

TEXT-TO-IMAGE GENERATIVE MODELING WITHOUT LANGUAGE PRETRAINING

Keshigeyan Chandrasegaran^{1,2} Yue Zhao¹ Armin W. Thomas²
 Ehsan Adeli¹ Jiajun Wu¹ Juan Carlos Niebles¹ Li Fei-Fei¹ Michael Poli^{1,2}
¹Stanford University ²Radical Numerics

ABSTRACT

Modern text-to-image (T2I) systems use internet-scale language pretraining for text-conditioned image generation. *We ask whether language pretraining is necessary for strong T2I generation.* We introduce Halcyon, a family of early-fusion multimodal autoregressive models pretrained entirely from scratch on GPIC, a dataset of 100M image-text pairs. Halcyon-0.52B outperforms LlamaGen-XL/8 (48.0 vs. 63.3 gFD) with **18.2** $\times$ less pretraining compute. Further, Halcyon-1.6B outperforms JiT-T2I-XXL/8++ (35.3 vs. 41.0 gFD) with **275** $\times$ less pretraining compute. Controlled experiments show that text-only pretraining does not improve T2I generation. By replacing caption-dropped T2I examples with I2T sequences, Halcyon preserves unconditional image training while learning image captioning under the same pretraining budget, with competitive T2I generation quality. We show that strong T2I systems can be trained from scratch without internet-scale language pretraining. Code and models are available at <https://cs.stanford.edu/~keshik/halcyon>.




Figure 1: **Text-to-image generation without internet-scale language pretraining.** **Left:** Generation quality on GPIC (Chandrasegaran et al., 2026) (gFD, log scale) vs. total pretraining FLOPs required to build each T2I system, including language model, image tokenizer, and image generator pretraining. Each $\star$ marks total pretraining FLOPs and gFD, while the donut breaks down language and vision compute. **Halcyon (Ours)**, trained entirely from scratch, achieves lower gFD with substantially lower total pretraining compute compared to: (i) *LlamaGen* (Sun et al., 2024), late-fusion autoregressive models with 0.78–0.90B image generators and pretrained FLAN-T5 encoders (0.11–1.22B); (ii) *JiT-T2I* (Li & He, 2026; Ma et al., 2026), late-fusion diffusion models with a 1.10B image generator and pretrained Qwen3-1.7B; and (iii) *Qwen3-0.6B-CPT*, an autoregressive model continually pretrained on image–text pairs, initialized from a pretrained Qwen3-0.6B language model (Yang et al., 2025a). **Right:** 256 $\times$ 256 samples generated by Halcyon-4B with prompts shown for two samples. Best viewed in color.

1 INTRODUCTION

Text-to-image (T2I) generation maps natural-language descriptions to images. Modern T2I systems have an ingrained reliance on internet-scale language pretraining¹ for semantic conditioning of image generators. This dependence is realized via two designs, distinguished by how and when text tokens interact with image tokens: (1) *Late fusion*, the predominant design, uses two separate models: a large model is pretrained on text, and the resulting text embeddings are used to condition a separate image generator through cross- or joint attention (Saharia et al., 2022; Chen et al., 2024; Black Forest Labs, 2025; Zhao et al., 2026); (2) *Early fusion* removes the separate text encoder while retaining language pretraining within a single multimodal model trained on text-only and image-text data, with captions used as a prefix to condition image generation (Team et al., 2024a; Wang et al., 2026; Zhou et al., 2025; Liu et al., 2026). Across both designs, language pretraining dominates the compute budget, accounting for over 80% of total pretraining FLOPs in representative T2I systems (Sun et al., 2024; Liu et al., 2026). This raises two issues. First, its contribution to T2I generation has not been systematically quantified. Second, language pretraining cost is often excluded, understating the total cost of building a T2I system.

In this work, we show that strong T2I generation can be achieved without language pretraining. We introduce *Halcyon*, a family of early-fusion autoregressive models that learn text conditioning directly from image–text pairs using next-token prediction. We quantify the contribution of language pretraining to T2I generation using GPIC (Chandrasegaran et al., 2026), a dataset of 100M image–text pairs, through three experimental axes: (1) **From-scratch pretraining**: we train all components from scratch, including the text and image tokenizers and the generative model; (2) **Data composition and scheduling**: under a fixed pretraining budget, we control the mixture and schedule of text-only, T2I, and I2T examples and measure their effect on T2I generation quality; and (3) **Compute comparisons**: we pretrain representative T2I baselines on GPIC and compare generation quality as a function of total pretraining compute, decomposed into language and vision. Through these experiments, we quantify the utility of language pretraining for T2I generation.

Our results show that *Halcyon* achieves strong T2I generation without language pretraining, with generation Fréchet distance (gFD) consistently improving as model size scales (Figure 1). Particularly, *Halcyon*-0.52B achieves 48.0 gFD, outperforming LlamaGen-XL/8 (FLAN-T5-XL 1.22B + 0.90B image generator; 63.3 gFD) using **18.2× less** total pretraining compute. *Halcyon*-1.6B achieves 35.3 gFD, outperforming JiT-T2I-XXL/8++ (Qwen3-1.7B + 1.10B image generator; 41.0 gFD) using **275× less** total pretraining compute. Our data composition and scheduling experiments further show that external text-only pretraining does not improve T2I generation. Further, we show that *Halcyon* can learn image captioning by replacing caption-dropped T2I examples with I2T sequences, while maintaining competitive T2I generation quality under the same pretraining budget. We provide a simple recipe for pretraining T2I models without language pretraining. Our work prompts rethinking the role of generic language pretraining for multimodal generation.

2 RELATED WORK

Modern T2I systems use large-scale language pretraining through two designs. In *late fusion*, the predominant paradigm, a language model is pretrained separately on text-only data, and its embeddings condition a visual generator. Imagen (Saharia et al., 2022), Fluid Fan et al. (2025) and PixArt- α (Chen et al., 2024) use T5-XXL (Chung et al., 2024), LlamaGen-T2I (Sun et al., 2024) uses T5-XL, SANA (Xie et al., 2025) uses Gemma-2B (Team et al., 2024b) and PixelGen (Ma et al., 2026) uses Qwen3-1.7B. In *early fusion*, the separate text encoder is removed, but text-only




Figure 2: **Language vs. vision pretraining tokens.** Language tokens count text-only pretraining. Emu3 and Tuna-2 values are estimated.

¹Throughout this paper, *language pretraining* refers to pretraining on web-scale text-only corpora, excluding image-paired captions.

pretraining is incorporated into multimodal pretraining. Chameleon (Team et al., 2024a), Transfusion (Zhou et al., 2025), and Emu3 (Wang et al., 2026) train jointly on text-only and image-text data, while Tuna-2 (Liu et al., 2026) and SimpleAR (Wang et al., 2025) initialize from pretrained Qwen models. Figure 2 compares the text and vision tokens used to pretrain these T2I systems. Systems with language pretraining consume $10^{12} \sim 10^{14}$ text tokens, often exceeding the vision budget. In contrast, Halcyon uses no text-only language pretraining; its only language supervision comes from caption tokens paired with images, more than two orders of magnitude lower than the language-pretraining budgets shown. We note that early T2I systems such as DALL-E, CogView, and GLIDE learned text conditioning directly from paired image-text data, without language pretraining (Ramesh et al., 2021; Ding et al., 2021; Nichol et al., 2022). However, the role of language pretraining has been debated: Imagen (Saharia et al., 2022) found benefits from scaling pretrained text encoders, whereas Parti (Yu et al., 2022b) found that text-encoder pretraining only marginally reduced T2I generation loss. Similarly, Zhang et al. (2024) found limited benefits from initializing autoregressive T2I models with pretrained language models. Nevertheless, modern T2I systems have largely shifted toward pretrained language models for text conditioning. We quantify the role of language pretraining and show that T2I models trained entirely from scratch, without language pretraining, can outperform modern T2I systems relying on language pretraining.

3 APPROACH

Halcyon is a decoder-only Transformer with a joint text-image vocabulary. We model text and image tokens as a single sequence using next-token prediction, following prior multimodal early-fusion works (Team et al., 2024a; Wang et al., 2026). All components, including the text tokenizer, image tokenizer, and multimodal generative model are trained from scratch. Below, we describe the tokenizers, sequence format, training objective, and pretraining data composition and scheduling.

3.1 TOKENIZATION

Images. We train a discrete image tokenizer on GPIC images, following the encoder-decoder design of Cosmos Tokenizer (Agarwal et al., 2025) with finite scalar quantization (FSQ) (Mentzer et al., 2024). The tokenizer downsamples images by a factor of 8 along each spatial dimension, with a codebook size of 64,000. Each 256×256 image is represented as $N = 1,024$ **image tokens** $\mathbf{x} = (x_1, \dots, x_N)$ obtained by flattening the 32×32 grid in raster order.

Text. We train a byte-level BPE tokenizer (Sennrich et al., 2016) from scratch on GPIC captions and OpenWebText Gokaslan & Cohen (2019), with a vocabulary size of 16,384. Each caption is represented as $\mathbf{c} = (c_1, \dots, c_L)$.

Joint vocabulary. Text tokens, image tokens, and special tokens form a single vocabulary of size 80,640, comprising 16,384 text tokens, 64,000 image tokens, and 256 special tokens. The model uses a shared embedding table and output head. We use `[SOV]` and `[EOV]` to mark the beginning and end of an image, `[EOS]` to mark the end of a sequence, and `[PAD]` for sequence padding.

3.2 SEQUENCE FORMAT AND TRAINING OBJECTIVE

We use four sequence formats for pretraining shown in Figure 3. Text-only sequences are packed, while image-containing sequences are padded to a fixed length with `[PAD]` tokens. We do not pack multiple image-caption pairs into a sequence. For image-only (I) sequences, image tokens are predicted from `[SOV]`. For image-to-text (I2T) sequences, image tokens are predicted first, followed by caption tokens conditioned on the image. For text-to-image (T2I) sequences, image tokens are predicted conditioned on the caption. For classifier-free guidance (CFG) (Ho & Salimans, 2022), we adopt the standard practice of dropping the caption from a fraction p_{drop} of T2I examples to train the unconditional branch. At inference, the conditional and unconditional image-token logits are combined as $\ell_{\text{cfg}} = \ell_u + \gamma(\ell_c - \ell_u)$, where γ is the guidance scale. We note that I2T sequences share the same image prefix as unconditional T2I sequences, and therefore also train the unconditional image branch used for CFG. We use weighted next-token cross-entropy, $\mathcal{L} = (\mathcal{L}_{\text{text}} + \lambda_{\text{img}} \mathcal{L}_{\text{img}})/M$, where $\mathcal{L}_{\text{text}} = -\sum_{i \in \mathcal{T}} \log p_{\theta}(s_i | s_{<i})$, $\mathcal{L}_{\text{img}} = -\sum_{i \in \mathcal{I}} \log p_{\theta}(s_i | s_{<i})$, and $M = |\mathcal{T}| + |\mathcal{I}|$. Here, $\mathcal{T}$ and $\mathcal{I}$ denote the non-padding text and image tokens respectively, with special tokens included in $\mathcal{T}$. We set $\lambda_{\text{img}} = 0.5$ to prevent image tokens dominating the training loss similar to Emu3 (Wang et al., 2026).




Figure 3: **Sequence formats.** Text-only pretraining tokens, caption tokens, and image tokens occupy one joint vocabulary alongside special tokens. The top two formats are unimodal, while the bottom two represent paired image-caption examples. Rows I and I2T share the same $[\text{SOV}] \times [\text{EOV}]$ prefix, which is also obtained when dropping the caption from a T2I sequence used for classifier-free guidance. Image-caption sequences are padded to a fixed length, with $[\text{PAD}]$ positions excluded from the training loss. Box widths are schematic and not proportional to token counts.

3.3 PRETRAINING DATA COMPOSITION AND SCHEDULING

To measure the utility of language pretraining for T2I generation, we control the composition and scheduling of text-only, T2I, and I2T examples under a **fixed** pretraining budget. Figure 4 illustrates the five configurations. Unless otherwise specified, caption dropout is applied to T2I examples during training at rate p_{drop} .

- **Baseline T2I.** We train exclusively on T2I examples.
- **Staged: Language $\rightarrow$ T2I.** We allocate the first α fraction of training to text-only pretraining, followed by T2I for the remaining $1 - \alpha$ fraction. We choose α such that the text-only pretraining tokens exceed the Chinchilla-optimal token budget for the model size (Hoffmann et al., 2022). This setup studies language-first pretraining schedules used in early-fusion systems such as Tuna-2 (Liu et al., 2026). Separately, we evaluate continual pre-training recipe from an internet-scale pretrained LLM by expanding its vocabulary with image tokens and continuing training on paired image-text data.
- **Staged: I2T $\rightarrow$ T2I.** We allocate the first β fraction of training to I2T, followed by T2I for the remaining $1 - \beta$ fraction.
- **Mixing: T2I + I2T.** We interleave T2I and I2T examples throughout training, allocating a fraction η to I2T and $1 - \eta$ to T2I.
- **Mixing: T2I without caption dropout + I2T.** We set $\eta = p_{\text{drop}}$ and disable caption dropout. This preserves the baseline’s conditional and unconditional image-training tokens while additionally providing I2T supervision.




Figure 4: **Pretraining data composition and scheduling described in Section 3.3.** The four token types are: $\text{I} | \text{T}$ image tokens conditioned on text; $\text{I} | \emptyset$ image tokens without text conditioning; $\text{T} | \text{I}$ text tokens conditioned on images; and $\text{T} | \emptyset$ text-only tokens. Vertical boundaries denote training phases, while horizontal bands show the token composition within each phase. Bands represent interleaved examples. Band height and width are schematic.

4 EXPERIMENTS

4.1 EXPERIMENTAL SETUP

Dataset and pretraining protocol. We pretrain on GPIC (Chandrasegaran et al., 2026), a dataset of approximately 100M image–text pairs, at 256×256 resolution. All configurations are trained for one epoch on GPIC with a global batch size of 1,024 and an identical warmup-stable-decay (WSD) learning-rate schedule (Hu et al., 2024a). For language pretraining, we use ClimbMix (Diao et al., 2025) packed without padding. Model details and training details are provided in Appendix A.

Tokenizer training and evaluation. • *Image tokenizer.* We train our image tokenizer for one epoch using reconstruction and perceptual losses, followed by an additional half epoch with adversarial loss using a PatchGAN discriminator. Compared to Cosmos-DI-8x8 (Agarwal et al., 2025), we increase the decoder width from 128 to 192. To assess reconstruction quality, we compare our tokenizer with pretrained LLamaGen-VQ-16x16 and Cosmos-DI-8x8 tokenizers. We include Real-vs-Real scores as a reference. Our tokenizer outperforms Cosmos-DI-8x8 across all metrics (Table 1). As with prior discrete tokenizers, reconstruction fidelity is lower for small text and fine glyphs (Team et al., 2024a; Sun et al., 2024; Agarwal et al., 2025).

Table 1: **Reconstruction results on GPIC Test-50K** (256×256). Our tokenizer, trained from scratch on GPIC, outperforms Cosmos-DI-8x8 on all metrics.

Tokenizer	rFD ↓	PSNR ↑	SSIM ↑	LPIPS ↓
LLamaGen-VQ-16x16	129.87	20.18	0.560	0.146
Cosmos-DI-8x8	37.93	25.38	0.776	0.153
Halcyon	18.86	25.61	0.792	0.062
Real-vs-Real	7.59	–	–	–

• *Text tokenizer.* Despite having a $9.3\times$ smaller vocabulary than Qwen3, our tokenizer achieves comparable compression on GPIC test captions (4.774 vs. 4.763 bytes/token), outperforming the FLAN-T5 tokenizer (4.152 bytes/token; Table 2). On held-out OpenWebText, our tokenizer achieves 4.086 bytes/token. OpenWebText is used only to train the BPE tokenizer, not the multimodal generative models. Full tokenizer training details are provided in Table 14.

Table 2: **Text tokenizer comparison on GPIC test captions.** Higher bytes/token and lower tokens/caption indicate better compression.

Tokenizer	Vocab.	Bytes/token ↑	Tok./caption ↓
FLAN-T5	32,100	4.152	51.1
Qwen3	151,669	4.763	44.6
Halcyon	16,384	4.774	44.5

Model configurations. We train *Halcyon* at five sizes: 0.1B, 0.25B, 0.52B, 1.6B, and 4B parameters. Model architecture details are provided in Table 11. We use a sequence length of 1280 for all experiments. We use $p_{drop} = 0.1$ unless otherwise specified.

Late-fusion methods. Late fusion is the predominant design choice for modern T2I systems. Selecting baselines for this study is non-trivial: existing systems differ in pretraining objective, language conditioning, image representation, and training setup, while our study requires public pretraining code and models that enable controlled interventions to study the role of language pretraining. We select two representative families, LlamaGen (Sun et al., 2024) and JiT-T2I (Li & He, 2026; Ma et al., 2026). They cover (i) autoregressive and diffusion-based generation, (ii) encoder-decoder and decoder-only language models of different sizes, and (iii) image sequence lengths of 256 and 1,024. For both families, we use the authors’ official codebases and follow their training configurations unless stated otherwise. Below, we describe our setup and summarize them in Table 3.

• **LlamaGen (Sun et al., 2024).** This method uses a FLAN-T5-XL text encoder for language modeling, a discrete image tokenizer, and an autoregressive image generator.

- **Default (LLamaGen-XL).** We use the original LlamaGen configuration: FLAN-T5-XL, the pretrained image tokenizer producing 256 image tokens from a codebook of 16,384 entries, and the 0.78B image generator.
- **Use *Halcyon* image tokenizer (LlamaGen-XL/8).** Starting with the default configuration, we replace the image tokenizer with our GPIC-trained tokenizer. This increases the visual vocabulary size by $4\times$, the image sequence length from 256 to 1,024, and the image generator size from 0.78B to 0.90B.
- **Use smaller text encoder (LlamaGen-XL-T5-Base/8).** Starting with LlamaGen-XL/8, we replace FLAN-T5-XL (1.22B) with FLAN-T5-base (0.11B) to study the effect of language model size.

Table 3: **Experiment configurations for late-fusion methods.** LlamaGen and JiT-T2I configurations, varying the language model, image sequence length, and perceptual supervision. LlamaGen and Halcyon are autoregressive; JiT is diffusion-based. Pretraining tokens report the total language and vision tokens used during training; Halcyon uses no pretrained language model, and all its language tokens come from GPIC captions. We report encoder parameters for FLAN-T5.

Configuration	Language model	Params		Pretraining tokens		Perceptual loss
		Language	Vision	Language	Vision	
LlamaGen-XL	FLAN-T5-XL	1.22B	0.78B	~1T	25.6B	✓
LLamaGen-XL/8	FLAN-T5-XL	1.22B	0.90B	~1T	102.4B	✓
LLamaGen-XL-T5-Base/8	FLAN-T5-Base	0.11B	0.90B	~1T	102.4B	✓
JiT-T2I-XXL/16	Qwen3	1.70B	1.10B	36T	25.6B	
JiT-T2I-XXL/16 ++	Qwen3	1.70B	1.10B	36T	25.6B	✓
JiT-T2I-XXL/8 ++	Qwen3	1.70B	1.10B	36T	102.4B	✓
Halcyon (Ours)	—	0.1–4B		4.5B	102.4B	✓

• **JiT-T2I (Li & He, 2026; Ma et al., 2026).** This method uses pretrained Qwen3-1.7B (Yang et al., 2025a) as the language model and a pixel-space flow-matching image generator.

- **Default (JiT-T2I-XXL/16).** We use the original JiT-T2I-XXL/16 configuration: Qwen3-1.7B and no image tokenizer, with 16×16 image patches producing 256 image tokens.
- **Add perceptual loss (JiT-T2I-XXL/16 ++).** Starting with the default configuration, we add LPIPS loss (Zhang et al., 2018) to train the image generator. This tests the effect of perceptual supervision used to also train the image tokenizers in Halcyon and LlamaGen.
- **Use 8×8 patches (JiT-T2I-XXL/8 ++).** Starting with JiT-T2I-XXL/16 ++, we reduce the image patch size from 16×16 to 8×8 , increasing the image sequence length from 256 to 1,024, matching Halcyon and the modified LlamaGen configurations.

Early-fusion methods. Existing early-fusion models, including Chameleon (Team et al., 2024a) and Emu3 (Wang et al., 2026), incorporate language pretraining into multimodal generative modeling. We study the role of language pretraining in two controlled settings. First, we vary the composition and scheduling of text-only, T2I, and I2T sequences within Halcyon under a fixed pretraining budget (Section 3.3). Second, we study continual pretraining by initializing from Qwen3-0.6B-Base, pretrained on 36T text tokens, expanding its vocabulary with image tokens, and continuing pretraining on GPIC using the baseline T2I recipe. For comparison, we train the same architecture and vocabulary from scratch under an identical T2I pretraining budget. We disable weight decay for embedding parameters when vocabulary size is large (216K).

Evaluation protocol. We evaluate all models at 256×256 resolution on three benchmarks: • **GPIC (primary metrics).** We follow the GPIC evaluation protocol using the GPIC Test-50K captions², reporting generation Fréchet distance (gFD), Precision, Recall, Density, Coverage, and MMD. • **GenEval (Ghosh et al., 2023).** We evaluate compositional generation using the rephrased prompts from PixelGen (Ma et al., 2026)³, removing generic phrases such as “an image of” and “the image depicts”. • **DPG-Bench (Hu et al., 2024b).** We evaluate complex prompt adherence on DPG-Bench using the original prompts. • **Token-level NLL.** For Halcyon, we report negative log-likelihood (NLL) on GPIC Test-50K for four settings: text-conditioned images ($\mathcal{I} \mid \mathcal{T}$), unconditional images ($\mathcal{I} \mid \emptyset$), image-conditioned text ($\mathcal{T} \mid \mathcal{I}$), and unconditional text ($\mathcal{T} \mid \emptyset$). • **Image captioning.** We evaluate image–text alignment using Qwen3-VL-Reranker-8B scores over GPIC Test-50K images and generated captions, treating the image as the query and the caption as the document.

For each model, we sweep guidance scales in increments of 0.5 on GPIC Validation-50K⁴ and use the scale achieving the lowest gFD for all three benchmarks. We report the sampling and guidance settings for each model in Appendix A.

²hf.co/datasets/stanford-vision-lab/gpic/blob/main/gpic_test_50k.jsonl

³github.com/Zehong-Ma/PixelGen/blob/main/evaluations/geneval/

⁴hf.co/datasets/stanford-vision-lab/gpic/blob/main/gpic_val_50k.jsonl

4.2 RESULTS: COMPARISON WITH LATE-FUSION T2I METHODS

We report quantitative results in Table 4, including compositional generation benchmarks (GenEval and DPG-Bench) in Table 5. We include generated samples in Figure 5 and Appendix C. Figure 1 compares generation quality as a function of total pretraining FLOPs, with the exact compute breakdown reported in Table 15. Below, we discuss our key findings:

Table 4: Comparison of Halcyon with late-fusion T2I methods on GPIC. Without language pretraining, Halcyon achieves stronger T2I generation than the late-fusion baselines: Halcyon-0.52B outperforms LlamaGen-XL/8 (48.0 vs. 63.3 gFD), while Halcyon-1.60B outperforms JiT-T2I-XXL/8++ (35.3 vs. 41.0 gFD). Vision parameter counts refer to the generative models, excluding image tokenizers (100M for LlamaGen and 126M for Halcyon). The CFG scale for each model is reported in Appendix Table 16.

Model	Language	Vision	gFD ↓	Prec. ↑	Rec. ↑	Dens. → 1	Cov. ↑	MMD ↓
LlamaGen-XL	1.22B	0.78B	147.01	0.914	0.715	0.925	0.822	0.174
LLamaGen-XL/8	1.22B	0.90B	63.29	0.948	0.827	1.058	0.925	0.041
LLamaGen-XL-T5-Base/8	0.11B	0.90B	66.69	0.941	0.822	1.016	0.924	0.045
JiT-T2I-XXL/16	1.70B	1.10B	67.92	0.934	0.822	0.979	0.923	0.038
JiT-T2I-XXL/16 ++	1.70B	1.10B	48.86	0.947	0.861	1.022	0.939	0.022
JiT-T2I-XXL/8 ++	1.70B	1.10B	40.97	0.949	0.880	1.049	0.952	0.016
Halcyon (Ours)	0.10B		144.16	0.924	0.636	1.029	0.853	0.131
	0.25B		72.55	0.946	0.796	1.050	0.922	0.050
	0.52B		48.02	0.953	0.858	1.064	0.938	0.027
	1.60B		35.28	0.955	0.891	1.071	0.955	0.019
	4.00B		28.43	0.951	0.905	1.034	0.960	0.014

Effect of image representation and perceptual supervision. For LlamaGen-XL, replacing its original image tokenizer with our GPIC-trained tokenizer increases the image sequence length from 256 to 1,024 and reduces gFD (147.0 $\rightarrow$ 63.3). For JiT-T2I-XXL/16, adding LPIPS perceptual supervision reduces gFD (67.9 $\rightarrow$ 48.9). Reducing the image patch size from 16×16 to 8×8 increases the image sequence length from 256 to 1,024 and further reduces gFD (48.9 $\rightarrow$ 41.0). We use LlamaGen-XL/8 and JiT-T2I-XXL/8++ for subsequent comparisons with Halcyon.

Effect of language encoder size. We study whether a larger pretrained language encoder improves T2I generation. For LlamaGen, replacing FLAN-T5-XL with FLAN-T5-base reduces the language encoder size from 1.22B to 0.11B while retaining the same image tokenizer and 0.90B image generator. The resulting gFD changes only marginally (63.3 $\rightarrow$ 66.7), despite a $13.5\times$ reduction in language-model training compute⁵. This provides initial evidence that an order-of-magnitude increase in language pretraining compute yields only marginal improvements in T2I generation. We next investigate whether strong T2I generation can be achieved without language pretraining altogether.

Table 5: Comparison of Halcyon with late-fusion T2I methods on compositional generation benchmarks.

Model	GPIC gFD ↓	GenEval ↑	DPG-Bench ↑
LlamaGen-XL	147.01	30.03	65.40
LLamaGen-XL/8	63.29	30.48	67.08
LLamaGen-XL-T5-Base/8	66.69	28.48	64.67
JiT-T2I-XXL/16	67.92	36.27	72.69
JiT-T2I-XXL/16 ++	48.86	41.79	76.58
JiT-T2I-XXL/8 ++	40.97	43.17	76.85
Halcyon-0.10B	144.16	27.31	66.36
Halcyon-0.25B	72.55	35.72	72.66
Halcyon-0.52B	48.02	44.91	76.84
Halcyon-1.60B	35.28	53.48	77.90
Halcyon-4.00B	28.43	56.39	77.90

T2I generation without language pretraining. We now address the central question of this work: *can a T2I model trained entirely from scratch match or surpass strong late-fusion systems that rely on pretrained language models?* Halcyon-0.52B, a single multimodal model trained entirely from scratch, achieves 48.0 gFD, outperforming LlamaGen-XL/8 with pretrained FLAN-T5-XL (63.3) and approaching JiT-T2I-XXL/8++ with pretrained Qwen3-1.7B (41.0). Increasing Halcyon to 1.6B parameters further reduces gFD to 35.3, surpassing both late-fusion methods. At 4B parameters, gFD decreases to 28.4. On compositional generation benchmarks, Halcyon-0.52B outper-

⁵Pretraining and finetuning the full FLAN-T5-XL encoder-decoder model requires 9.06×10^{21} FLOPs vs. 6.69×10^{20} FLOPs for FLAN-T5-base (Chung et al., 2024)

forms JiT-T2I-XXL/8++ on GenEval (44.9 vs. 43.2) and matches its DPG-Bench score (76.84 vs. 76.85). At 1.6B, Halcyon surpasses JiT-T2I-XXL/8++ on both GenEval (53.5 vs. 43.2) and DPG-Bench (77.9 vs. 76.9). Moreover, Halcyon-1.6B achieves these results with $275\times$ less total pre-training compute than JiT-T2I-XXL/8++. Together, these results demonstrate that strong T2I generation capabilities can be learned from paired image-text data alone, without language pretraining, while substantially reducing total pretraining compute.

Takeaway. Halcyon demonstrates that T2I models trained entirely from scratch, without internet-scale language pretraining, can outperform strong late-fusion methods while requiring substantially less total pretraining compute.




Figure 5: **Comparison of generated images.** From left to right: JiT-T2I-XXL/8++, LlamaGen-XL/8, and Halcyon models from 0.1B to 4B parameters. Image quality improves as Halcyon scales up. Additional samples are provided in Appendix C.

4.3 RESULTS: PRETRAINING DATA COMPOSITION AND SCHEDULING

External language pretraining does not improve T2I generation. We investigate this through two controlled interventions. First, under a fixed pretraining budget, allocating $\alpha=0.1, 0.2, 0.5$ to text-only pretraining reduces paired image-text training and worsens gFD from 48.0 to 51.6, 52.8, and 65.1, respectively; conditional image-token NLL also increases (Table 6). To isolate the effect of language pretraining from reduced T2I training, we additionally evaluate language pretraining followed by a full T2I epoch ($1.5\times$ total compute), but gFD does not improve (48.0 vs. 49.2). Second, we investigate continual pretraining from Qwen3-0.6B-Base, pretrained on 36T text tokens. Despite this extensive language pretraining, Qwen3-0.6B-CPT performs similar to the same architecture trained from random initialization (47.8 vs. 46.6 gFD). Reducing the peak learning rate from 7.5×10^{-4} to 1×10^{-4} for continual pretraining also fails to improve T2I generation (Table 7). Together, these experiments show that external text-only pretraining provides no benefit to T2I generation under the evaluated settings.

Table 6: **Effect of language pretraining on Halcyon-0.52B.** Under a fixed pretraining budget, we measure the impact of allocating compute to external text-only pretraining on T2I generation. Each language-pretraining stage exceeds the Chinchilla reference of 20 text tokens per parameter. The $\dagger$ configuration preserves a full T2I epoch after additional language pretraining ($1.5\times$ total pretraining compute), but does not improve gFD. Token counts exclude padding.

Configuration	α	Text tokens (B)		Image tokens (B)		NLL $\downarrow$				gFD $\downarrow$
		$\mathcal{T} \mid \emptyset$	$\mathcal{T} \mid \mathcal{I}$	$\mathcal{I} \mid \emptyset$	$\mathcal{I} \mid \mathcal{T}$	$\mathcal{T} \mid \emptyset$	$\mathcal{T} \mid \mathcal{I}$	$\mathcal{I} \mid \emptyset$	$\mathcal{I} \mid \mathcal{T}$	
Baseline	–	4.05	0	10.24	92.16	1.433	18.414	8.428	8.366	48.02
Staged: Language $\rightarrow$ T2I	0.10	16.44	0	9.22	82.94	1.447	19.098	8.442	8.382	51.56
	0.20	28.83	0	8.19	73.73	1.451	17.197	8.445	8.388	52.76
	0.50	65.80	0	5.12	46.08	1.454	22.850	8.505	8.451	65.13
Staged: Language $\rightarrow$ T2I †	0.33	68.03	0	10.24	92.16	1.445	18.809	8.428	8.371	49.22

Table 7: **Continual pretraining from a language model.** Qwen3-0.6B-CPT is initialized from Qwen3-0.6B-Base (pretrained on 36T text tokens), while Qwen3-0.6B is randomly initialized. Both share the same Transformer architecture as Halcyon-0.52B but use a larger vocabulary. We additionally evaluate continual pretraining with a lower peak learning rate (1×10^{-4} vs. 7.5×10^{-4}) to test whether aggressive weight updates diminish the benefits of pretrained language weights. Neither configuration improves T2I generation over random initialization. All models are trained for one epoch on GPIC and evaluated with CFG 6.0.

Model	Vocab.	gFD $\downarrow$	Prec. $\uparrow$	Rec. $\uparrow$	Dens. $\rightarrow 1$	Cov. $\uparrow$	MMD $\downarrow$
Halcyon-0.52B	80,640	48.02	0.953	0.858	1.064	0.938	0.027
Qwen3-0.6B	216,000	46.61	0.953	0.864	1.067	0.943	0.026
Qwen3-0.6B-CPT	216,000	47.81	0.949	0.870	1.041	0.940	0.029
Qwen3-0.6B-CPT (lower LR)	216,000	63.62	0.951	0.823	1.073	0.937	0.043

I2T supervision as an additional task does not improve T2I generation. We investigate whether adding I2T supervision as a separate task through staged or mixed training improves T2I generation under a fixed pretraining budget (Table 8). At an I2T allocation of 0.1, both schedules achieve gFD comparable to the T2I baseline (48.4 and 48.2 vs. 48.0). Increasing the allocation to 0.5 reduces conditional image tokens by half ($92.2\text{B} \rightarrow 46.1\text{B}$), with staged training outperforming mixing (54.2 vs. 76.0 gFD), although neither improves upon the baseline. Conversely, mixing substantially reduces conditional text NLL compared to staging (1.009 vs. 20.437), indicating better image captioning. These results suggest that allocating part of the pretraining budget to I2T as a separate task does not benefit T2I generation, despite improving captioning.

Table 8: **Effect of I2T supervision and scheduling on T2I generation for Halcyon-0.52B.** Under a fixed pretraining budget, staged and mixed training use identical token counts at each I2T allocation, differing only in scheduling. At 10%, both schedules preserve comparable T2I generation, while at 50%, staged training achieves substantially lower gFD. However, only mixing substantially reduces captioning NLL, namely $\text{NLL}(\mathcal{T} \mid \mathcal{I})$. NLL and gFD are evaluated on GPIC Test-50K; token counts exclude padding.

Configuration	β (or η)	Text tokens (B)		Image tokens (B)		NLL $\downarrow$				gFD $\downarrow$
		$\mathcal{T} \mid \emptyset$	$\mathcal{T} \mid \mathcal{I}$	$\mathcal{I} \mid \emptyset$	$\mathcal{I} \mid \mathcal{T}$	$\mathcal{T} \mid \emptyset$	$\mathcal{T} \mid \mathcal{I}$	$\mathcal{I} \mid \emptyset$	$\mathcal{I} \mid \mathcal{T}$	
Baseline	–	4.05	0	10.24	92.16	1.433	18.414	8.428	8.366	48.02
Staged: I2T $\rightarrow$ T2I	$\beta = 0.1$	3.65	0.45	19.46	82.94	1.440	18.415	8.426	8.366	48.35
	$\beta = 0.5$	2.03	2.25	56.32	46.08	1.435	20.437	8.421	8.370	54.22
Mixing: T2I + I2T	$\eta = 0.1$	3.65	0.45	19.46	82.94	1.443	1.125	8.414	8.362	48.18
	$\eta = 0.5$	2.03	2.25	56.32	46.08	1.475	1.009	8.453	8.417	75.97

Replacing caption dropout with I2T supervision enables image captioning under the same pretraining budget. Rather than allocating a separate training budget to I2T, we replace caption-dropped T2I examples ($p_{\text{drop}} = 0.1$) with I2T examples ($p_{\text{drop}} = 0, \eta = 0.1$). Since I2T sequences share the $[\text{SOV}] x_1 \cdots x_N [\text{EOV}]$ prefix, this preserves the conditional and unconditional image-token counts under the same pretraining budget. Across all model sizes, conditional text NLL decreases by more than an order of magnitude (e.g., $18.414 \rightarrow 1.102$ at 0.52B), while gFD, GenEval, and DPG-Bench remain comparable to the baseline (Table 9). Caption alignment also increases with model size ($0.090 \rightarrow 0.559$), with qualitative examples from Halcyon-4.00B shown in Figure 6. Although caption alignment remains below that of GPIC ground truth captions, its consistent improvement with model size demonstrates the potential to learn additional multimodal capabilities while maintaining competitive T2I generation under the same pretraining budget.

Table 9: **Replacing caption dropout with I2T supervision across model sizes.** Replacing caption-dropped T2I examples with I2T preserves the pretraining budget and image-token counts, enabling image captioning while maintaining comparable T2I generation. We report mean Qwen3-VL-Reranker-8B alignment scores over 50K generated image–caption pairs from GPIC Test-50K, using the image as the query and caption as the document. GPIC Test-50k image–caption pairs obtain a mean alignment score of 0.891.

Size	η	p_{drop}	NLL $\downarrow$				T2I generation			Captioning
			$\mathcal{T} \mid \emptyset$	$\mathcal{T} \mid \mathcal{I}$	$\mathcal{I} \mid \emptyset$	$\mathcal{I} \mid \mathcal{T}$	gFD $\downarrow$	GenEval $\uparrow$	DPG-Bench $\uparrow$	Alignment $\uparrow$
0.10B	–	0.1	1.601	19.066	8.774	8.727	144.16	27.31	66.36	–
	0.1	0	1.598	1.410	8.771	8.724	144.14	27.66	68.89	0.090
0.25B	–	0.1	1.497	21.472	8.566	8.509	72.55	35.72	72.66	–
	0.1	0	1.491	1.218	8.559	8.503	70.44	36.54	73.07	0.207
0.52B	–	0.1	1.433	18.414	8.428	8.366	48.02	44.91	76.84	–
	0.1	0	1.435	1.102	8.419	8.358	46.48	44.60	76.14	0.315
1.6B	–	0.1	1.370	20.426	8.303	8.235	35.28	53.48	77.90	–
	0.1	0	1.372	0.977	8.292	8.228	35.91	53.51	78.21	0.467
4B	–	0.1	1.338	16.738	8.193	8.122	28.43	56.39	77.90	–
	0.1	0	1.331	0.898	8.204	8.125	29.30	55.33	78.61	0.559

Ground truth Caption	A black and white dog with a happy expression sits on concrete, looking up.	A rider in a helmet sits on an orange dirt bike on a muddy forest trail.	A purple flower with spiky petals grows on a green stem. The background is blurred grass and stems.	A person in a red and blue Spider-Man costume crouches on a tiled floor. They point forward with one hand, wearing the iconic mask and suit. The setting appears to be indoors, possibly a convention or event space.
Ground truth Image				
Generated Caption	A black and white dog sits on a paved surface, looking up with its mouth open. A gray hose lies behind it.	A motocross rider in a colorful suit and helmet rides an orange dirt bike on a muddy trail. Trees and dry grass surround the path, with a yellow marker visible on the right side.	A purple flower with long, spiky petals stands in the foreground. Its green stem rises from a blurred background of grass and foliage.	A person in a Spider-Man costume crouches on a tiled floor.

Figure 6: **Image captioning examples from Halcyon-4B.** For each example, the ground-truth caption is shown above the input image, and the generated caption is shown below. The model learns image captioning by replacing caption-dropped T2I examples with I2T supervision, using only 10% of the pretraining budget for captioning.

5 DISCUSSION AND CONCLUSION

Implications of language pretraining for T2I generation. Our results challenge the prevailing reliance on internet-scale language pretraining for T2I generation. In late-fusion systems such as LlamaGen, frozen pretrained language representations condition a separately trained image generator. While language pretraining may improve these representations, it does not directly learn the mapping from text to images. In contrast, Halcyon jointly learns text representations and image generation from paired image–text data. This difference may help explain our results, although our comparisons do not isolate its effect. Similarly, Parti (Yu et al., 2022a) found that text-encoder pretraining provided only marginal improvements for its 3B model. For early-fusion models, text-only pretraining optimizes $p(\mathcal{T} \mid \emptyset)$, which may not provide a useful initialization for $p(\mathcal{I} \mid \mathcal{T})$ when image tokens come from an independently trained tokenizer. Consistent with this hypothesis, our randomly initialized Qwen3 model achieves lower validation NLL (vision) at every evaluated checkpoint (Appendix Figure 10) and lower final gFD than its language-pretrained counterpart. Importantly, reusing pretrained language models does not eliminate their upstream pretraining cost. Beyond T2I, our findings motivate reconsidering the role of modality-specific pretraining and the choice between early and late fusion in other multimodal settings, including biological modalities.

Model scaling and classifier-free guidance. We examine how Halcyon scales from 0.1B to 4B parameters. Figure 7 shows that both image and text validation NLL decrease consistently with model size, with larger models achieving lower NLL throughout pretraining. This demonstrates improvements in modeling both modalities as model capacity increases, despite training entirely from scratch on paired image–text data. We further examine how the optimal CFG scale varies with model size (Figure 8). As model size increases from 0.1B to 4B, the guidance scale minimizing gFD on GPIC Validation-50K decreases from 7.5 to 4.5, while gFD improves from 142.7 to 28.4. This suggests that larger models require less guidance to achieve their best generation quality, although CFG remains beneficial across all evaluated model sizes.

Caption length and text conditioning. Since Halcyon learns text conditioning directly from paired image–text data, we investigate how generation quality varies with the length of the caption. We evaluate all five model sizes using short, medium, and long captions from GPIC Test-1M (Table 10). As model size increases from 0.1B to 4B, gFD decreases across all caption lengths. Medium




Figure 7: **Scaling behavior of Halcyon across model sizes.** Validation negative log-likelihood (NLL) for image tokens (left) and text tokens (right) throughout pretraining on GPIC. As model size increases from 0.1B to 4B, both image and text NLL consistently decrease, with larger models achieving lower NLL throughout training. All models are pretrained from scratch under the same training schedule.

and long captions consistently achieve lower gFD than short captions, with comparable performance between medium and long captions from 0.52B onward. These results suggest that more detailed text descriptions improve T2I generation quality, although longer captions do not consistently outperform medium-length captions.




Figure 8: CFG sweep on GPIC validation. The optimal scale drops from 7.5 (0.1B) to 4.5 (4B) as gFD improves from 142.7 to 28.4; these scales are used on Test-50K.

Table 10: **Effect of caption length on T2I generation.** We report gFD (↓) for Halcyon across model sizes using 50K examples with short, medium, and long captions from GPIC Test-1M. Test-50K results are included for comparison. Across all model sizes, medium and long captions achieve lower gFD than short captions, with comparable performance between medium and long captions from 0.52B onward. Generation quality consistently improves with model size across all caption lengths.

Model	Test-50K	Short	Medium	Long
0.1B	144.16	149.03	139.27	147.85
0.25B	72.55	77.18	70.14	72.53
0.52B	48.02	51.99	46.12	46.83
1.6B	35.28	38.59	34.21	34.33
4B	28.43	31.48	27.91	27.77

Limitations and future directions. While our experiments use GPIC, which is, to our knowledge, the largest high-quality, centrally hosted, publicly available dataset and benchmark for visual generation, extending these findings to other datasets and higher resolutions remains important. Our controlled data composition and scheduling experiments use Halcyon-0.52B, while the pretrained initialization comparison uses Qwen3-0.6B; Extensions to larger model sizes and pretraining budgets would help establish the generality of our findings. Although our late-fusion baselines include diffusion-based generators, our controlled interventions focus on autoregressive models. Extending these experiments to Transfusion-style architectures that combine autoregressive text modeling with diffusion-based image generation is an important direction for future work. Finally, text-to-video systems such as Wan (Wan et al., 2025), CogVideoX (Yang et al., 2025b), and Mochi (Team, 2024)

also rely on large pretrained language encoders for video generation. Investigating whether T2V models can similarly learn text conditioning directly from paired video–text data, without language pretraining, is a promising direction toward efficient video generation systems.

Conclusion. We show that strong text-to-image generation does not require internet-scale language pretraining. *Halcyon*, trained entirely from scratch on paired image–text data, outperforms strong late-fusion baselines while using substantially less total pretraining compute. Our controlled experiments show that external text-only pretraining does not improve T2I generation under the evaluated settings. Furthermore, replacing caption-dropped T2I examples with I2T supervision enables image captioning while maintaining competitive generation quality under the same pretraining budget. Together, these findings suggest that visually grounded language supervision from image–text pairs is sufficient for strong T2I generation, challenging the view that internet-scale language pretraining is a prerequisite.

AI USE STATEMENT

In this work, we used generative AI tools to assist with implementing methods. We did not use generative AI tools to propose or refine hypotheses, design or provide feedback on research methodology or experiments, support qualitative or thematic data analysis, or interpret results. Generating synthetic datasets, developing theoretical models or conceptual frameworks, formulating mathematical claims, providing critical ingredients for proving mathematical claims, assisting in the writing of proofs, assisting with translation, and cleaning or reformatting datasets were not applicable to our work. Additionally, we used generative AI tools to create or edit software code, search for and identify relevant literature and information, draft and edit parts of the manuscript, and improve readability. We reviewed all AI-assisted work. AI-generated or edited code was inspected and tested for correctness, and information and references suggested by AI tools were verified against the original sources. We take responsibility for the final content of this work, including text, claims, code, and artifacts produced with the aid of generative AI.

ETHICS STATEMENT

Our work studies text-to-image generative modeling, which can be used to generate realistic visual content and may therefore enable misuse. By providing a significantly lower-cost recipe for training strong T2I models without internet-scale language pretraining, our work may also lower the barrier to developing capable generative systems, including for malicious use. *Halcyon* and the baseline models are pretrained on GPIC and may inherit biases or limitations present in its images and captions. Our work does not involve human-subject studies or the collection of a new dataset. We do not identify additional ethical risks specific to our method beyond those associated with visual generative models.

REPRODUCIBILITY STATEMENT

We provide the model architecture, tokenization, training objective, data composition and scheduling, baseline modifications, evaluation protocol, and compute accounting in the main paper, with full architecture, optimization, sampling, and guidance details in the Appendix. GPIC and the evaluation protocol used in this work are publicly available. Our model checkpoints are released.

ACKNOWLEDGMENTS

We thank Radical Numerics for providing compute for this project. We thank Willie Neiswanger, Garyk Brixi, Tiange Xiang, Wenlong Huang, Mark Endo, Federico Berto, Zane Durante, Bailey Trang, Yanchen Wang, and Suzannah Wistreich for feedback on the manuscript. We also thank Suchir Agarwal for valuable discussions, the broader Radical Numerics research team for their feedback, and Alessandro Moro for infrastructure support.

REFERENCES

- Niket Agarwal, Arslan Ali, Maciej Bala, Yogesh Balaji, Erik Barker, Tiffany Cai, Prithvijit Chattopadhyay, Yongxin Chen, Yin Cui, Yifan Ding, et al. Cosmos world foundation model platform for physical ai. *arXiv preprint arXiv:2501.03575*, 2025.
- Black Forest Labs. FLUX.2. <https://github.com/black-forest-labs/flux2>, 2025.
- Keshigeyan Chandrasegaran, Kyle Sargent, Suchir Agarwal, Michael Jang, Michael Poli, Juan Carlos Niebles, Justin Johnson, Jiajun Wu, and Li Fei-Fei. Gpic: A giant permissive image corpus for visual generation. *arXiv preprint arXiv:2605.30341*, 2026.
- Junsong Chen, Jincheng Yu, Chongjian Ge, Lewei Yao, Enze Xie, Zhongdao Wang, James Kwok, Ping Luo, Huchuan Lu, and Zhenguo Li. Pixart- α : Fast training of diffusion transformer for photorealistic text-to-image synthesis. In *ICLR*, 2024.
- Hyung Won Chung, Le Hou, Shayne Longpre, Barret Zoph, Yi Tay, William Fedus, Yunxuan Li, Xuezhi Wang, Mostafa Dehghani, Siddhartha Brahma, et al. Scaling instruction-finetuned language models. *JMLR*, 25(70):1–53, 2024.
- Shizhe Diao, Yu Yang, Yonggan Fu, Xin Dong, Dan Su, Markus Kliegl, Zijia Chen, Peter Belcak, Yoshi Suhara, Hongxu Yin, et al. Nemotron-climb: Clustering-based iterative data mixture bootstrapping for language model pre-training. *NeurIPS*, 2025.
- Ming Ding, Zhuoyi Yang, Wenyi Hong, Wendi Zheng, Chang Zhou, Da Yin, Junyang Lin, Xu Zou, Zhou Shao, Hongxia Yang, et al. Cogview: Mastering text-to-image generation via transformers. *NeurIPS*, 2021.
- Lijie Fan, Tianhong Li, Siyang Qin, Yuanzhen Li, Chen Sun, Michael Rubinstein, Deqing Sun, Kaiming He, and Yonglong Tian. Fluid: Scaling autoregressive text-to-image generative models with continuous tokens. In *ICLR*, 2025.
- Dhruba Ghosh, Hannaneh Hajishirzi, and Ludwig Schmidt. Geneval: An object-focused framework for evaluating text-to-image alignment. *NeurIPS*, 2023.
- Aaron Gokaslan and Vanya Cohen. Openwebtext corpus. <http://Skyllion007.github.io/OpenWebTextCorpus>, 2019.
- Jonathan Ho and Tim Salimans. Classifier-free diffusion guidance. *arXiv preprint arXiv:2207.12598*, 2022.
- Jordan Hoffmann, Sebastian Borgeaud, Arthur Mensch, Elena Buchatskaya, Trevor Cai, Eliza Rutherford, Diego de Las Casas, Lisa Anne Hendricks, Johannes Welbl, Aidan Clark, et al. Training compute-optimal large language models. In *NeurIPS*, 2022.
- Shengding Hu, Yuge Tu, Xu Han, Ganqu Cui, Chaoqun He, Weilin Zhao, Xiang Long, Zhi Zheng, Yewei Fang, Yuxiang Huang, Xinrong Zhang, Zhen Leng Thai, Chongyi Wang, Yuan Yao, Chenyang Zhao, Jie Zhou, Jie Cai, Zhongwu Zhai, Ning Ding, Chao Jia, Guoyang Zeng, dahai li, Zhiyuan Liu, and Maosong Sun. MiniCPM: Unveiling the potential of small language models with scalable training strategies. In *CoLM*, 2024a.
- Xiwei Hu, Rui Wang, Yixiao Fang, Bin Fu, Pei Cheng, and Gang Yu. Ella: Equip diffusion models with llm for enhanced semantic alignment. *arXiv preprint arXiv:2403.05135*, 2024b.
- Tianhong Li and Kaiming He. Back to basics: Let denoising generative models denoise. In *CVPR*, 2026.
- Zhiheng Liu, Weiming Ren, Xiaoke Huang, Shoufa Chen, Tianhong Li, Mengzhao Chen, Yatai Ji, Sen He, Jonas Schult, Belinda Zeng, et al. Tuna-2: Pixel embeddings beat vision encoders for multimodal understanding and generation. *arXiv preprint arXiv:2604.24763*, 2026.
- Zehong Ma, Ruihan Xu, and Shiliang Zhang. Pixelgen: Improving pixel diffusion with perceptual supervision. *arXiv preprint arXiv:2602.02493*, 2026.

- Fabian Mentzer, David Minnen, Eirikur Agustsson, and Michael Tschannen. Finite scalar quantization: Vq-vae made simple. In *ICLR*, 2024.
- Alexander Quinn Nichol, Prafulla Dhariwal, Aditya Ramesh, Pranav Shyam, Pamela Mishkin, Bob McGrew, Ilya Sutskever, and Mark Chen. GLIDE: Towards photorealistic image generation and editing with text-guided diffusion models. In *ICML*, 2022.
- Aditya Ramesh, Mikhail Pavlov, Gabriel Goh, Scott Gray, Chelsea Voss, Alec Radford, Mark Chen, and Ilya Sutskever. Zero-shot text-to-image generation. In *ICML*, 2021.
- Chitwan Saharia, William Chan, Saurabh Saxena, Lala Li, Jay Whang, Emily L Denton, Kamyar Ghasemipour, Raphael Gontijo Lopes, Burcu Karagol Ayan, Tim Salimans, et al. Photorealistic text-to-image diffusion models with deep language understanding. *NeurIPS*, 2022.
- Rico Sennrich, Barry Haddow, and Alexandra Birch. Neural machine translation of rare words with subword units. In *ACL*, 2016.
- Peize Sun, Yi Jiang, Shoufa Chen, Shilong Zhang, Bingyue Peng, Ping Luo, and Zehuan Yuan. Autoregressive model beats diffusion: Llama for scalable image generation. *arXiv preprint arXiv:2406.06525*, 2024.
- Chameleon Team, Mingda Chen, Jacob Kahn, and Shang-Wen Li. Chameleon: Mixed-modal early-fusion foundation models. *arXiv preprint arXiv:2405.09818*, 2024a.
- Gemma Team, Morgane Riviere, Shreya Pathak, Pier Giuseppe Sessa, Cassidy Hardin, Surya Bhupatiraju, Léonard Hussenot, Thomas Mesnard, Bobak Shahriari, Alexandre Ramé, et al. Gemma 2: Improving open language models at a practical size. *arXiv preprint arXiv:2408.00118*, 2024b.
- Genmo Team. Mochi 1. <https://github.com/genmoai/models>, 2024.
- Team Wan, Ang Wang, Baole Ai, Bin Wen, Chaojie Mao, Chen-Wei Xie, Di Chen, Fei Wu, Haiming Zhao, Jianxiao Yang, et al. Wan: Open and advanced large-scale video generative models. *arXiv preprint arXiv:2503.20314*, 2025.
- Junke Wang, Zhi Tian, Xun Wang, Xinyu Zhang, Weilin Huang, Zuxuan Wu, and Yu-Gang Jiang. SimpleAR: Pushing the frontier of autoregressive visual generation through pretraining, sft, and rl. *arXiv preprint arXiv:2504.11455*, 2025.
- Xinlong Wang, Yufeng Cui, Jinsheng Wang, Fan Zhang, Yuezhe Wang, Xiaosong Zhang, Zhengxiong Luo, Quan Sun, Zhen Li, Yuqi Wang, et al. Multimodal learning with next-token prediction for large multimodal models. *Nature*, 650(8101):327–333, 2026.
- Enze Xie, Junsong Chen, Junyu Chen, Han Cai, Haotian Tang, Yujun Lin, Zhekai Zhang, Muyang Li, Ligeng Zhu, Yao Lu, and Song Han. Sana: Efficient high-resolution image synthesis with linear diffusion transformer. In *ICLR*, 2025.
- An Yang, Anfeng Li, Baosong Yang, Beichen Zhang, Binyuan Hui, Bo Zheng, Bowen Yu, Chang Gao, Chengen Huang, Chenxu Lv, et al. Qwen3 technical report. *arXiv preprint arXiv:2505.09388*, 2025a.
- Zhuoyi Yang, Jiayan Teng, Wendi Zheng, Ming Ding, Shiyu Huang, Jiazheng Xu, Yuanming Yang, Wenqi Hong, Xiaohan Zhang, Guanyu Feng, et al. Cogvideox: Text-to-video diffusion models with an expert transformer. In *ICLR*, 2025b.
- Jiahui Yu, Yuanzhong Xu, Jing Yu Koh, Thang Luong, Gunjan Baid, Zirui Wang, Vijay Vasudevan, Alexander Ku, Yinfei Yang, Burcu Karagol Ayan, Ben Hutchinson, Wei Han, Zarana Parekh, Xin Li, Han Zhang, Jason Baldridge, and Yonghui Wu. Scaling autoregressive models for content-rich text-to-image generation. *TMLR*, 2022a. URL <https://openreview.net/forum?id=AFDcYJKhND>.
- Jiahui Yu, Yuanzhong Xu, Jing Yu Koh, Thang Luong, Gunjan Baid, Zirui Wang, Vijay Vasudevan, Alexander Ku, Yinfei Yang, Burcu Karagol Ayan, et al. Scaling autoregressive models for content-rich text-to-image generation. *TMLR*, 2022b.

- Richard Zhang, Phillip Isola, Alexei A Efros, Eli Shechtman, and Oliver Wang. The unreasonable effectiveness of deep features as a perceptual metric. In *CVPR*, 2018.
- Yuhui Zhang, Brandon McKinzie, Zhe Gan, Vaishaal Shankar, and Alexander T Toshev. Pre-trained language models do not help auto-regressive text-to-image generation. In *EMNLP*, 2024.
- Bing Zhao, Chenfei Wu, Deqing Li, Hao Meng, Jiahao Li, Jie Zhang, Jingren Zhou, Junyang Lin, Kaiyuan Gao, Kuan Cao, et al. Qwen-image-2.0 technical report. *arXiv preprint arXiv:2605.10730*, 2026.
- 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. In *ICLR*, 2025.

APPENDIX

A Training Details	17
B Sampling Details	17
C Samples	18

A TRAINING DETAILS

We provide model architecture and training details for all `Halcyon` models in Table 11.

Table 11: **Halcyon configurations.**

	0.1B	0.25B	0.52B	1.6B	4B
architecture					
backbone	Qwen3 (Yang et al., 2025a), decoder-only				
layers	14	21	28	28	36
hidden dim	512	768	1024	2048	2560
FFN dim	1536	2304	3072	6144	9728
query / KV heads	8 / 4	12 / 6	16 / 8	16 / 8	32 / 8
head dim	128				
non-embedding params	55.1M	185.8M	440.4M	1409.3M	3634M
embedding params	41.3M	61.9M	82.6M	165.2M	206.4M
total params	96.4M	247.7M	523.0M	1574.5M	3840M
vocabulary	80,640 (16,384 text + 64,000 image + 256 special)				
tied embeddings	yes				
image size	256				
image tokens	1,024				
max caption tokens	254				
sequence length	1,280				
training					
training steps	97,632 (1 epoch on GPIC)				
global batch size	1,024				
warmup steps	1,000				
peak learning rate	1e-3	7.5e-4	7.5e-4	7.5e-4	7.5e-4
lr schedule	linear warmup, constant to step 49,595, cooldown to peak/10				
optimizer	AdamW, $\beta_1, \beta_2 = 0.9, 0.95$				
weight decay	0.1				
gradient clipping	1.0				
init std	0.02				
image loss weight λ_{img}	0.5				
caption dropout p_{drop}	0.1				
precision	bf16				
sampling					
cfg	7.5	6.0	6.0	5.0	4.5

B SAMPLING DETAILS

All models generate one 256×256 image per caption for the 50K GPIC Test-50K captions. **Xanthus.** We sample the 1,024 image tokens autoregressively after the caption prefix and `[SOV]`, using temperature 1.0 with top-k and top-p disabled and a fixed random seed. CFG is applied only to image-token logits, $\ell_{\text{cfg}} = \ell_u + \gamma(\ell_c - \ell_u)$, where the unconditional branch uses the `[SOV]` prefix without the caption. For each model size, we select γ to minimize gFD on the GPIC validation set (Figure 8) and use the same scale for all configurations of that size (Tables 5–7 and 12). **Baselines.** For LlamaGen and JiT-T2I, we use the official sampling code and settings and identify the best hyper-parameters using the GPIC validation set.

Table 12: **JiT-T2I configurations.**

	JiT-T2I-XXL/16	JiT-T2I-XXL/16++	JiT-T2I-XXL/8++
<i>language model (frozen)</i>			
text encoder		Qwen3-1.7B	
feature dim		2048	
max caption tokens		300	
<i>image generator</i>			
backbone	JiT-T2I (pixel space, no image tokenizer)		
blocks (image / text)		16 / 4	
hidden dim		1536	
attention heads		24	
patch bottleneck dim		256	
params		1.10B	
patch size	16	16	8
image tokens	256	256	1,024
<i>training</i>			
objective	flow matching, linear schedule		
timestep sampling	logit-normal ($\mu=-0.8, \sigma=0.8$), $t_\epsilon=0.05$		
LPIPS weight	–	0.1	0.1
LPIPS timestep gate	–	$t \geq 0.3$	$t \geq 0.3$
training steps		390,600	
global batch size		256	
peak learning rate		1e-4	
lr schedule		4,000 warmup, constant, 4,000 cooldown	
optimizer		AdamW, $\beta_1, \beta_2 = 0.9, 0.95$	
weight decay		0	
gradient clipping		1.0	
EMA decay		0.9999	
caption dropout		0.1	
precision		bf16	
<i>sampling</i>			
sampler	Euler (ODE), 50 steps, EMA weights		
CFG scale	6.00	5.00	4.00

C SAMPLES

We compare generated samples across the four GPIC caption types: tag (Figure 11), short (Figure 12), medium (Figure 13), and long (Figure 14). Each row shows JiT-T2I-XXL/8++, LlamaGen-XL/8, and Xanthus from 0.1B to 4B for the same caption, sampled with the settings in Appendix B. Sample quality and prompt adherence improve as *Halcyon* scales, with larger models capturing finer details such as clothing, object attributes, and scene layout. Rendering legible text remains challenging for all models at 256×256 resolution.

Table 13: LlamaGen configurations.

	LlamaGen-XL	LlamaGen-XL/8	LlamaGen-XL-T5-Base/8
<i>language model (frozen)</i>			
text encoder	FLAN-T5-XL	FLAN-T5-XL	FLAN-T5-Base
encoder params	1.22B	1.22B	0.11B
feature dim	2048	2048	768
max caption tokens		300	
<i>image tokenizer</i>			
tokenizer	VQ-16 (pretrained)	Ours (GPIC-trained)	
spatial downsampling	16×		8×
codebook size	16,384		64,000
image tokens	256		1,024
<i>image generator</i>			
backbone		GPT-XL (Llama-style, decoder-only)	
layers / hidden dim / heads		36 / 1280 / 20	
FFN dim		3584 (SwiGLU)	
params	0.78B	0.90B	0.90B
sequence length	556	1,324	1,324
<i>training</i>			
training steps	390,600	390,600	390,600
global batch size		256	
peak learning rate		1e-4	
lr schedule		3,906 warmup, then constant	
optimizer		AdamW, $\beta_1, \beta_2 = 0.9, 0.95$	
weight decay		0.05	
gradient clipping		1.0	
dropout / token dropout		0.1 / 0.1	
caption dropout		0.1	
precision		bf16	
<i>sampling</i>			
CFG scale	6.5	6.5	5.5
temperature / top- k / top- p	1.0 / 1000 / off		1.0 / off / off

Table 14: Tokenizer training configurations.

Image tokenizer	
architecture	Cosmos (Agarwal et al., 2025) encoder-decoder, FSQ (Mentzer et al., 2024)
decoder width	192 (128 in Cosmos)
spatial downsampling	8×
codebook size	64,000 (FSQ levels [8, 8, 8, 5, 5, 5])
Params	136M
tokens per image	1,024 (32 × 32, raster order)
losses (stage 1)	reconstruction (1.0) + LPIPS (0.1)
losses (stage 2)	+ PatchGAN adversarial (5e-3)
training data	GPIC train images, 1.5 epochs (1.0 stage 1 / 0.5 stage 2)
iterations	242,172
batch size	640
learning rate	1e-4, constant
optimizer	Adam, $\beta_1, \beta_2 = 0.9, 0.999$
weight decay	0.001
gradient clipping	1.0
EMA	yes
hardware	1 × 8 × H100 node
Text tokenizer	
algorithm	byte-level BPE (Sennrich et al., 2016)
vocabulary size	16,384
training data	GPIC train captions (21.3 GB) + OpenWebText (11.9 GB)
caption length	mean 44.5, p50 36, p90 64, p99 182 tokens

Table 15: **Total pretraining compute of T2I systems.** We report total pretraining FLOPs, decomposed into language, image-generator, and image-tokenizer compute, for all configurations shown in Fig. 1. Vision compute includes both image-generator and image-tokenizer pretraining. For late-fusion methods, language FLOPs include pretraining of the language model. For Halcyon, language FLOPs correspond to the text-token portion of joint multimodal pretraining, without external language pretraining. All image generators are pretrained on GPIC for one epoch at 256×256 resolution. Generation Fréchet distance (gFD) is evaluated on GPIC Test-50K; lower is better.

Method	Total FLOPs	Language		Vision			gFD ↓
		FLOPs	(%)	Generator FLOPs	Tokenizer FLOPs	(%)	
LlamaGen-XL	9.32×10^{21}	9.06×10^{21}	97.17	1.19×10^{20}	1.45×10^{20}	2.83	147.17
LlamaGen-XL/8	9.68×10^{21}	9.06×10^{21}	93.58	4.76×10^{20}	1.45×10^{20}	6.42	63.52
LlamaGen-XL-T5-Base/8	1.29×10^{21}	6.69×10^{20}	51.87	4.76×10^{20}	1.45×10^{20}	48.13	66.24
JiT-T2I-XXL/16	3.67×10^{23}	3.67×10^{23}	99.95	1.69×10^{20}	0	0.05	68.67
JiT-T2I-XXL/16++	3.67×10^{23}	3.67×10^{23}	99.95	1.69×10^{20}	0	0.05	55.70
JiT-T2I-XXL/8++	3.68×10^{23}	3.67×10^{23}	99.82	6.76×10^{20}	0	0.18	43.22
Halcyon-0.1B	2.05×10^{20}	1.47×10^{19}	7.15	5.93×10^{19}	1.31×10^{20}	92.85	144.16
Halcyon-0.25B	3.21×10^{20}	3.78×10^{19}	11.74	1.53×10^{20}	1.31×10^{20}	88.26	72.55
Halcyon-0.52B	5.33×10^{20}	7.97×10^{19}	14.96	3.22×10^{20}	1.31×10^{20}	85.04	48.02
Halcyon-1.6B	1.34×10^{21}	2.40×10^{20}	17.90	9.69×10^{20}	1.31×10^{20}	82.10	35.28
Halcyon-4B	3.08×10^{21}	5.85×10^{20}	19.00	2.36×10^{21}	1.31×10^{20}	81.00	28.43

Table 16: **Classifier-free guidance (CFG) scales for T2I evaluation.** We report the CFG scale selected for each model based on the lowest gFD on GPIC Validation-50K. Sweeps are shown in Figures 8 and 9.

Model	p_{drop}	η	CFG Scale
LlamaGen-XL	0.1	—	6.5
LlamaGen-XL/8	0.1	—	6.5
LlamaGen-XL-T5-Base/8	0.1	—	5.5
JiT-T2I-XXL/16	0.1	—	6.0
JiT-T2I-XXL/16 ++	0.1	—	5.0
JiT-T2I-XXL/8 ++	0.1	—	4.0
Halcyon-0.10B	0.1	—	7.5
	0.0	0.1	7.5
Halcyon-0.25B	0.1	—	6.0
	0.0	0.1	6.0
Halcyon-0.52B	0.1	—	6.0
	0.0	0.1	6.0
Halcyon-1.6B	0.1	—	5.0
	0.0	0.1	5.2
Halcyon-4B	0.1	—	4.5
	0.0	0.1	4.6




Figure 9: **Classifier-free guidance (CFG) sweeps for late-fusion T2I baselines.** We evaluate LlamaGen variants (left) and JiT-T2I variants (right) across CFG scales on GPIC Validation-50K. Stars indicate the scale achieving the lowest gFD for each configuration, with the corresponding gFD and CFG scale reported in the figure. The selected CFG scales are used for all subsequent evaluations on GPIC Test-50K, GenEval, and DPG-Bench.




Figure 10: **Effect of language pretraining on validation NLL.** Validation NLL for image tokens conditioned on text ($\mathcal{I} \mid \mathcal{T}$, left) and text tokens without conditioning ($\mathcal{T} \mid \emptyset$, right) during training of Qwen3-0.6B with random versus language-pretrained initialization. Random initialization consistently achieves lower image NLL, while language-pretrained initialization achieves lower text NLL.




Figure 11: **Comparison of generated images with “tag” captions.** From left to right: JiT-T2I-XXL/8++, LlamaGen-XL/8, and Halcyon from 0.1B to 4B. Image quality improves as Halcyon scales up.




Figure 12: **Comparison of generated images with “short” captions.** From left to right: JiT-T2I-XXL/8++, LlamaGen-XL/8, and Halcyon from 0.1B to 4B. Image quality improves as Halcyon scales up.




Figure 13: **Comparison of generated images with “medium” captions.** From left to right: JiT-T2I-XXL/8++, LlamaGen-XL/8, and Halcyon from 0.1B to 4B. Image quality improves as Halcyon scales up.




Figure 14: **Comparison of generated images with “long” captions.** From left to right: JiT-T2I-XXL/8+,, LlamaGen-XL/8, and Halcyon from 0.1B to 4B. Image quality improves as Halcyon scales up.