

Unleashing the Potential of Large Language Models for Text-to-Image Generation Through Autoregressive Representation Alignment

Xing Xie^{1,2*}, Jiawei Liu^{1*}, Ziyue Lin³, Huijie Fan^{1†}, Zhi Han¹, Yandong Tang¹, Liangqiong Qu^{3†}

¹State Key Laboratory of Robotics and Intelligent Systems, Shenyang Institute of Automation, Chinese Academy of Sciences.

²University of Chinese Academy of Sciences.

³School of Computing and Data Science, The University of Hong Kong.

{xiexing, liujiawei, fanhuijie, hanzhi, ytang}@sia.cn, ziyue_lin@connect.hku.hk, liangqqu@hku.hk

Abstract

We present Autoregressive Representation Alignment (ARRA), a new training framework that unlocks global-coherent text-to-image generation in autoregressive LLMs without architectural modifications. Different from prior works that require complex architectural redesigns, ARRA aligns LLM’s hidden states with visual representations from external visual foundational models via a global visual alignment loss and a hybrid token, $\langle \text{HYBNEXT} \rangle$. This token enforces dual constraints: local next-token prediction and global semantic distillation, enabling LLMs to implicitly learn spatial and contextual coherence while retaining their original autoregressive paradigm. Extensive experiments validate ARRA’s plug-and-play versatility. When training T2I LLMs from scratch, ARRA reduces FID by 16.6% (ImageNet), 12.0% (LAION-COCO) for autoregressive LLMs like LlamaGen, without modifying original architecture and inference mechanism. For training from text-generation-only LLMs, ARRA reduces FID by 25.5% (MIMIC-CXR), 8.8% (DeepEyeNet) for advanced LLMs like Chameleon. For domain adaptation, ARRA aligns general-purpose LLMs with specialized models (e.g., BioMedCLIP), achieving an 18.6% FID reduction over direct fine-tuning on medical imaging (MIMIC-CXR). These results demonstrate that training objective redesign, rather than architectural modifications, can resolve cross-modal global coherence challenges. ARRA offers a complementary paradigm for advancing autoregressive models.

Code — <https://github.com/HKU-HealthAI/ARRA>

1 Introduction

Large language models (LLMs) (Achiam et al. 2023; Yang et al. 2025; Guo et al. 2025) and subsequent multimodal LLMs (MLLMs) (Team et al. 2023; Wang et al. 2024; Liu et al. 2024b; Zhu et al. 2025) have revolutionized the field of generative AI. These models, built on the autoregressive (AR) paradigm, show remarkable scalability and generalization capabilities in complex understanding tasks through a simple yet powerful next-token prediction framework.

*These authors contributed equally.

†Corresponding author.

Copyright © 2026, Association for the Advancement of Artificial Intelligence (www.aaai.org). All rights reserved.

Inspired by the success of LLMs, some researchers seek to replicate autoregressive “next-token prediction” paradigm for text-to-image generation. This paradigm is directly adopted in DALL-E (Ramesh et al. 2021), Parti (Yu et al. 2022), and LlamaGen (Sun et al. 2024) for image generation, where images are treated as sequences of discrete tokens. However, although “next-token prediction” paradigm excels in language tasks, where local dependencies naturally align with sequential structure, it struggles to bridge the significant cross-modal gap between language and images. As shown in Fig. 1(c), optimizing only for local next-token prediction forces the model to focus on isolated token-level features, neglecting the global coherence required for spatially structured visual content. This sometimes leads to fragmented parts on generated images, such as misaligned ribs in X-rays, where fine-grained details fail to harmonize into a unified whole. It can also cause semantic mismatches, as shown in Fig. 1(e), where global information is not maintained, leading to inconsistencies in the generated images.

Recognizing this limitation, recent efforts aim to inject global constraints into autoregressive frameworks to fully unlock the potential of LLMs in image generation (Zhou et al. 2024; Xie et al. 2024; Tian et al. 2024). Bidirectional attention mechanisms are introduced in Transfusion (Zhou et al. 2024) and Show-O (Xie et al. 2024) to model global image structure through patch diffusion and mask token modeling, respectively. These methods achieve promising results in generating high-quality images and demonstrate the potential of LLMs for multimodal generation. However, they rely on architectural modifications, such as cross-modal attention layers or grafted diffusion modules. While effective, such adaptations often deviate from standard LLM frameworks, limiting their compatibility with pretrained LLMs that excel under pure autoregressive paradigms. For instance, repurposing an off-the-shelf LLM for text-to-image generation would require retraining these modified components, losing benefits of existing scaling laws and generalization capabilities. This practical constraint raises a critical question: *Can we unlock the full potential of LLMs for text-to-image generation without altering the original architecture or inference mechanism?*

We address this by proposing Autoregressive Representation Alignment (ARRA), a novel training framework that redefines how LLMs learn text-to-image generation. Differ-

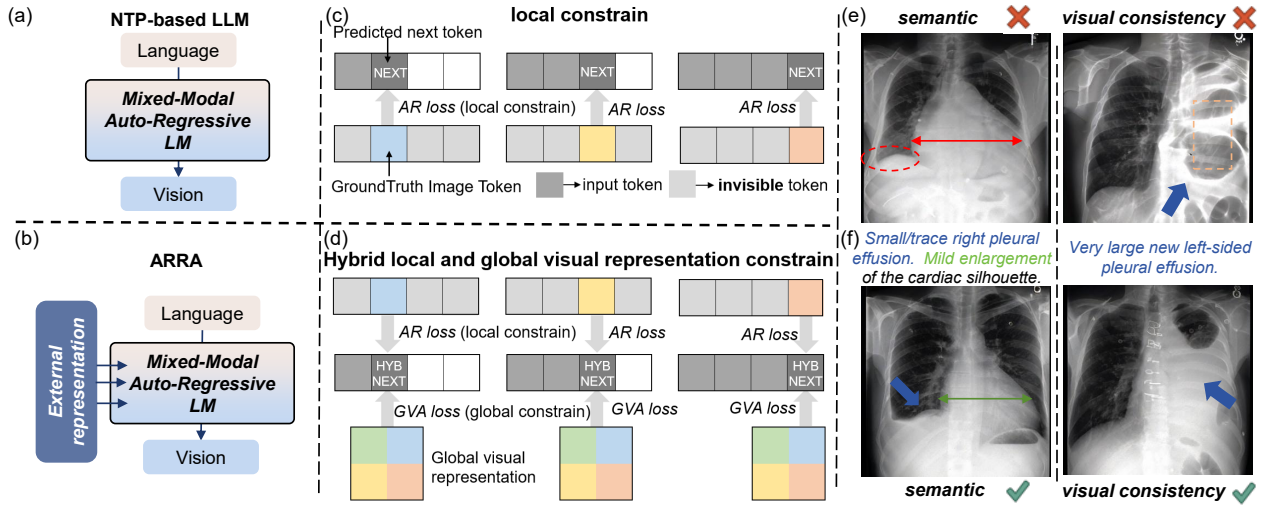


Figure 1: ARRA enables high-quality text-to-image generation through a redefined training objective that promotes global coherence. (a)(c) Traditional next-token prediction (NTP)-based LLMs rely solely on the autoregressive loss (AR loss) of the next token $\langle \text{NEXT} \rangle$ for local constraints. (b)(d) ARRA constructs a hybrid token $\langle \text{HYBNEXT} \rangle$, which is aligned by introducing external global visual representation, ensuring that $\langle \text{HYBNEXT} \rangle$ is constrained both locally by the AR loss and globally by the global visual alignment loss (GVA loss). (e)(f) ARRA demonstrates advantages in semantic consistency and visual continuity.

ent from prior works that modify architectures (e.g., adding attention layers or diffusion modules), ARRA preserves the original LLM framework while injecting global constraints directly into the training objective. Our key insight is simple: global coherence does not require architectural complexity; it can instead be achieved through a redefined training paradigm. Specifically, ARRA augments the standard autoregressive loss with a **global visual alignment loss** that aligns the LLM’s latent representations with semantic guidance from pretrained foundational models (Fig. 2). To bridge local and global learning, we introduce a hybrid token, $\langle \text{HYBNEXT} \rangle$, which serves as a bidirectional anchor. Locally, it predicts the next token via standard codebook indices. Globally, its latent embedding aligns with compressed visual features extracted from external models (e.g., BioMedCLIP (Zhang et al. 2023) or MedSAM (Ma et al. 2024)) via our novel global visual alignment loss. By distilling rich semantic features from external models into the $\langle \text{HYBNEXT} \rangle$ token during training, ARRA enables autoregressive sequences to implicitly learn global structure. More crucially, this alignment occurs only during training, leaving the LLM’s inference process untouched and preserving its inference-time efficiency.

Our experiments demonstrate the versatility of ARRA framework, achieving improvements without architectural modifications. ARRA supports three key capabilities on AR LLMs: (1) ARRA enhances training T2I LLMs from scratch. Applied to LlamaGen (Sun et al. 2024) with varying parameter scales, it consistently improves generation performance and exhibits strong scalability. (2) ARRA effectively transforms pretrained text-generation-only LLMs into T2I generators. When integrated into LLMs without image generation capabilities, such as Chameleon (Team 2024), ARRA yields steady improvements. (3) ARRA facilitates adaptation of general-purpose generative models to

special domains. By integrating domain-specific priors into LLMs with image generation capabilities, such as Lumina-mGPT (Liu et al. 2024a), ARRA substantially outperforms direct fine-tuning. These capabilities confirm the plug-and-play flexibility of ARRA framework.

The main contributions are summarized below:

- (i) We propose Autoregressive Representation Alignment, a novel training framework that redefines how LLMs learn text-to-image generation by decoupling global structure learning from model design. By aligning training objective with external representations, ARRA resolves local dependency limitations in LLMs while retaining original architectures and inference efficiency.
- (ii) We introduce the $\langle \text{HYBNEXT} \rangle$ token, a novel mechanism that bridges local next-token prediction with global semantic alignment via distillation from external models (e.g., BioMedCLIP or MedSAM), enabling implicit learning of spatial and contextual relationships.
- (iii) We provide a detailed experimental analysis, offering comprehensive insights into the selection of alignment tokens, aggregation strategy, and external representations. These findings serve as a practical guide for the effective utilization of representation alignment.
- (iv) ARRA offers plug-and-play flexibility, facilitating training T2I LLMs from scratch, transforming pretrained text-generation-only LLMs into T2I generators, and adapting general generative models to specific domains, all without architectural modifications. These capabilities are validated on both natural and medical image generation tasks using advanced AR models.

2 Related Work

2.1 Visual Generation Models

Diffusions. The success of diffusion models has revolutionized the image generation paradigm (Qu et al. 2015, 2019;

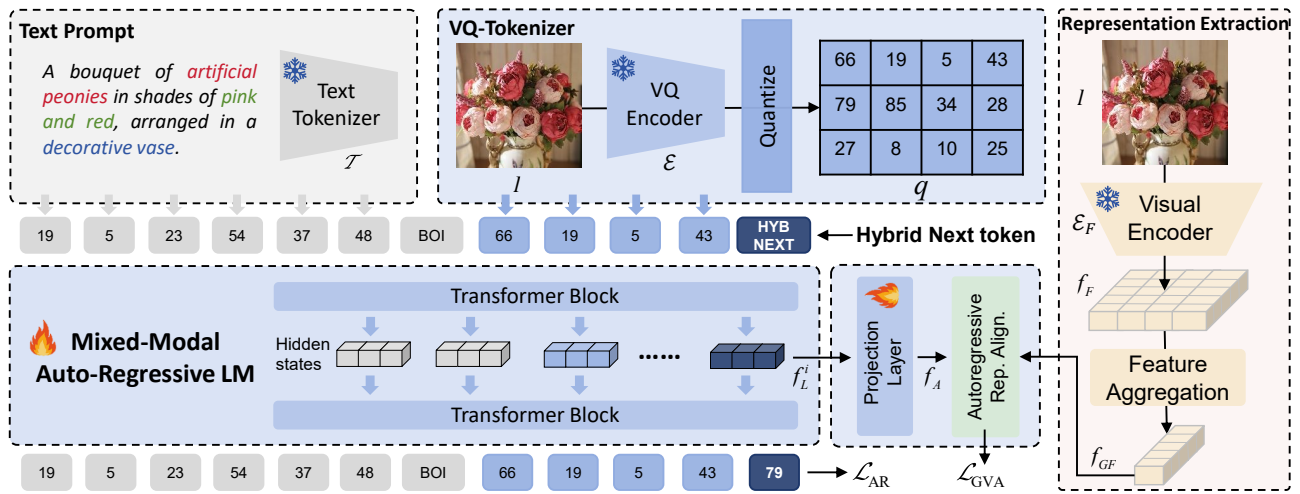


Figure 2: Proposed ARRA Framework. We define the next token predicted in the autoregressive sequence as the ‘‘HYBRID next token’’, denoted as $\langle \text{HYBNEXT} \rangle$. During training, $\langle \text{HYBNEXT} \rangle$ is constrained not only **locally** by the autoregressive loss \mathcal{L}_{AR} from the next token prediction and LLM codebook matching, but also **globally** through visual alignment loss \mathcal{L}_{GVA} , which modulates its hidden states using externally well-trained representations. We extract visual representations from a pretrained foundational model and further aggregate these features to obtain semantically enriched representations for alignment.

Rombach et al. 2022a; Liu et al. 2024c; Xie et al. 2025). DiTs (Peebles and Xie 2023) show strong scalability by replacing or integrating U-Net with Transformers. This inspires subsequent models such as SD3 (Esser et al. 2024), Imagen3 (Baldrige et al. 2024), which achieve new state-of-the-art performance in image generation. Recent work REPA (Yu et al. 2025) explores enhancing diffusion models with external representations. It utilizes patch-wise representations alignment, matching each patch-level hidden state of the diffusion transformer with corresponding patch tokens from an external encoder. This alignment improves generation performance without modifying architecture. However, REPA’s strategy is not directly compatible with AR models. AR models generate image tokens sequentially and do not produce all patch tokens simultaneously during training, making patch-wise alignment infeasible. Different from this, ARRA introduces a novel hybrid token to bridge local next-token prediction and global alignment, enabling effective integration of the alignment mechanism into AR architectures. For further discussion, please refer to Appendix 2.

Autoregressive models. Early pioneering works, VQ-VAE (Van Den Oord, Vinyals et al. 2017), VQ-GAN (Esser, Rombach, and Ommer 2021) and DALL-E (Ramesh et al. 2021), demonstrated the potential of AR models for image generation. Subsequent works such as RQ-Transformer (Lee et al. 2022) also follows a raster-scan manner and enhances image generation performance through extra scales or stacked codes. Recent works achieve performance comparable to diffusion models by employing scale modeling (Tian et al. 2024) and eliminating vector quantization (Li et al. 2024a; Fan et al. 2025). Additionally, masked prediction autoregressive models (Chang et al. 2022; Li et al. 2024b) employ BERT-like (Devlin et al. 2019) masked prediction modeling, improving generation efficiency and quality.

2.2 LLMs for Text-to-Image Generation

Currently, researchers pay attention to LLM-based text-to-image generation models, aiming to replicate the success of LLMs in language tasks. Early works leverage diffusion models as tools to extend LLMs (Dong et al. 2024; Ge et al. 2024). These works utilize LLMs as feature extractors to guide diffusion generators for visual generation. Such models are complex to design and do not fully unleash the generative potential of the LLM for visual generation. Recent works (Sun et al. 2024; Team 2024; Lu et al. 2024; Liu et al. 2024a) attempt to unify text and image modeling within a single LLM, enhancing generation performance through tokenizer optimization (Sun et al. 2024), early fusion modeling (Team 2024), and flexible resolution modeling (Liu et al. 2024a). These works discretize both text and images into tokens, which are then fed into the LLM for sequence modeling by next token prediction. However, the local constraints provided by next-token prediction struggle to bridge the cross-modal gap between language and images.

3 Proposed Method

We aim to achieve high-quality image generation without altering the next-token prediction paradigm of LLMs. We argue that the inherent cross-domain gaps in LLMs pose significant challenges for image generation tasks when the model lacks the ability to learn global features. To address this, we propose an Autoregressive Representation Alignment (ARRA) framework (see Fig. 2), which leverages the inherent representation capabilities of well-pretrained foundation models to facilitate the training of complex text-to-image autoregressive generation. Our framework is applied only during training, without affecting inference. It can enable the generation of high-quality images with exceptional semantic consistency in a cost-effective manner.

3.1 Overview

Our goal is to train an autoregressive model \mathcal{M}_θ leveraging representations derived from an external foundation visual encoder \mathcal{E}_F . \mathcal{M}_θ takes a text prompt T as input and generates a target image I . During training, T and I are first tokenized into token sequences s_T and s_I , respectively. These token sequences are then used to train the transformer-based autoregressive model \mathcal{M}_θ . Meanwhile, foundation visual encoder \mathcal{E}_F encodes I into the global visual representation f_{GF} , which is used to align with the feature f_A extracted from x_t by \mathcal{M}_θ . During image generation, the alignment module is removed, and image tokens are generated by \mathcal{M}_θ through next-token prediction. Finally, the output image tokens are decoded into pixel space by an image decoder to produce the target image I . We describe the autoregressive modeling process in Section 3.2 and detail our autoregressive representation alignment framework in Section 3.3.

3.2 Modeling via Next-Token Prediction

The autoregressive architecture comprises two core components: (1) A transformer-based autoregressive model \mathcal{M}_θ for probabilistic modeling of token sequences. (2) A VQ-based model (Esser, Rombach, and Ommer 2021) with encoder \mathcal{E} , quantizer \mathcal{Q} and decoder \mathcal{D} for transformation between image pixels and discrete token sequences. Formulation involves the following two parts:

Tokenization. In order to apply the next-token prediction modeling in the image domain, it is first required to convert the continuous 2D image pixels into discrete sequences. This process consists of two steps: (1) 2D image pixels to 2D image tokens, and (2) 2D image tokens to a 1D token sequences. Specifically, given a image $I \in \mathbb{R}^{H \times W \times 3}$, we first obtain the image feature map $f = \mathcal{E}(I) \in \mathbb{R}^{h \times w \times d}$, where $h = H/c$, $w = W/c$, d is the dimension of the codes, c denotes a compression factor. Subsequently, we convert f into discrete tokens by $q = \mathcal{Q}(f) \in \mathcal{Z}^{h \times w}$, where the quantizer $\mathcal{Q}(\cdot)$ maps each vector $f^{(i,j)}$ in the image feature map f to the code index $q^{(i,j)}$ of its nearest vector $z^{(i,j)}$ in the codebook \mathcal{Z} . The image tokens q are then reshaped into a 1D token sequence $s_I = \{x_1^I, x_2^I, x_3^I, \dots, x_n^I\}$ with a length of $h \cdot w$, arranged according to the raster scan order.

For text prompt T , we obtain the discrete sequence through $s_T = \mathcal{T}(T) = \{x_1^T, x_2^T, x_3^T, \dots, x_n^T\}$, where $\mathcal{T}(\cdot)$ denotes a text tokenizer.

Next token prediction modeling. We combine text sequences s_T and image sequences s_I to obtain discrete tokens $x = \{x_1, x_2, x_3, \dots, x_n\}$, where x_n is an integer from a tokenizer’s vocabulary V . The next-token prediction paradigm posits the probability of current token x_t depends only on its prefix $(x_1, x_2, x_3, \dots, x_{t-1})$. The likelihood of sequence modeling can be expressed as:

$$p(x) = \prod_{t=1}^n p(x_t | x_1, x_2, \dots, x_{t-1}). \quad (1)$$

The autoregressive model \mathcal{M}_θ formulates the generative task as predicting the distribution of the next token and optimizes the likelihood $p_\theta(x)$ through cross-entropy (CE) loss:

$$\mathcal{L}_{AR}(\theta) = \mathbb{E}_{x_t} [-\log p_\theta(x_t | x_{<t})]. \quad (2)$$

During training, the autoregressive model \mathcal{M}_θ relies on the previous tokens $x_{<t}$ to predict the next token x_t , where the hidden state of x_t in the i -th layer of \mathcal{M}_θ is denoted by f_L^i . In the last layer, the hidden state f_L^{-1} is then passed through the LLM head to compute the probability distribution p_θ for x_t . **In the original autoregressive model, p_θ is constrained only by the local context of a single token (i.e., x_t), lacking the ability to capture global information.** This limitation restricts model’s capacity to capture complex cross-modal relationships. To address this problem, we propose Autoregressive Representation Alignment.

3.3 Autoregressive Representation Alignment

We align the visual representations extracted from the pre-trained foundational model with LLM’s hidden states and investigate impact of different alignment strategies. The goal of alignment is to enable the hidden state of autoregressive transformer to acquire external global representations, providing meaningful guidance for reconstructing image.

Pre-trained Visual Representation Extraction. Let \mathcal{E}_F be a pretrained foundation model’s visual encoder and I be a target image. We encode I as a visual representation by $f_F = \mathcal{E}_F(I) \in \mathbb{R}^{N \times D}$, where N, D denotes the embedding length and dimension of f_F . f_F is aggregated to the global visual representation f_{GF} , i.e., $f_{GF} = \text{agg}(f_F) \in \mathbb{R}^{1 \times D}$, where $\text{agg}(\cdot)$ denotes a feature aggregation operation. This aggregation operation fully extracts global information in the features and facilitates alignment with the hidden states of autoregressive model. For CLIP series, inspired by (Raghu et al. 2021), we use the $\langle \text{CLS} \rangle$ token representation from the Transformer-based visual encoder as global visual representation. For SAM series, which lack a $\langle \text{CLS} \rangle$ token, we instead apply average pooling over all patch features for feature aggregation operation.

Hybrid Next Token. We define the next token predicted by LLM sequence in our framework as “*HYBRID next token*”, denoted as $\langle \text{HYBNEXT} \rangle$. Unlike a “locally constrained token” in previous autoregressive models that are solely constrained by the LLM codebook, our $\langle \text{HYBNEXT} \rangle$ can fully incorporate external, well-trained global visual representations, making it a “globally and locally constrained token”.

Global Visual Representation Alignment. We obtain the hidden state f_L^i of $\langle \text{HYBNEXT} \rangle$ from the autoregressive model \mathcal{M}_θ . The hidden state f_L^i is converted to $f_A \in \mathbb{R}^{1 \times D}$ by a projection layer \mathcal{A}_ϕ to align with the global visual representation $f_{GF} \in \mathbb{R}^{1 \times D}$, i.e., $f_A = \mathcal{A}_\phi(f_L^i)$, where \mathcal{A}_ϕ is a two-layer MLP. Representation alignment is achieved through a **Global Visual Alignment loss** \mathcal{L}_{GVA} , which maximizes the similarity between the projected feature f_A and the global visual representation f_{GF} :

$$\mathcal{L}_{GVA}(\theta, \phi) = \text{sim}(f_A, f_{GF}). \quad (3)$$

where $\text{sim}(\cdot, \cdot)$ denotes cosine similarity loss. This alignment enables the $\langle \text{HYBNEXT} \rangle$ token to learn global visual representation, bridging the cross-modal gap and making the token prediction process more reliable.

Therefore, the autoregressive model can be jointly optimized through the following composite loss function:

$$\mathcal{L}_{ARRA}(\theta, \phi) = \mathcal{L}_{AR}(\theta) + \lambda \mathcal{L}_{GVA}(\theta, \phi). \quad (4)$$

λ serves as a balancing hyperparameter that controls the relative importance of the alignment objective. Experimentally, we set $\lambda = 1$.

3.4 Versatile ARRA for Diverse Scenarios

Our ARRA framework is flexible and plug-and-play, supporting different training scenarios and LLM frameworks. Therefore, we provide three representative model variants:

(1) **ARRA-Base**. It trains an LLM from scratch with random initialization, supporting settings where no pretrained models are available and showcasing ARRA’s ability to learn multimodal alignment from the ground up.

(2) **ARRA**. It initializes with a pretrained LLM that has strong text generation capabilities, enabling efficient extension to text-to-image tasks with text-generation-only LLMs.

(3) **ARRA-Adapt**. It builds on pretrained LLMs with both text and image generation capabilities, allowing adaptation to specialized domains such as medical imaging by leveraging special-domain priors.

4 Experimental Analysis and Results

We first perform a comprehensive component analysis of the ARRA framework to investigate how different design choices, such as alignment mechanism, feature aggregation strategy, and external encoder selection, impact alignment performance. Based on analysis, we finalize the ARRA framework and compare the three model variants, ARRA-Base, ARRA, and ARRA-Adapt, with advanced generative models to evaluate their performance and adaptability.

4.1 Experimental Setup

Datasets. We evaluate our model on both natural and medical image datasets. For natural images, we conduct evaluations on the ImageNet (Deng et al. 2009) dataset and a 2.4M high-quality subset of LAION-COCO (Schuhmann et al. 2022). For medical imaging, we use the MIMIC-CXR (Johnson et al. 2019) and DeepEyeNet (Huang et al. 2021) datasets. Detailed preprocessing is provided in Appendix 3. **Implementation Details.** We adopt LlamaGen (Sun et al. 2024), Chameleon 7B (Team 2024) and Lumina-mGPT 7B (Liu et al. 2024a) as respective baseline models for the ARRA-Base, ARRA, and ARRA-Adapt model variants, respectively, to verify the effectiveness of our framework.

4.2 Component Analysis of Alignment

To better understand the core design principles of ARRA and build an effective model, we analyze how different alignment component designs affect the effectiveness of representation alignment in autoregressive generation. We focus on the following key questions:

- Alignment mechanism: Does token-level alignment (our proposed `<HYBNEXT>`) outperform fixed-position alignment (`<REP>`) in preserving global visual constraints during autoregressive generation? (Table 1)
- Feature aggregation strategy: How do different types of features extracted from the same visual encoder affect the generation performance? (Table 2)

- Visual encoder selection on alignment: How does external encoder selection (general vs. specialized and cross-modal vs. vision-only) impact generation? (Table 3)

Target token	FID ↓	MS-SSIM ↑	CLIP-Score ↑
<code><REP></code>	4.85	0.410	0.4527
<code><HYBNEXT></code>	4.15	0.422	0.4576
w/o align.	5.10	0.401	0.4518
<code><REP></code>	6.26	0.401	0.4499
<code><HYBNEXT></code>	5.30	0.405	0.4532
w/o align.	7.11	0.383	0.4460

Table 1: Impact of different alignment mechanism selection on MIMIC-CXR. Gray represents ARRA-adapt model, and white represents ARRA model.

Alignment mechanism. We compare two strategies for integrating visual representations: (1) aligning features to a fixed `<REP>` token at the start of the generated sequence, and (2) aligning to the hidden state of `<HYBNEXT>`, a hybrid token interleaved at every generation step. As shown in Table 1, `<HYBNEXT>` yields superior performance. We argue that `<HYBNEXT>` allows for comprehensive traversal of every token during training sampling, ensuring that each token is effectively constrained by external global representations. In contrast, `<REP>` suffers from the “attention sink” (Liu et al. 2024a; Xiao et al. 2023) effect, where attention to the fixed token decays over long sequences, leading to degraded outputs. This leads to a key insight: **Takeaway 1.** *Aligning visual representations to a hybrid token interleaved at each generation step `<HYBNEXT>` is more effective than using a fixed token `<REP>`, as it prevents attention decay and ensures consistent constraint by external representations.*

CLS	Avgpool	FID ↓	MS-SSIM ↑	CLIP-Score ↑
✓		5.30	0.405	0.4532
	✓	6.56	0.387	0.4434
✓	✓	5.93	0.385	0.4465

Table 2: Impact of using different feature aggregation strategies on MIMIC-CXR with ARRA.

Feature aggregation strategy. To investigate how different types of features extracted from the visual encoder affect the generation performance, we understand three strategies for aggregating features from a single visual encoder: (1) the `[CLS]` token representation, (2) average pooling of all image patch representations, and (3) a concatenation of both representations. As shown in Table 2, the `[CLS]` token representation yields optimal performance. We attribute this superiority to the `[CLS]` token’s ability to aggregate global visual information through self-attention mechanisms. This ability provides a compact yet comprehensive representation for cross-modal alignment. This finding consistent with (Raghu et al. 2021), which establishes that `[CLS]` tokens capture global representations while patch tokens focus on local features. This global-local distinction suggests that autoregressive image generation models particularly benefit

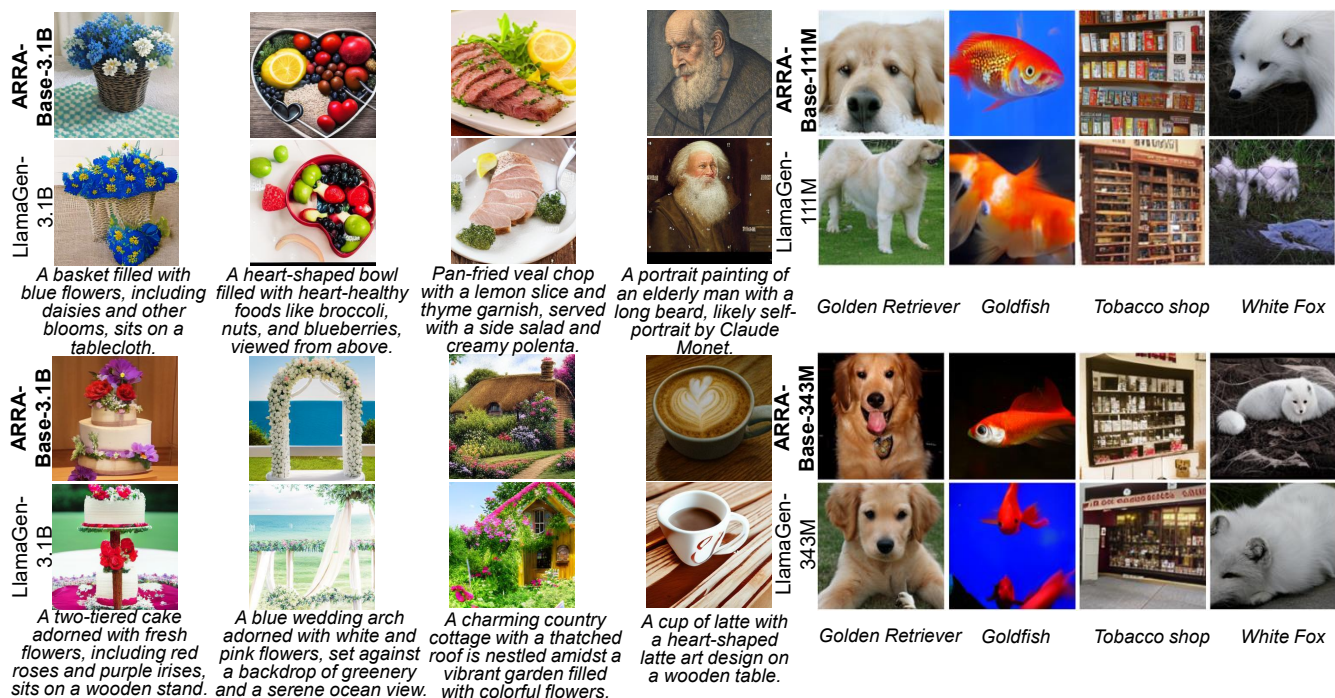


Figure 3: ARRA-Base improves the generation of LlamaGen. Left are text-conditional image generation results on the LAION-COCO dataset, and right are category-conditional image generation results on the Imagenet dataset.

from external global semantic representations as guidance. These results lead to a key insight: **Takeaway 2.** *The [CLS] token representation in foundation models effectively aggregates global visual information, providing comprehensive guidance for cross-modal alignment.*

Target Rep.	FID ↓	MS-SSIM ↑	CLIP-Score ↑
BioMedCLIP (Zhang et al. 2023)	4.15	0.422	0.4576
Med-SAM (Ma et al. 2024)	4.08	0.398	0.4542
CLIP-L (Radford et al. 2021a)	4.63	0.407	0.4519
w/o align.	5.10	0.401	0.4518
BioMedCLIP (Zhang et al. 2023)	5.30	0.405	0.4532
Med-SAM (Ma et al. 2024)	6.54	0.384	0.4465
CLIP-L (Radford et al. 2021a)	5.16	0.394	0.4450
w/o align.	7.11	0.383	0.4460

Table 3: Impact of alignment with representation extracted from different encoders on MIMIC-CXR.

Visual encoder selection on alignment. To evaluate how visual encoder selection impacts cross-modal alignment and generation quality, we conduct experiments with three encoders: BioMedCLIP (Zhang et al. 2023) (domain-specific cross-modal encoders), CLIP (Radford et al. 2021b) (general-purpose cross-modal encoders), and Med-SAM (Ma et al. 2024) (domain-specific pure visual encoders). As shown in Table 3, all pretrained encoders improve generation performance compared to without alignment. This improvement arises from global constraints imposed by external representations, which regularize the autoregressive generation process. Notably, for training from a text-generation-only LLM (ARRA), BioMedCLIP and CLIP

demonstrate superior performance. Their cross-modal training helps bridge the gap between text and image modalities, enabling LLMs to learn what to generate (semantics) before how to generate (pixel details). Conversely, when fine-tuning pretrained LLMs with image-generation capabilities (ARRA-Adapt), domain-specific encoders like BioMedCLIP and MedSAM dominate. BioMedCLIP injects medical-specific semantics, while MedSAM provides structural priors (e.g., organ shapes) through its segmentation-focused features. These results lead to a key insight: **Takeaway 3.** *When the LLM lacks image generation capabilities, cross-modal encoders are crucial for semantic grounding. However, for LLMs with image generation capabilities, domain-specific encoders are more effective, as they provide fine-grained features needed for domain-specific adaptation.*

4.3 Main Comparison Results

Based on analysis, we finalize ARRA framework with CLIP and BioMedCLIP for alignment on natural and medical datasets, respectively. We use [HYBNEXT] and [CLS] as our alignment strategy. We compare ARRA framework under the three proposed model variants (ARRA-Base, ARRA, ARRA-Adapt) to demonstrate its superiority and versatility. **Facilitating training T2I LLMs from scratch: performance comparison of ARRA-Base.** We evaluate the generation performance and model scalability of ARRA-Base variant, which is trained from scratch, on large-scale datasets. We first conduct experiments on ImageNet, using LlamaGen (Sun et al. 2024) with 111M and 343M parameters as baselines. We then train models on a 2.4M

ImageNet				LAION-COCO					
Model	111M		343M		Model	343M	775M	1.4B	3.1B
	FID ↓	IS ↑	FID ↓	IS ↑		(FID, CLIP-Score)	(FID, CLIP-Score)	(FID, CLIP-Score)	(FID, CLIP-Score)
LlamaGen	5.47	193.6	4.33	286.6	LlamaGen	(12.78, 0.2470)	(11.79, 0.2512)	(11.57, 0.2510)	(11.88, 0.2605)
ARRA-Base	5.08	201.2	3.61	262.6	ARRA-Base	(11.67, 0.2500)	(10.81, 0.2571)	(10.50, 0.2589)	(10.45, 0.2615)

Table 4: Quantitative performance comparison of ARRA-Base and LlamaGen models with different scales on ImageNet datasets and 2.4M subset of LAION-COCO datasets, respectively.

Method	Chest-Xray: MIMIC-CXR				Fundus: DeepEyeNet		
	FID ↓	IS ↑	MS-SSIM ↑	CLIP-Score ↑	FID ↓	MS-SSIM ↑	CLIP-Score ↑
Original SD V2-I (Rombach et al. 2022b)	71.68	2.365	0.128	0.2333	166.45	0.141	0.2164
DreamBooth SD (Ruiz et al. 2023)	60.40	2.269	0.270	0.3693	80.70	0.352	0.3061
MINIM (Wang et al. 2025)	15.62	2.468	0.317	0.4423	59.71	0.333	0.2862
UniXGen (Lee et al. 2024a)	30.75	2.437	0.361	0.4128	-	-	-
LLM-CXR (Lee et al. 2024b)	5.88	2.134	0.395	0.4374	-	-	-
Chameleon (Team 2024)	7.11	2.498	0.383	0.4460	38.37	0.341	0.3234
ARRA	5.30	2.587	0.405	0.4532	35.01	0.376	0.3354
ARRA-Adapt	4.15	2.746	0.422	0.4576	34.70	0.392	0.3389

Table 5: Compare different methods with ARRA and ARRA-Adapt on MIMIC-CXR and DeepEyeNet datasets, respectively.

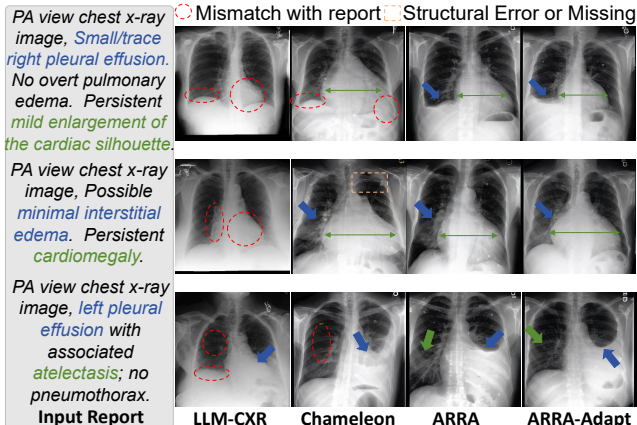


Figure 4: Visual comparison of different methods with ARRA and ARRA-Adapt on the MIMIC-CXR dataset.

high-quality subset of LAION-COCO and employ LlamaGen variants of increasing sizes (343M to 3.1B parameters) as baselines. As shown in Table 4 and Fig. 3, ARRA-Base outperforms the baseline model across all scales. The FID decreased by 16.6% and 12.0% on the ImageNet and LAION-COCO datasets, respectively. Meanwhile, as the model size increased, the performance remained steadily improved (FID decreased from 11.67 to 10.45), demonstrating ARRA maintains strong scalability of autoregressive model. These results lead to a key insight: **Takeaway 4. ARRA-Base enables efficient training of T2I LLMs from scratch, while preserving strong model scalability.**

Boosting pretrained LLMs for T2I generation and domain adaptation: performance of ARRA and ARRA-Adapt. We evaluate ARRA and ARRA-Adapt, initialized with text-generation-only and general image generation pretrained LLMs, respectively, on the MIMIC-CXR and Deep-

EyeNet datasets. Both variants are benchmarked against state-of-the-art diffusion models (Stable Diffusion (Rombach et al. 2022b), DreamBooth (Ruiz et al. 2023), MINIM (Wang et al. 2025)) and autoregressive models (LLM-CXR (Lee et al. 2024b), UniXGen (Lee et al. 2024a), Chameleon (Team 2024)). As shown in Table 5, our model achieves superior performance in both visual quality and semantic alignment. Compared to the chameleon, ARRA achieves a 25.5% and 8.8% reduction in FID on the MIMIC-CXR and DeepEyeNet datasets, respectively. On the MIMIC-CXR dataset, ARRA surpasses the best-performing baseline, LLM-CXR, reducing the FID from 5.88 to 5.30 and increasing the CLIP-Score from 0.4374 to 0.4532. The ARRA-Adapt variant delivers even stronger results, achieving an FID of 4.15 and a CLIP-Score of 0.4576. In addition, as shown in Fig. 4, our models demonstrate superior alignment with fine-grained clinical details, such as lesion location and severity. These results lead to a key insight: **Takeaway 5. ARRA enables a more effective transformation of text-generated-only LLMs into T2I generators, while ARRA-Adapt substantially improves domain adaptation and aligns general image-generation LLMs with specialized fields more effectively, both outperforming baseline approaches.**

5 Conclusion

We propose Autoregressive Representation Alignment (ARRA), a framework that enhances autoregressive image generation by injecting external visual representations during training. This approach enriches global semantic understanding while maintaining the model’s original autoregressive paradigm during generation. Experiments on natural and medical image generation tasks demonstrate ARRA’s versatility, offering a cost-effective framework for training autoregressive text-to-image generation models. Our work bridges the gap between multimodal domains and provides novel insights into modeling unified multimodal generation.

Acknowledgments

This work was supported by the National Natural Science Foundation of China under Grant (62306253, 61873259, T2596040, T2596045), the National Key Research and Development Program of China under Grant 2024YFB4707700, the Early Career Fund (27207025), the National Natural Science Foundation of China under Grant U23A20343, the Guangdong Natural Science Fund-General Program (2024A1515010233), the China Postdoctoral Science Foundation under Grant Number 2025M781669, and the Fundamental Research Project of SIA (2025JC1K05).

References

- Achiam, J.; Adler, S.; Agarwal, S.; Ahmad, L.; Akkaya, I.; Aleman, F. L.; Almeida, D.; Altschmidt, J.; Altman, S.; Anadkat, S.; et al. 2023. Gpt-4 technical report. *arXiv preprint arXiv:2303.08774*.
- Baldrige, J.; Bauer, J.; Bhutani, M.; Brichtova, N.; Bunner, A.; Chan, K.; Chen, Y.; Dieleman, S.; Du, Y.; Eaton-Rosen, Z.; et al. 2024. Imagen 3. *arXiv preprint arXiv:2408.07009*.
- Chang, H.; Zhang, H.; Jiang, L.; Liu, C.; and Freeman, W. T. 2022. Maskgit: Masked generative image transformer. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 11315–11325.
- Deng, J.; Dong, W.; Socher, R.; Li, L.-J.; Li, K.; and Fei-Fei, L. 2009. ImageNet: A large-scale hierarchical image database. In *2009 IEEE Conference on Computer Vision and Pattern Recognition*, 248–255.
- Devlin, J.; Chang, M.-W.; Lee, K.; and Toutanova, K. 2019. Bert: Pre-training of deep bidirectional transformers for language understanding. In *Proceedings of the 2019 conference of the North American chapter of the association for computational linguistics: human language technologies, volume 1 (long and short papers)*, 4171–4186.
- Dong, R.; Han, C.; Peng, Y.; Qi, Z.; Ge, Z.; Yang, J.; Zhao, L.; Sun, J.; Zhou, H.; Wei, H.; Kong, X.; Zhang, X.; Ma, K.; and Yi, L. 2024. DreamLLM: Synergistic Multimodal Comprehension and Creation. In *The Twelfth International Conference on Learning Representations*.
- Esser, P.; Kulal, S.; Blattmann, A.; Entezari, R.; Müller, J.; Saini, H.; Levi, Y.; Lorenz, D.; Sauer, A.; Boesel, F.; Podell, D.; Dockhorn, T.; English, Z.; Lacey, K.; Goodwin, A.; Marek, Y.; and Rombach, R. 2024. Scaling Rectified Flow Transformers for High-Resolution Image Synthesis. *arXiv:2403.03206*.
- Esser, P.; Rombach, R.; and Ommer, B. 2021. Taming transformers for high-resolution image synthesis. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 12873–12883.
- Fan, L.; Li, T.; Qin, S.; Li, Y.; Sun, C.; Rubinstein, M.; Sun, D.; He, K.; and Tian, Y. 2025. Fluid: Scaling autoregressive text-to-image generative models with continuous tokens.
- Ge, Y.; Zhao, S.; Zhu, J.; Ge, Y.; Yi, K.; Song, L.; Li, C.; Ding, X.; and Shan, Y. 2024. Seed-x: Multimodal models with unified multi-granularity comprehension and generation. *arXiv preprint arXiv:2404.14396*.
- Guo, D.; Yang, D.; Zhang, H.; Song, J.; Zhang, R.; Xu, R.; Zhu, Q.; Ma, S.; Wang, P.; Bi, X.; et al. 2025. Deepseek-r1: Incentivizing reasoning capability in llms via reinforcement learning. *arXiv preprint arXiv:2501.12948*.
- Huang, J.-H.; Yang, C.-H. H.; Liu, F.; Tian, M.; Liu, Y.-C.; Wu, T.-W.; Lin, I.-H.; Wang, K.; Morikawa, H.; Chang, H.; Tegner, J.; and Worring, M. 2021. DeepOpht: Medical Report Generation for Retinal Images via Deep Models and Visual Explanation. In *Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision (WACV)*, 2442–2452.
- Johnson, A. E.; Pollard, T. J.; Greenbaum, N. R.; Lungren, M. P.; Deng, C.-y.; Peng, Y.; Lu, Z.; Mark, R. G.; Berkowitz, S. J.; and Horng, S. 2019. MIMIC-CXR-JPG, a large publicly available database of labeled chest radiographs. *arXiv preprint arXiv:1901.07042*.
- Lee, D.; Kim, C.; Kim, S.; Cho, M.; and Han, W.-S. 2022. Autoregressive image generation using residual quantization. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 11523–11532.
- Lee, H.; Kim, W.; Kim, J.-H.; Kim, T.; Kim, J.; Sunwoo, L.; Choi, E.; et al. 2024a. Vision-Language Generative Model for View-Specific Chest X-ray Generation. In *Conference on Health, Inference, and Learning*, 280–296. PMLR.
- Lee, S.; Kim, W. J.; Chang, J.; and Ye, J. C. 2024b. LLM-CXR: instruction-finetuned LLM for CXR image understanding and generation. In *12th International Conference on Learning Representations, ICLR 2024*.
- Li, T.; Tian, Y.; Li, H.; Deng, M.; and He, K. 2024a. Autoregressive Image Generation without Vector Quantization. *arXiv preprint arXiv:2406.11838*.
- Li, T.; Tian, Y.; Li, H.; Deng, M.; and He, K. 2024b. Autoregressive Image Generation without Vector Quantization. *arXiv:2406.11838*.
- Liu, D.; Zhao, S.; Zhuo, L.; Lin, W.; Qiao, Y.; Li, H.; and Gao, P. 2024a. Lumina-mgpt: Illuminate flexible photorealistic text-to-image generation with multimodal generative pretraining. *arXiv preprint arXiv:2408.02657*.
- Liu, H.; Li, C.; Li, Y.; and Lee, Y. J. 2024b. Improved baselines with visual instruction tuning. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 26296–26306.
- Liu, J.; Wang, Q.; Fan, H.; Wang, Y.; Tang, Y.; and Qu, L. 2024c. Residual denoising diffusion models. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2773–2783.
- Lu, J.; Clark, C.; Lee, S.; Zhang, Z.; Khosla, S.; Marten, R.; Hoiem, D.; and Kembhavi, A. 2024. 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*, 26439–26455.
- Ma, J.; He, Y.; Li, F.; Han, L.; You, C.; and Wang, B. 2024. Segment Anything in Medical Images. *Nature Communications*, 15: 654.

- Peebles, W.; and Xie, S. 2023. Scalable diffusion models with transformers. In *Proceedings of the IEEE/CVF international conference on computer vision*, 4195–4205.
- Qu, L.; Tian, J.; Han, Z.; and Tang, Y. 2015. Pixel-wise orthogonal decomposition for color illumination invariant and shadow-free image. *Optics express*, 23(3): 2220–2239.
- Qu, L.; Wang, S.; Yap, P.-T.; and Shen, D. 2019. Wavelet-based semi-supervised adversarial learning for synthesizing realistic 7T from 3T MRI. In *International conference on medical image computing and computer-assisted intervention*, 786–794. Springer.
- Radford, A.; Kim, J. W.; Hallacy, C.; Ramesh, A.; Goh, G.; Agarwal, S.; Sastry, G.; Askell, A.; Mishkin, P.; Clark, J.; Krueger, G.; and Sutskever, I. 2021a. Learning Transferable Visual Models From Natural Language Supervision. arXiv:2103.00020.
- Radford, A.; Kim, J. W.; Hallacy, C.; Ramesh, A.; Goh, G.; Agarwal, S.; Sastry, G.; Askell, A.; Mishkin, P.; Clark, J.; et al. 2021b. Learning transferable visual models from natural language supervision. In *International conference on machine learning*, 8748–8763. PmLR.
- Raghu, M.; Unterthiner, T.; Kornblith, S.; Zhang, C.; and Dosovitskiy, A. 2021. Do vision transformers see like convolutional neural networks? *Advances in neural information processing systems*, 34: 12116–12128.
- Ramesh, A.; Pavlov, M.; Goh, G.; Gray, S.; Voss, C.; Radford, A.; Chen, M.; and Sutskever, I. 2021. Zero-shot text-to-image generation. In *International conference on machine learning*, 8821–8831. Pmlr.
- Rombach, R.; Blattmann, A.; Lorenz, D.; Esser, P.; and Ommer, B. 2022a. High-Resolution Image Synthesis with Latent Diffusion Models. arXiv:2112.10752.
- Rombach, R.; Blattmann, A.; Lorenz, D.; Esser, P.; and Ommer, B. 2022b. High-resolution image synthesis with latent diffusion models. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 10684–10695.
- Ruiz, N.; Li, Y.; Jampani, V.; Pritch, Y.; Rubinstein, M.; and Aberman, K. 2023. Dreambooth: Fine tuning text-to-image diffusion models for subject-driven generation. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 22500–22510.
- Schuhmann, C.; Beaumont, R.; Vencu, R.; Gordon, C.; Wightman, R.; Cherti, M.; Coombes, T.; Katta, A.; Mullis, C.; Wortsman, M.; et al. 2022. Laion-5b: An open large-scale dataset for training next generation image-text models. *Advances in neural information processing systems*, 35: 25278–25294.
- Sun, P.; Jiang, Y.; Chen, S.; Zhang, S.; Peng, B.; Luo, P.; and Yuan, Z. 2024. Autoregressive Model Beats Diffusion: Llama for Scalable Image Generation. arXiv:2406.06525.
- Team, C. 2024. Chameleon: Mixed-modal early-fusion foundation models. arXiv preprint arXiv:2405.09818.
- Team, G.; Anil, R.; Borgeaud, S.; Alayrac, J.-B.; Yu, J.; Soricut, R.; Schalkwyk, J.; Dai, A. M.; Hauth, A.; Millican, K.; et al. 2023. Gemini: a family of highly capable multimodal models. arXiv preprint arXiv:2312.11805.
- Tian, K.; Jiang, Y.; Yuan, Z.; Peng, B.; and Wang, L. 2024. Visual autoregressive modeling: Scalable image generation via next-scale prediction. arXiv preprint arXiv:2404.02905.
- Van Den Oord, A.; Vinyals, O.; et al. 2017. Neural discrete representation learning. *Advances in neural information processing systems*, 30.
- Wang, J.; Wang, K.; Yu, Y.; Lu, Y.; Xiao, W.; Sun, Z.; Liu, F.; Zou, Z.; Gao, Y.; Yang, L.; et al. 2025. Self-improving generative foundation model for synthetic medical image generation and clinical applications. *Nature Medicine*, 31(2): 609–617.
- Wang, X.; Zhang, X.; Luo, Z.; Sun, Q.; Cui, Y.; Wang, J.; Zhang, F.; Wang, Y.; Li, Z.; Yu, Q.; et al. 2024. Emu3: Next-token prediction is all you need. arXiv preprint arXiv:2409.18869.
- Xiao, G.; Tian, Y.; Chen, B.; Han, S.; and Lewis, M. 2023. Efficient streaming language models with attention sinks. arXiv preprint arXiv:2309.17453.
- Xie, J.; Mao, W.; Bai, Z.; Zhang, D. J.; Wang, W.; Lin, K. Q.; Gu, Y.; Chen, Z.; Yang, Z.; and Shou, M. Z. 2024. Show-o: One single transformer to unify multimodal understanding and generation. arXiv preprint arXiv:2408.12528.
- Xie, X.; Liu, J.; Fan, H.; Han, Z.; Tang, Y.; and Qu, L. 2025. DVG-Diffusion: Dual-View Guided Diffusion Model for CT Reconstruction from X-Rays. arXiv preprint arXiv:2503.17804.
- Yang, A.; Li, A.; Yang, B.; Zhang, B.; Hui, B.; Zheng, B.; Yu, B.; Gao, C.; Huang, C.; Lv, C.; et al. 2025. Qwen3 technical report. arXiv preprint arXiv:2505.09388.
- Yu, J.; Xu, Y.; Koh, J. Y.; Luong, T.; Baid, G.; Wang, Z.; Vasudevan, V.; Ku, A.; Yang, Y.; Ayan, B. K.; Hutchinson, B.; Han, W.; Parekh, Z.; Li, X.; Zhang, H.; Baldridge, J.; and Wu, Y. 2022. Scaling Autoregressive Models for Content-Rich Text-to-Image Generation. arXiv:2206.10789.
- Yu, S.; Kwak, S.; Jang, H.; Jeong, J.; Huang, J.; Shin, J.; and Xie, S. 2025. Representation Alignment for Generation: Training Diffusion Transformers Is Easier Than You Think. In *International Conference on Learning Representations*.
- Zhang, S.; Xu, Y.; Usuyama, N.; Xu, H.; Bagga, J.; Tinn, R.; Preston, S.; Rao, R.; Wei, M.; Valluri, N.; et al. 2023. BiomedCLIP: a multimodal biomedical foundation model pretrained from fifteen million scientific image-text pairs. arXiv preprint arXiv:2303.00915.
- Zhou, C.; Yu, L.; Babu, A.; Tirumala, K.; Yasunaga, M.; Shamis, L.; Kahn, J.; Ma, X.; Zettlemoyer, L.; and Levy, O. 2024. Transfusion: Predict the next token and diffuse images with one multi-modal model. arXiv preprint arXiv:2408.11039.
- Zhu, J.; Wang, W.; Chen, Z.; Liu, Z.; Ye, S.; Gu, L.; Tian, H.; Duan, Y.; Su, W.; Shao, J.; et al. 2025. InternV13: Exploring advanced training and test-time recipes for open-source multimodal models. arXiv preprint arXiv:2504.10479.