv preprint arXiv:2505.14683*, 2025. 2, 3, 5
- [7] Rinon Gal, Yuval Alaluf, Yuval Atzmon, Or Patashnik, Amit H Bermano, Gal Chechik, and Daniel Cohen-Or. An image is worth one word: Personalizing text-to-image generation using textual inversion. *arXiv preprint arXiv:2208.01618*, 2022. 5
- [8] Yuying Ge, Sijie Zhao, Jinguo Zhu, Yixiao Ge, Kun Yi, Lin Song, Chen Li, Xiaohan Ding, and Ying Shan. Seed-x: Multimodal models with unified multi-granularity comprehension and generation. *arXiv preprint arXiv:2404.14396*, 2024. 2, 3
- [9] Haoran Hao, Jiaming Han, Changsheng Li, Yu-Feng Li, and Xiangyu Yue. Rap: Retrieval-augmented personalization for multimodal large language models. In *Proceedings of the Computer Vision and Pattern Recognition Conference*, pages 14538–14548, 2025. 2, 3, 5
- [10] Zecheng He, Bo Sun, Felix Juefei-Xu, Haoyu Ma, Ankit Ramchandani, Vincent Cheung, Siddharth Shah, Anmol Kalia, Harihar Subramanyam, Alireza Zareian, et al. Imagine yourself: Tuning-free personalized image generation. *arXiv preprint arXiv:2409.13346*, 2024. 3
- [11] Yang Jiao, Haibo Qiu, Zequn Jie, Shaoxiang Chen, Jingjing Chen, Lin Ma, and Yu-Gang Jiang. Unitoken: Harmonizing multimodal understanding and generation through unified visual encoding. In *Proceedings of the Computer Vision and Pattern Recognition Conference*, pages 3600–3610, 2025. 2, 3
- [12] Tianwei Lin, Wenqiao Zhang, Sijing Li, Yuqian Yuan, Binhe Yu, Haoyuan Li, Wangui He, Hao Jiang, Mengze Li, Xiaohui Song, et al. Healthgpt: A medical large vision-language model for unifying comprehension and generation via heterogeneous knowledge adaptation. *arXiv preprint arXiv:2502.09838*, 2025. 2
- [13] Jiasen Lu, Christopher Clark, Sangho Lee, Zichen Zhang, Savya Khosla, Ryan Marten, Derek Hoiem, and Aniruddha Kembhavi. Unified-io 2: Scaling autoregressive multimodal models with vision language audio and action. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 26439–26455, 2024. 2, 3
- [14] Zheqi Lv, Wenqiao Zhang, Shengyu Zhang, Kun Kuang, Feng Wang, Yongwei Wang, Zhengyu Chen, Tao Shen, Hongxia Yang, Beng Chin Ooi, et al. Duet: A tuning-free device-cloud collaborative parameters generation framework for efficient device model generalization. In *Proceedings of the ACM Web Conference 2023*, pages 3077–3085, 2023. 3
- [15] Chuofan Ma, Yi Jiang, Junfeng Wu, Jihan Yang, Xin Yu, Zehuan Yuan, Bingyue Peng, and Xiaojuan Qi. Unitok: A unified tokenizer for visual generation and understanding. *arXiv preprint arXiv:2502.20321*, 2025. 3
- [16] Yiyang Ma, Xingchao Liu, Xiaokang Chen, Wen Liu, Chengyue Wu, Zhiyu Wu, Zizheng Pan, Zhenda Xie, Haowei Zhang, Liang Zhao, et al. Janusflow: Harmonizing autoregression and rectified flow for unified multimodal understanding and generation. *arXiv preprint arXiv:2411.07975*, 2024. 2, 3
- [17] Thao Nguyen, Haotian Liu, Yuheng Li, Mu Cai, Utkarsh Ojha, and Yong Jae Lee. Yo’llava: Your personalized language and vision assistant. *Advances in Neural Information Processing Systems*, 37:40913–40951, 2024. 2, 3, 5
- [18] Thao Nguyen, Krishna Kumar Singh, Jing Shi, Trung Bui, Yong Jae Lee, and Yuheng Li. Yo’chameleon: Personalized vision and language generation. In *Proceedings of the Computer Vision and Pattern Recognition Conference*, pages 14438–14448, 2025. 2, 3, 5
- [19] Kaihang Pan, Weile Chen, Haiyi Qiu, Qifan Yu, Wendong Bu, Zehan Wang, Yun Zhu, Juncheng Li, and Siliang Tang. Wiseddit: Benchmarking cognition-and creativity-informed image editing. *arXiv preprint arXiv:2512.00387*, 2025. 2
- [20] Xichen Pan, Satya Narayan Shukla, Aashu Singh, Zhuokai Zhao, Shlok Kumar Mishra, Jialiang Wang, Zhiyang Xu, Jiu-hai Chen, Kungpeng Li, Felix Juefei-Xu, et al. Transfer between modalities with metaqueries. *arXiv preprint arXiv:2504.06256*, 2025. 3
- [21] 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. 3, 5
- [22] Alireza Salemi, Sheshera Mysore, Michael Bendersky, and Hamed Zamani. Lamp: When large language models meet personalization. In *Proceedings of the 62nd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 7370–7392, 2024. 3
- [23] Jing Shi, Wei Xiong, Zhe Lin, and Hyun Joon Jung. Instantbooth: Personalized text-to-image generation without test-time finetuning. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 8543–8552, 2024. 3
- [24] Chameleon Team. Chameleon: Mixed-modal early-fusion foundation models. *arXiv preprint arXiv:2405.09818*, 2024. 2
- [25] Shengbang Tong, David Fan, Jiachen Li, Yunyang Xiong, Xinlei Chen, Koustuv Sinha, Michael Rabbat, Yann LeCun, Saining Xie, and Zhuang Liu. Metamorph: Multimodal understanding and generation via instruction tuning. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 17001–17012, 2025. 2
- [26] Xinlong Wang, Xiaosong Zhang, Zhengxiong Luo, Quan Sun, Yufeng Cui, Jinsheng Wang, Fan Zhang, Yueze Wang, Zhen Li, Qiying Yu, et al. Emu3: Next-token prediction is all you need. *arXiv preprint arXiv:2409.18869*, 2024. 3
- [27] Chengyue Wu, Xiaokang Chen, Zhiyu Wu, Yiyang Ma, Xingchao Liu, Zizheng Pan, Wen Liu, Zhenda Xie, Xingkai Yu, Chong Ruan, et al. Janus: Decoupling visual encoding for unified multimodal understanding and generation. *arXiv preprint arXiv:2410.13848*, 2024. 2
- [28] Chenfei Wu, Jiahao Li, Jingren Zhou, Junyang Lin, Kaiyuan Gao, Kun Yan, Sheng-ming Yin, Shuai Bai, Xiao Xu, Yilei Chen, et al. Qwen-image technical report. *arXiv preprint arXiv:2508.02324*, 2025. 3
- [29] Fangtai Wu, Mushui Liu, Weijie He, Zhao Wang, and Yunlong Yu. Dcoar: Deep concept injection into unified autoregressive models for personalized text-to-image generation. *arXiv preprint arXiv:2508.07341*, 2025. 3
- [30] Yecheng Wu, Zhuoyang Zhang, Junyu Chen, Haotian Tang, Dacheng Li, Yunhao Fang, Ligeng Zhu, Enze Xie, Hongxu Yin, Li Yi, et al. Vila-u: a unified foundation model integrating visual understanding and generation. *arXiv preprint arXiv:2409.04429*, 2024. 3
- [31] Jinheng Xie, Zhenheng Yang, and Mike Zheng Shou. Showo2: Improved native unified multimodal models. *arXiv preprint arXiv:2506.15564*, 2025. 2, 3
- [32] Hu Ye, Jun Zhang, Sibio Liu, Xiao Han, and Wei Yang. Ip-adapt: Text compatible image prompt adapter for text-to-image diffusion models. *arXiv preprint arXiv:2308.06721*, 2023. 3
- [33] Yuqian Yuan, Ronghao Dang, Long Li, Wentong Li, Dian Jiao, Xin Li, Deli Zhao, Fan Wang, Wenqiao Zhang, Jun Xiao, et al. Eoc-bench: Can mllms identify, recall, and forecast objects in an egocentric world? *arXiv preprint arXiv:2506.05287*, 2025. 3
- [34] Yuqian Yuan, Hang Zhang, Wentong Li, Zesen Cheng, Boqiang Zhang, Long Li, Xin Li, Deli Zhao, Wenqiao Zhang, Yueting Zhuang, et al. Videorefer suite: Advancing spatial-temporal object understanding with video llm. In *Proceedings of the Computer Vision and Pattern Recognition Conference*, pages 18970–18980, 2025.
- [35] Yuqian Yuan, Wenqiao Zhang, Xin Li, Shihao Wang, Kehan Li, Wentong Li, Jun Xiao, Lei Zhang, and Beng Chin Ooi. Pixelrefer: A unified framework for spatio-temporal object referring with arbitrary granularity. *arXiv preprint arXiv:2510.23603*, 2025.
- [36] Wenqiao Zhang, Haochen Shi, Siliang Tang, Jun Xiao, Qiang Yu, and Yueting Zhuang. Consensus graph representation learning for better grounded image captioning. In *Proceedings of the AAAI Conference on Artificial Intelligence*, pages 3394–3402, 2021. 3
- [37] Wenqiao Zhang, Lei Zhu, James Hallinan, Shengyu Zhang, Andrew Makmur, Qingpeng Cai, and Beng Chin Ooi. Boostmis: Boosting medical image semi-supervised learning with adaptive pseudo labeling and informative active annotation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 20666–20676, 2022. 3
- [38] Wenqiao Zhang, Tianwei Lin, Jiang Liu, Fangxun Shu, Haoyuan Li, Lei Zhang, He Wanggui, Hao Zhou, Zheqi Lv, Hao Jiang, et al. Hyperllava: Dynamic visual and language expert tuning for multimodal large language models. *arXiv preprint arXiv:2403.13447*, 2024. 3
- [39] Wenqiao Zhang, Zheqi Lv, Hao Zhou, Jia-Wei Liu, Juncheng Li, Mengze Li, Yunfei Li, Dongping Zhang, Yueting Zhuang, and Siliang Tang. Revisiting the domain shift and sample uncertainty in multi-source active domain transfer. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 16751–16761, 2024. 3
- [40] Xinjie Zhang, Jintao Guo, Shanshan Zhao, Minghao Fu, Lunhao Duan, Jiakui Hu, Yong Xien Chng, Guo-Hua Wang, Qing-Guo Chen, Zhao Xu, et al. Unified multimodal understanding and generation models: Advances, challenges, and opportunities. *arXiv preprint arXiv:2505.02567*, 2025. 2
- [41] Haoyu Zheng, Wenqiao Zhang, Zheqi Lv, Yu Zhong, Yang Dai, Jianxiang An, Yongliang Shen, Juncheng Li, Dongping Zhang, Siliang Tang, et al. Makima: Tuning-free multi-attribute open-domain video editing via mask-guided attention modulation. *arXiv preprint arXiv:2412.19978*, 2024. 2
- [42] Chunting Zhou, Lili Yu, Arun Babu, Kushal Tirumala, Michihiro Yasunaga, Leonid Shamis, Jacob Kahn, Xuezhe Ma, Luke Zettlemoyer, and Omer Levy. Transfusion: Predict the next token and diffuse images with one multi-modal model. *arXiv preprint arXiv:2408.11039*, 2024. 2, 3