

UniT: Unified Multimodal Chain-of-Thought Test-time Scaling

Leon Liangyu Chen^{1,2} Haoyu Ma² Zhipeng Fan² Ziqi Huang^{2,3} Animesh Sinha² Xiaoliang Dai²
 Jialiang Wang² Zecheng He² Jianwei Yang² Chunyuan Li² Junzhe Sun² Chu Wang²
 Serena Yeung-Levy¹ Felix Juefei-Xu²

¹Stanford University ²Meta Superintelligence Labs ³Nanyang Technological University

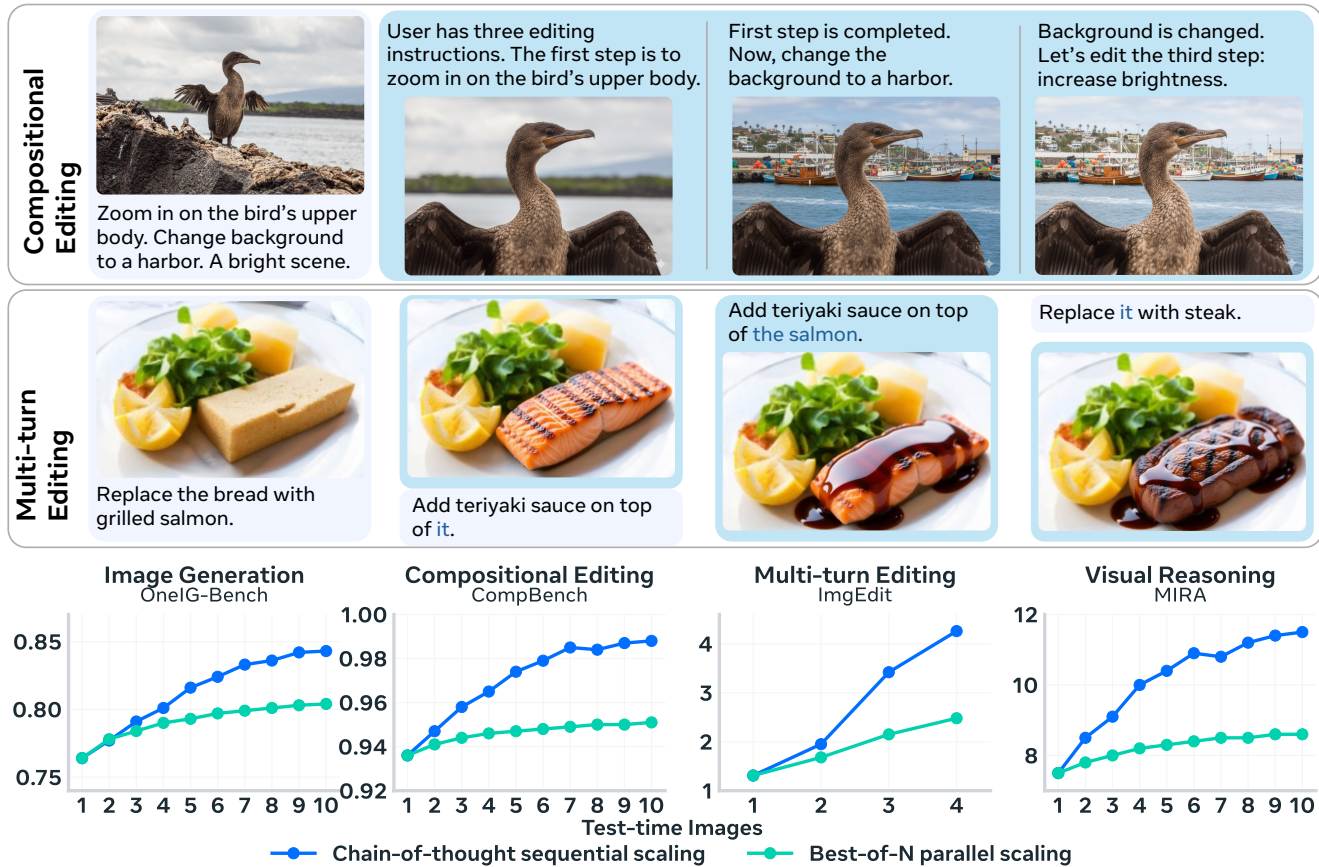


Figure 1. **Multimodal chain-of-thought enables test-time scaling through emergent cognitive behaviors.** We propose the UniT framework for unified multimodal models, which induces subgoal decomposition for compositional tasks and unlocks content understanding and memory for multi-turn editing. Controlling the number of test-time images, chain-of-thought sequential scaling outperforms best-of-N parallel scaling across generation and reasoning benchmarks. User input Model output

Abstract

Unified models can handle both multimodal understanding and generation within a single architecture, yet they typically operate in a single pass without iteratively refining their outputs. Many multimodal tasks, especially those involving complex spatial compositions, multiple interacting objects, or evolving instructions, require decomposing instructions, verifying intermediate results, and making itera-

tive corrections. While test-time scaling (TTS) has demonstrated that allocating additional inference compute for iterative reasoning substantially improves language model performance, extending this paradigm to unified multimodal models remains an open challenge.

We introduce UniT, a framework for multimodal chain-of-thought test-time scaling that enables a single unified model to reason, verify, and refine across multiple rounds. UniT combines agentic data synthesis, unified model train-

ing, and flexible test-time inference to elicit cognitive behaviors including verification, subgoal decomposition, and content memory. Our key findings are: (1) unified models trained on short reasoning trajectories generalize to longer inference chains at test time; (2) sequential chain-of-thought reasoning provides a more scalable and compute-efficient TTS strategy than parallel sampling; (3) training on generation and editing trajectories improves out-of-distribution visual reasoning. These results establish multimodal test-time scaling as an effective paradigm for advancing both generation and understanding in unified models.

1. Introduction

Unified multimodal models [15, 66, 80] aim to merge vision, language, and more modalities into a single architecture capable of both understanding and generation. Unlike modular pipelines where separate models handle perception, verification, and generation, unified models handle all modalities within one coherent conversation, enabling richer cross-modal grounding, continuous contextual tracking, and seamless interleaving of understanding and generation. However, despite this potential, existing unified models still operate mostly in a *single-pass* mode: they produce an output once, without explicit mechanisms for evaluating, reflecting on, or refining their predictions. This limitation becomes fundamental for tasks that intrinsically require multi-step reasoning and self-correction, such as compositional generation, multi-turn editing, and complex visual reasoning - settings where both humans and AI naturally benefit from extended reasoning.

Recent advances in language models have demonstrated that test-time scaling (TTS)—allocating additional computational resources during inference through extended chain-of-thought reasoning, verification, and iterative refinement [13, 46, 55]—enables substantial performance gains on complex reasoning tasks in mathematics [12], coding [7], and logic [57]. Early work on multimodal chain-of-thought has shown similar benefits for single-round visual understanding and generation [10, 18, 26, 53, 69]. Yet, test-time scaling for unified multimodal models remains largely unexplored. The challenge is nontrivial: test-time scaling requires capabilities that currently scatter across specialized models (image generation models for generation, vision-language models for verification, image editing models for refinement). Bridging this gap requires a unified framework that systematically integrates data synthesis, model training, and inference mechanisms for multimodal test-time scaling. This motivates the central question:

How to enable scalable multimodal inference that allows unified models to iteratively generate, reflect, and refine?

We introduce **UniT**, a unified framework for multimodal chain-of-thought test-time scaling. Scalable multimodal inference requires the tight integration of three components: (i) **Agentic data synthesis** to induce cognitive behaviors through multi-round trajectories. We develop an automated pipeline (Fig. 2) where vision-language models iteratively critique and image editing models refine generated images with explicit chain-of-thought reasoning. This naturally produces training data exhibiting three critical cognitive behaviors [19]: **verification**—evaluating outputs against instructions; **subgoal decomposition**—breaking complex instructions into sequential planning steps; **content memory**—maintaining understanding of visual content across rounds through unified multimodal context. (ii) **Unified model training** to enable the model to internalize multimodal reasoning patterns. We collect approximately 12K multi-round trajectories and fine-tune the Bagel unified multimodal model [15] for 700 H100 hours, enabling it to perform both understanding and refinement without switching models. (iii) **Multimodal test-time scaling** at inference with flexible computational budget. The trained model performs all reasoning, generation, and refinement iteratively through explicit multimodal chain-of-thought thinking, allocating more rounds to more challenging tasks.

The synergy of these components enables the model to act as a single, coherent multimodal reasoner capable of self-evaluation and iterative improvement. The UniT framework exhibits strong test-time scaling behavior (Fig. 1) with emergent capabilities. Most notably, *models trained on shorter reasoning trajectories (averaging 3.6 rounds) effectively generalize to longer inference chains at test time (averaging 4.7 rounds)* (Fig. 4), echoing patterns previously seen only in text-only models [55]. Furthermore, chain-of-thought sequential scaling substantially outperforms best-of-N parallel sampling, achieving comparable performance with $2.5\times$ less computational cost (Fig. 1). This demonstrates that iterative refinement with explicit reasoning provides more efficient use of inference compute than parallel sampling. Critically, UniT achieves 5.56% improvement on CompBench multi-object editing, 2.95 human preference scores on ImgEdit multi-turn editing, and 10.34% on OneIG instruction following compared to single-pass generation. Moreover, it improves out-of-distribution visual reasoning on MIRA by 53.33%, *establishing multimodal chain-of-thought test-time scaling as a unified paradigm that benefits both generation and comprehension.*

We summarize our contributions as follows:

- **Unified multimodal test-time scaling.** We propose **UniT**, a unified framework for multimodal chain-of-thought test-time scaling, integrating agentic data synthesis, unified model training, and test-time scaling mechanisms.
- **Emergent extrapolation to longer reasoning chains.**

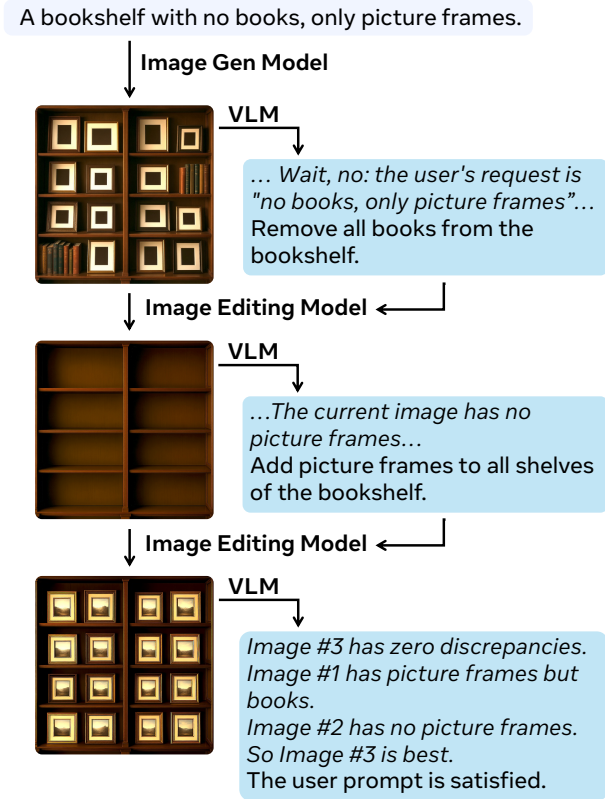


Figure 2. **Agentic framework for synthesizing chain-of-thought training data.** Starting from a user prompt, an image generation model generates an initial image. A vision-language model then performs verification - evaluating whether the output satisfies the prompt. When unsatisfactory, the VLM engages in explicit sub-goal decomposition through thinking tokens, planning concrete improvements, and rewriting editing instructions. This iterative loop continues until verification succeeds, generating multi-turn reasoning trajectories that teach unified models to refine outputs through test-time computation. The explicit reasoning traces of the three models capture how cognitive behaviors emerge from the interplay between generation, verification, and planning.

We demonstrate that models trained on shorter trajectories generalize to longer inference chains at test time, extrapolating beyond the training distribution.

- **Broad improvements across multimodal tasks.** UniT achieves substantial gains on compositional generation/editing, multi-turn editing, and visual reasoning, establishing chain-of-thought test-time scaling as a unified paradigm for both generation and understanding tasks.

2. Related Works

Test-time scaling. Test-time scaling allocates additional computation during inference to improve model performance. We distinguish two primary approaches: parallel and sequential methods. Parallel scaling generates multi-

ple independent candidates and selects the best via criteria such as best-of-N sampling [4, 34] or majority voting [29], typically guided by outcome reward models [1, 70]. Sequential scaling [24, 33, 56] enables iterative refinement where models critique and improve outputs across multiple rounds. Self-refinement approaches [43] exemplify this by having models explicitly reason about deficiencies and produce progressively better solutions. Tree-based methods such as Monte Carlo Tree Search [11, 37, 75, 78] and REBASE [68] occupy a middle ground, using process reward models [35, 62, 64] to guide structured search. Recent breakthroughs including o1 [46] and DeepSeek-R1 [13] demonstrate that reinforcement learning enables effective utilization of extended inference computation. Budget forcing [45, 56] trains models to produce reasoning chains of controllable cost by varying computational budgets during training. While most test-time scaling research focuses on text-only reasoning, recent work has explored search-based methods for image and video generation [23, 36]. However, unified multimodal test-time scaling interleaving text and images remains largely unexplored, which we address in this work.

Unified multimodal models. Unified models that jointly handle understanding and generation have attracted substantial interest. Autoregressive approaches [9, 40, 50, 59, 63, 66] extend next-token prediction to both text and discrete image tokens. Additional diffusion methods [16, 48, 60, 61] augment language models with external diffusion modules for image generation. Unified integrated transformers [15, 41, 54, 72, 79] deeply integrate language modeling and diffusion within single architectures. Our work builds upon Bagel [15], pretrained on large-scale interleaved text-image sequences. Our framework generalizes to all three paradigms as they naturally handle interleaved multimodal inputs and outputs.

Multimodal chain-of-thought. Chain-of-thought reasoning [65] has proven effective for text-based problem-solving, motivating extensions to multimodal tasks. Visual chain-of-thought methods [17, 25, 27, 39, 53, 76] incorporate visual representations into reasoning steps for multimodal understanding. Recent work explores interleaved reasoning [21, 26] across text and visual modalities. Uni-CoT [49] further demonstrates a unified model that couples macro- and micro-level reasoning for vision-language understanding, but it does not study compute scaling or iterative editing. In text-to-image generation, studies investigate whether explicit reasoning improves generation quality [15, 18, 22, 31, 69]. Reflection-based approaches [10, 67, 82] iteratively critique and refine generated images. Our work differs by focusing on both semantic correctness and visual quality through test-time scaled refinement, while demonstrating that multimodal chain-of-thought benefits both generation and understanding tasks as

a unified paradigm.

3. Method

We extend test-time compute scaling from text-only reasoning to unified multimodal models. As illustrated in Fig. 2, we develop an agentic framework to automatically collect chain-of-thought training data, then control test-time computational budget through iterative refinement rounds. Sequential chain-of-thought scaling outperforms parallel best-of-N sampling (Fig. 1) while inducing emergent cognitive behaviors—verification and subgoal decomposition.

Key distinction. The multi-model agentic framework described in Sec. 3.1 is used *solely for synthesizing training data*. At inference time (Sec. 3.3), we use only the single unified BAGEL model [15], which performs all planning, generation, reflection, and refinement operations without external models.

3.1. Multimodal Chain-of-Thought Data

Agentic data collection pipeline. Our automated pipeline synthesizes multimodal chain-of-thought trajectories through iterative reflection-editing (Fig. 2):

1. **Prompt generation:** Llama-4-Scout-17B-16E [44] generates 20K diverse prompts covering compositional attributes, spatial relations, and complex multimodal tasks based on [58].
2. **Initial generation:** Flux Pro [3] produces initial images from prompts. For complex prompts, the VLM (Qwen3-VL) decomposes the prompts into subgoals and executes the first step in initial generation.
3. **Reflection:** Qwen3-VL [51] evaluates whether the image satisfies the prompt. If not, it generates explicit chain-of-thought reasoning, identifying deficiencies, planning improvements, and specifying editing instructions.
4. **Refinement:** Flux Kontext [32] or Qwen-Image-Edit [52] applies editing instructions.
5. **Iteration:** Steps 3-4 repeat until the VLM determines the output satisfies the prompt.

Cognitive behaviors. As demonstrated in Fig. 1, 2, this agentic pipeline naturally induces three critical cognitive behaviors: (i) **verification**—VLMs evaluate outputs against specifications to determine when refinement is needed; (ii) **subgoal decomposition**—complex compositional tasks are planned in sequential editing steps; (iii) **content memory**—the model maintains understanding of image content across refinement rounds through unified multimodal context.

Data filtering. We apply quality filtering to ensure training efficiency:

- **Length constraint:** Trajectories exceeding 8 rounds are removed to balance efficiency with reasoning depth.

Method	Alignment \uparrow					
	NP	T&P	Short	Medium	Long	Overall
Janus-Pro [8]	0.557	0.533	0.609	0.548	0.515	0.552
BLIP3-o [6]	0.719	0.671	0.754	0.712	0.674	0.706
Bagel [14]	0.776	0.734	0.782	0.769	0.759	0.764
Bagel+CoT [14]	0.798	0.767	0.824	0.793	0.767	0.790
UniT	0.853	0.844	0.859	0.849	0.812	0.843

Table 1. **Compositional generation, OneIG-Bench.** NP denotes the natural language prompt. T&P denotes the tag-based and phrase-based prompt. Short, Medium and Long represent the length of the prompts, where Short denotes the number of words is less than 30, Medium denotes the number between 30 and 60, and Long denotes the number exceeding 60. Bagel+CoT indicates Bagel with text-only chain-of-thought.

- **Quality regression:** Trajectories where the final image has worse instruction-following quality than any of the first three images (measured by Qwen3-VL) are removed.
- **Relevance filtering:** Individual rounds with editing prompts semantically irrelevant to the original task (measured by Llama-4-Scout-17B-16E) are removed.
- **Minimal visual changes:** Rounds with perceptual editing distance below LPIPS < 0.03 [74] between consecutive images are removed.
- **Benchmark deduplication:** Training prompts are deduplicated from evaluation benchmarks using 5-gram matching to prevent data leakage.

This filtering retains 12k high-quality trajectories for training.

3.2. Training and Inference

Training. We use Bagel [15], a unified multimodal architecture with understanding and generation capabilities, trained on the dataset from Sec. 3.1 for 700 H100 hours. To simulate user prompts for multi-turn editing, 10% of intermediate image editing instructions don’t require losses.

Inference. We adopt a framework that incorporates two complementary classifier-free guidance (CFG) schemes applied in a nested manner: (1) text CFG, conditioning with versus without the current text instruction, and (2) image CFG, conditioning with versus without all images in the generation history (including both the initial input image if exists and previously generated outputs). Formally, let v_t denote the fully conditional prediction, $v_{t,\text{unc}}$ the text-unconditional prediction, and $v_{i,\text{unc}}$ the image-unconditional prediction. We apply CFG sequentially: first text guidance $v_{\text{text}} = v_{t,\text{unc}} + s_t(v_t - v_{t,\text{unc}})$, then image guidance $v_{\text{final}} = v_{i,\text{unc}} + s_i(v_{\text{text}} - v_{i,\text{unc}})$, with scales $s_t=4.0$ and $s_i=2.0$. The nested application—where image guidance is applied on top of the text-guided prediction—enables independent control over prompt adherence and visual consistency. This strategy helps maintain strong alignment with text instructions while preserving structural

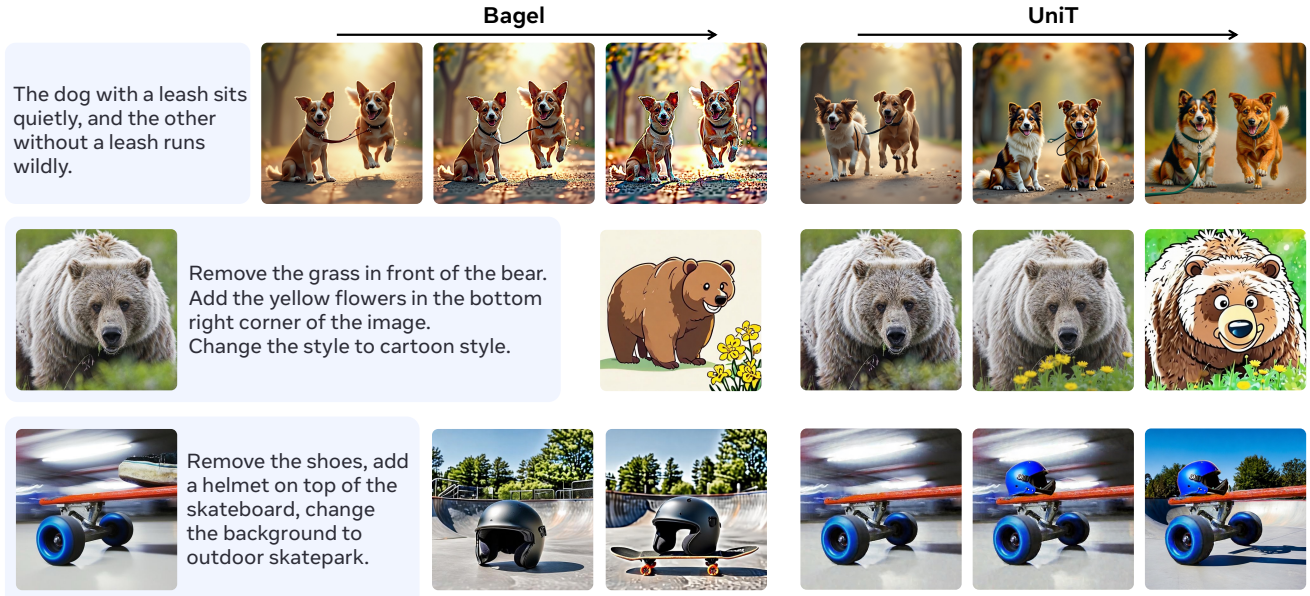


Figure 3. **UniT enables iterative refinement for compositional instructions through multimodal chain-of-thought reasoning.** UniT exhibits: (i) **error verification and correction**—identifying and fixing constraint violations that Bagel misses (top: correcting leash placement and dog action); (ii) **subgoal decomposition with subject consistency**—sequentially addressing instructions while maintaining subject identity across rounds (middle: preserving bear features through style transformation, bottom: skateboard consistency); (iii) **quality preservation**—maintaining visual fidelity through iterative refinement rather than degradation (top: reduced artifacts and halving).

coherence across multi-turn editing sequences, with notable benefits for generation quality and consistency in iterative refinement workflows.

The original untrained Bagel model can also be forced to reason chain-of-thought with our inference code. However, the image quality degrades quickly as context images scale and the Bagel model can’t verify its image outputs properly. It usually hallucinates visual content according to user prompts. Thus the untrained Bagel model is not feasible for practical multimodal chain-of-thought reasoning. Training is required to enable effective multimodal chain-of-thought reasoning.

3.3. Budget Forcing for Test-time Scaling

We adapt budget forcing [45] from text-only to multimodal test-time scaling. While text-based methods control reasoning tokens, we control image generation rounds, which dominate inference latency. At test time, the unified BAGEL model performs all operations autonomously—planning, generation, reflection, and refinement—without relying on external models.

Controlling computational budget. At inference, we specify computational budget C as the number of image generation rounds. Each round consists of textual chain-of-thought reasoning followed by image generation or editing. To enforce budget C :

- **Forcing extended reasoning:** If the model terminates be-




Figure 4. **Training vs. inference round distribution demonstrates beyond-training generalization.** The model is trained on trajectories averaging 3.6 refinement rounds, but effectively generalizes to longer inference chains averaging 4.7 rounds at test time. This distribution shift reveals the model’s emergent ability to extend inference beyond its training distribution, a key property of effective test-time scaling.

fore C rounds, we suppress EOS, append “Let’s edit the image”, wait for reasoning completion, then force image generation.

- **Budget constraint:** If the model generates more than C images, we use only the final image from round C .

This enables studying sequential chain-of-thought scaling (iterative refinement building on previous outputs) versus parallel best-of- N scaling (generating N independent images and selecting the best). We present detailed scaling analysis in Sec. 5.1.

Beyond-training generalization. Models trained on trajec-

Method	Content Memory	Content Understand	Version Backtrack	Overall
Bagel [14]	1.76	1.34	0.82	1.31
Bagel+CoT [14]	2.24	2.67	0.84	1.92
UniT	4.29	5.02	3.48	4.26

Table 2. **Multi-turn editing, ImgEdit.** Human evaluation score from 0-10, normalized over 3 turns of editing.

stories averaging 3.6 rounds generalize to longer inference chains averaging 4.7 rounds at test time (Fig. 4), exhibiting problem-solving capabilities beyond their training distribution. This establishes test-time compute as a general paradigm for unified multimodal models.

4. Experiments

We evaluate on text-to-image generation, compositional editing, multi-turn editing, and visual reasoning benchmarks. Chain-of-thought TTS achieves substantial gains across both generation and understanding tasks.

4.1. Experiment Settings

Evaluation setup. We evaluate with computational budgets $C=1$ to $C=10$ (maximum due to GPU memory), controlled via Sec. 3.3. For ImgEdit multi-turn editing, we apply C refinement rounds independently to each of three sequential turns, with maximum $C=4$ per turn due to memory constraints.

Baselines. We compare: (1) **Bagel**, base model without chain-of-thought; (2) **Bagel+CoT**, textual thinking only; (3) **UniT**, our full multimodal chain-of-thought with interleaved text and image reasoning. We report $C=10$ rounds ($C=4$ for ImgEdit) unless specified otherwise.

Human evaluation protocol. For ImgEdit multi-turn editing, 3 expert annotators with computer vision backgrounds independently rate outputs on a 0-10 scale across three criteria: content memory (tracking edits across turns), content understanding (correctly interpreting instructions), and version backtracking (maintaining coherence). We evaluate 100 randomly sampled test examples per method. Inter-annotator agreement is high (Krippendorff’s $\alpha = 0.82$), and final scores are averaged across annotators.

4.2. Compositional Generation and Editing

We evaluate on OneIG-Bench-EN [5] for compositional generation (Table 1), achieving 10.34% improvement over base model at $C=10$. On CompBench [30] multi-object compositional editing subset (Table 3), we achieve 5.56% improvement from $C=1$ to $C=10$. Fig. 1 demonstrates that iterative refinement with explicit reasoning enables better compositional understanding and generation.

Method	Foreground		Background			Overall \uparrow
	LC-T \uparrow	LC-I \uparrow	PSNR(dB) \uparrow	SSIM \uparrow	LPIPS \downarrow	
HQ-Edit [28]	19.163	0.757	12.987	0.412	0.421	0.007
UltraEdit [77]	20.022	0.795	22.326	0.719	0.164	0.609
AnyEdit [73]	19.875	0.809	22.789	0.697	0.129	0.622
SEED-X [20]	19.092	0.795	20.638	0.788	0.138	0.595
GoT [18]	19.919	0.804	21.296	0.826	0.127	0.668
StepIX-Edit [38]	20.213	0.828	22.696	0.873	0.089	0.844
FLUX.1 Kontext [32]	20.983	0.836	24.013	0.938	0.064	0.965
Qwen-Image-Edit [52]	21.058	0.836	21.927	0.810	0.121	0.782
Bagel [14]	20.434	0.842	24.370	0.917	0.069	0.936
Bagel+CoT [14]	20.658	0.846	24.691	0.926	0.065	0.956
UniT	21.127	0.854	25.442	0.942	0.055	0.988

Table 3. **Multi-object editing, CompBench.** LC-T denotes local CLIP scores between the edited foreground and the local description. LC-I refers to the CLIP image similarity between the foreground edited result and ground truth (GT) image. Overall scores are computed using min-max normalization for each metric.

4.3. Multi-Turn Editing

We evaluate on ImgEdit [71] multi-turn editing subset (Table 2), achieving 225.19% improvement from $C=1$ to $C=4$. This demonstrates that maintained context and reasoning chains are critical for multi-turn interactions. The model’s content memory through unified multimodal context enables coherent interactions as computational budget increases (Fig. 1).

4.4. Visual Reasoning

On MIRA [81] for out-of-distribution visual reasoning (Table 4), we achieve 53.33% improvement from $C=1$ to $C=10$ (Fig. 1). The remaining gap to frontier models (GPT-5, Qwen2.5-VL-72B) reflects base model capability differences: these models benefit from significantly larger scale and proprietary training data. Our key contribution is methodological, demonstrating that TTS successfully transfers to multimodal domains; as base unified models improve, the UniT framework directly benefits. This establishes multimodal chain-of-thought test-time scaling as a unified paradigm enhancing both generation and comprehension tasks. The cognitive behaviors induced by our agentic framework (Fig. 2)—verification, subgoal decomposition, and content memory—transfer from generation to understanding tasks.

4.5. Qualitative Results

Fig. 1 and Fig. 3 demonstrate how multimodal chain-of-thought test-time scaling progressively refines outputs through iterative reasoning, revealing the three cognitive behaviors—verification, subgoal decomposition, and content memory.

Compositional generation. For complex compositional prompts, Round 1 produces partial solutions with compositional errors—missing objects, incorrect attributes, or violated spatial constraints. The model’s explicit reasoning identifies deficiencies, and through subgoal decomposition,



Figure 5. **Chain-of-thought visual reasoning on MIRA.** The model decomposes the puzzle into subgoals (zoom in, identify patterns) before selecting the matching piece, demonstrating cognitive behaviors transferring from generation to understanding tasks.

breaks corrections into sequential steps. By Round 10, the output achieves precise alignment with all requirements, demonstrating systematic error correction through explicit reasoning.

Visual reasoning. For MIRA geometry tasks, the model’s chain-of-thought reveals how verification supports iterative problem-solving. Early rounds may produce incorrect analyses, but the reasoning shows self-critique, allowing identification of flaws and revision in subsequent rounds. Subgoal decomposition breaks complex reasoning into steps. These examples demonstrate that cognitive behaviors transfer from generation to understanding tasks.

5. Discussion

We analyze factors contributing to effective multimodal test-time scaling: sequential versus parallel scaling, cognitive behaviors, and data quality.

Experimental protocol. Unless specified otherwise, we use computational budget $C=10$ ($C=4$ for ImgEdit) and report: alignment score for OneIG-Bench, overall normalized score for CompBench, human evaluation score (0-10) for ImgEdit, and accuracy for MIRA.

5.1. Sequential vs. Parallel Scaling

We compare chain-of-thought sequential scaling against best-of- N parallel scaling. Following [45], sequential scaling builds on intermediate results through iterative refinement, while parallel scaling generates outputs independently.

Setup. For both approaches, we control the number of generated images ($C=N$) from 1 to 10. Sequential scaling uses budget forcing (Sec. 3.3). Parallel scaling generates N independent images and selects the best via HPSv3 [42].

Compute accounting. We use the number of generated images as our compute metric ($C=N$), which accurately reflects computational cost because: (i) image generation via diffusion models dominates wall-clock time—as noted in

Sec. 3.3, image generation rounds control inference latency; (ii) marginal text tokens from VLM reflection have negligible effect on latency compared to diffusion sampling; (iii) computational cost scales linearly with the number of images generated. We exclude selection costs (HPSv3 for parallel scaling, VLM verification for sequential scaling) from our comparison as these represent arbitrary model choices rather than fundamental algorithmic requirements. Under this metric, sequential scaling achieves comparable performance to parallel with 2.5× fewer generated images (e.g., $C=4$ sequential matches $N=10$ parallel on OneIG-Bench).

Results. Sequential scaling consistently outperforms parallel scaling across all tasks (Fig. 1). At $C=10$ ($C=4$ for ImgEdit), sequential achieves 4.85% improvement over parallel on OneIG-Bench, 3.89% on CompBench, 71.77% on ImgEdit, and 33.72% on MIRA. This advantage manifests through:

- **Steeper scaling slopes:** Sequential achieves larger performance improvements per additional image, indicating more efficient test-time compute use.
- **Sustained improvements:** Sequential shows continued gains up to $C=10$, while parallel plateaus after a few samples.

Iterative refinement with explicit chain-of-thought reasoning provides more effective test-time compute scaling than independent sampling. The advantage stems from sequential scaling accumulating successful edits and learning from previous iterations: each round builds on prior images with explicit CoT corrections, leveraging expanded textual context (reflections, plans). Parallel scaling generates independent samples without inter-sample learning, explaining why it plateaus earlier—it cannot systematically refine toward the target.

Latency considerations. Sequential and parallel scaling serve complementary use cases: sequential optimizes performance while parallel optimizes latency. Sequential scaling also benefits from unique acceleration techniques including speculative decoding, KV-cache reuse across rounds, and early stopping when the model determines satisfaction, which can significantly reduce the latency gap in practice.

Unified model vs. modular pipeline. At inference, UniT performs all reasoning, generation, and refinement within a single model, unlike the multi-model agentic pipeline used for data synthesis (Sec. 3.1). While the teacher pipeline (Flux Pro + Qwen3-VL) scores slightly higher due to frontier-scale components, UniT offers faster inference by eliminating inter-model communication overhead, seamless multimodal context within a single architecture, and practical single-model deployability.

Method	EG↑	PBR↑	ASLP↑	CT↑	Overall↑
GPT-5 [47]	14.5	29.9	10.8	17.9	16.5
Qwen2.5-VL (72B) [2]	14.5	21.7	11.1	8.6	13.1
Bagel [14]	9.7	7.9	3.5	12.3	7.5
Bagel+CoT [14]	8.0	9.5	4.8	14.2	9.2
UniT	12.5	11.2	6.1	16.3	11.5

Table 4. **Multimodal reasoning, MIRA**, with direct input. We report results across four reasoning categories: EG (Geometry), PBR (Physics), ASLP (Puzzles), and CT (Causal), along with the overall average score.

Configuration	OneIG Align (%)	CompBench Overall (%)	ImgEdit Score	MIRA Acc (%)
All behaviours	84.3	98.8	4.26	11.5
w/o Verification	81.2 (-3.1)	96.8 (-2.0)	3.55 (-0.71)	9.6 (-1.9)
w/o Subgoal Decomp.	80.5 (-3.8)	96.3 (-2.5)	3.75 (-0.51)	10.3 (-1.2)
w/o Content Memory	82.8 (-1.5)	97.8 (-1.0)	2.45 (-1.81)	10.8 (-0.7)

Table 5. **Cognitive behavior ablation**. Impact of removing verification, subgoal decomposition, or content memory from our agentic framework.

5.2. Ablation on Cognitive Behaviors

Setup. We train three ablated models, each removing one cognitive behavior: (1) **w/o Verification**, reflection stages do not evaluate quality; (2) **w/o Subgoal Decomposition**, planning stages removed; (3) **w/o Content Memory**, multimodal context not maintained across rounds.

Results. Table 5 shows task-specific sensitivities. Removing subgoal decomposition particularly hurts compositional tasks (3.8% and 2.5% drops on OneIG-Bench and CompBench), demonstrating its importance for planning multi-step generations. Most critically, removing content memory devastates multi-turn editing with 42.5% relative drop on ImgEdit (from 4.26 to 2.45), while minimally impacting single-turn tasks (1.0-1.5% drops). Removing verification most significantly impacts visual reasoning (1.9% drop on MIRA), highlighting its importance for validating reasoning steps. These patterns confirm each cognitive behavior serves a distinct functional role.

5.3. Data Quality Analysis

We analyze the impact of data quality by ablating individual filters from our curation pipeline (Sec. 3.1).

Setup. We train three ablated models, each removing one quality filter: (1) **w/o Quality regression filter**, including trajectories where refinement degrades quality; (2) **w/o Relevance filtering**, including rounds with semantically irrelevant editing prompts; (3) **w/o Minimal visual changes filter**, including rounds with negligible visual refinement.

Results. Table 6 shows task-specific sensitivities. Removing relevance filtering causes the largest degradation on compositional tasks (3.1% on OneIG-Bench, 2.5%

Data Configuration	OneIG Align (%)	CompBench Overall (%)	ImgEdit Score	MIRA Acc (%)
Full curated dataset	84.3	98.8	4.26	11.5
w/o Quality regression filter	82.5 (-1.8)	97.5 (-1.3)	3.30 (-0.96)	10.0 (-1.5)
w/o Relevance filtering	81.2 (-3.1)	96.3 (-2.5)	3.60 (-0.66)	10.3 (-1.2)
w/o Min. visual changes filter	83.5 (-0.8)	98.0 (-0.8)	3.10 (-1.16)	10.9 (-0.6)

Table 6. **Data quality ablation**. Impact of removing individual curation filters (Sec. 3.1) from the training data pipeline.

on CompBench), as off-topic edits undermine maintaining compositional constraints. Removing the minimal visual changes filter most significantly hurts multi-turn editing (1.16 points on ImgEdit), demonstrating that learning meaningful incremental progress is critical for sustained interactions. For visual reasoning (MIRA), removing the quality regression filter has the largest impact (1.5%), as learning from degraded trajectories impairs converging toward correct answers. Different quality dimensions matter for different capabilities; effective test-time scaling requires curating data along multiple axes.

6. Conclusion

We have presented a unified approach to multimodal chain-of-thought test-time scaling that extends inference-time compute from text-only reasoning to models handling both visual understanding and generation, establishing a paradigm that benefits both generation and comprehension across modalities. Our key contributions—an agentic framework that induces cognitive behaviors (verification, subgoal decomposition, content memory), budget forcing for beyond-training generalization, and evidence that sequential reasoning outperforms parallel sampling—yield substantial gains across compositional generation, multi-turn editing, and visual reasoning, demonstrating that iterative refinement through explicit reasoning unlocks significant performance improvements on complex multimodal tasks.

Acknowledgements

We thank Hao Tang, Mark Endo, Yuhui Zhang, and Alejandro Lozano for their helpful discussions and feedback.

References

- [1] Zachary Ankner, Mansheej Paul, Brandon Cui, Jonathan D. Chang, and Prithviraj Ammanabrolu. Critique-out-loud reward models, 2024. 3
- [2] Shuai Bai, Keqin Chen, Xuejing Liu, Jialin Wang, Wenbin Ge, Sibao Song, Kai Dang, Peng Wang, Shijie Wang, Jun Tang, Humen Zhong, Yuanzhi Zhu, Mingkun Yang, Zhaohai Li, Jianqiang Wan, Pengfei Wang, Wei Ding, Zheren Fu, Yiheng Xu, Jiabo Ye, Xi Zhang, Tianbao Xie, Zesen Cheng, Hang Zhang, Zhibo Yang, Haiyang Xu, and Jun-

- yang Lin. Qwen2.5-vl technical report. *arXiv preprint arXiv:2502.13923*, 2025. 8
- [3] Black Forest Labs. Flux 1.1 pro: State-of-the-art text-to-image generation model. <https://bfl.ai/models/flux-pro>, 2024. High-speed text-to-image generation with strong prompt adherence. 4
- [4] Bradley Brown, Jordan Juravsky, Ryan Ehrlich, Ronald Clark, Quoc V. Le, Christopher Ré, and Azalia Mirhoseini. Large language monkeys: Scaling inference compute with repeated sampling, 2024. 3
- [5] Yuhang Chang, Fei Fang, Wenjing Wang, Haochen Chen, Guangtao Zhang, et al. Oneig-bench: Omni-dimensional nuanced evaluation for image generation. *arXiv preprint arXiv:2506.07977*, 2025. 6
- [6] Jihai Chen, Zhiyang Xu, Xichen Pan, Yushi Hu, Can Qin, Tom Goldstein, Lifu Huang, Tianyi Zhou, Saining Xie, Silvio Savarese, et al. Blip3-o: A family of fully open unified multimodal models-architecture, training and dataset. *arXiv preprint arXiv:2505.09568*, 2025. 4
- [7] Mark Chen, Jerry Tworek, Heewoo Jun, Qiming Yuan, Henrique Ponde de Oliveira Pinto, Jared Kaplan, Harri Edwards, Yura Burda, Nicholas Joseph, Greg Brockman, et al. Evaluating large language models trained on code. *arXiv preprint arXiv:2107.03374*, 2021. 2
- [8] Xiaokang Chen, Zhiyu Wu, Xingchao Liu, Zizheng Pan, Wen Liu, Zhenda Xie, Xingkai Yu, and Chong Ruan. Janus-pro: Unified multimodal understanding and generation with data and model scaling. *arXiv preprint arXiv:2501.17811*, 2025. 4
- [9] Xiaokang Chen, Zhiyu Wu, Xingchao Liu, Zizheng Pan, Wen Liu, Zhenda Xie, Xingkai Yu, and Chong Ruan. Janus-pro: Unified multimodal understanding and generation with data and model scaling. *arXiv preprint arXiv:2501.17811*, 2025. 3
- [10] Steffi Chern et al. Thinking with generated images. *arXiv preprint arXiv:2505.22525*, 2025. 2, 3
- [11] Sehyun Choi, Tianqing Fang, Zhaowei Wang, and Yangqiu Song. Kcts: Knowledge-constrained tree search decoding with token-level hallucination detection, 2023. 3
- [12] Karl Cobbe, Vineet Kosaraju, Mohammad Bavarian, Mark Chen, Heewoo Jun, Lukasz Kaiser, Matthias Plappert, Jerry Tworek, Jacob Hilton, Reiichiro Nakano, Christopher Hesse, and John Schulman. Training verifiers to solve math word problems. *arXiv preprint arXiv:2110.14168*, 2021. 2
- [13] DeepSeek-AI, Daya Guo, Dejian Yang, Haowei Zhang, Junxiao Song, Ruoyu Zhang, Runxin Xu, Qihao Zhu, Shiron Ma, Peiyi Wang, et al. Deepseek-r1: Incentivizing reasoning capability in llms via reinforcement learning. *arXiv preprint arXiv:2501.12948*, 2025. 2, 3
- [14] Chaorui Deng, Deyao Zhu, Kunchang Li, Chenhui Gou, Feng Li, Zeyu Wang, Shu Zhong, Weihao Yu, Xiaonan Nie, Ziang Song, Guang Shi, and Haoqi Fan. Emerging properties in unified multimodal pretraining. *arXiv preprint arXiv:2505.14683*, 2025. 4, 6, 8
- [15] Chaorui Deng, Deyao Zhu, Kunchang Li, Chenhui Gou, Feng Li, Zeyu Wang, Shu Zhong, Weihao Yu, Xiaonan Nie, Ziang Song, Guang Shi, and Haoqi Fan. Emerging properties in unified multimodal pretraining. *arXiv preprint arXiv:2505.14683*, 2025. 2, 3, 4
- [16] Runpei Dong, Chunrui Han, Yuang Peng, Zekun Qi, Zheng Ge, Jinrong Yang, Liang Zhao, Jianjian Sun, Hongyu Zhou, Haoran Wei, Xiangwen Kong, Xiangyu Zhang, Kaisheng Ma, and Li Yi. Dreamllm: Synergistic multimodal comprehension and creation. In *The Twelfth International Conference on Learning Representations*, 2024. 3
- [17] Wenshan Fan, Shijie Chen, Jianqiang Zhang, and Donglin Liu. Visualization-of-thought: Eliciting spatial reasoning through sketch-based visual reasoning. *arXiv preprint arXiv:2404.03622*, 2024. 3
- [18] Rongyao Fang, Chengqi Duan, Kun Wang, Linjiang Huang, Hao Li, Shilin Yan, Hao Tian, Xingyu Zeng, Rui Zhao, and Jifeng Dai. Got: Unleashing reasoning capability of multimodal large language model for visual generation and editing. *arXiv preprint arXiv:2503.10639*, 2025. 2, 3, 6
- [19] Kanishk Gandhi, Ayush Chakravarthy, Anikait Singh, Nathan Lile, and Noah D Goodman. Cognitive behaviors that enable self-improving reasoners, or, four habits of highly effective stars. *arXiv preprint arXiv:2503.01307*, 2025. 2
- [20] Yuying Ge, Sijie Zhao, Jinguo Zhu, Yixiao Ge, Kun Yi, Lin Song, Chen Li, Xiaohan Ding, and Ying Shan. Seed-x: Multimodal models with unified multi-granularity comprehension and generation. *arXiv preprint arXiv:2404.14396*, 2024. 6
- [21] Jiawei Gu, Yunzhuo Hao, Huichen Will Wang, Linjie Li, Michael Qizhe Shieh, Yejin Choi, Ranjay Krishna, and Yu Cheng. Thinkmorph: Emergent properties in multimodal interleaved chain-of-thought reasoning. *arXiv preprint arXiv:2510.27492*, 2025. 3
- [22] Zeqi Gu, Markos Georgopoulos, Xiaoliang Dai, Marjan Ghazvininejad, Chu Wang, Felix Juefei-Xu, Kunpeng Li, Yujun Shi, Zecheng He, Zijian He, et al. Improving chain-of-thought efficiency for autoregressive image generation. *arXiv preprint arXiv:2510.05593*, 2025. 3
- [23] Haoran He, Jiajun Liang, Xintao Wang, Pengfei Wan, Di Zhang, Kun Gai, and Ling Pan. Scaling image and video generation via test-time evolutionary search. *arXiv preprint arXiv:2505.17618*, 2025. 3
- [24] Zhenyu Hou, Xin Lv, Rui Lu, Jiajie Zhang, Yujiang Li, Zijun Yao, Juanzi Li, Jie Tang, and Yuxiao Dong. Advancing language model reasoning through reinforcement learning and inference scaling, 2025. 3
- [25] Yushi Hu, Weijia Yin, Xingyu Shi, Anirudh Kumar, Ranjay Krishna, Dorsa Sadigh, and Trevor Darrell. Visual sketchpad: Sketching as a visual chain of thought for multimodal language models. *arXiv preprint arXiv:2406.09403*, 2024. 3
- [26] Wenxuan Huang, Shuang Chen, Zheyong Xie, Shaosheng Cao, Shixiang Tang, Yufan Shen, Qingyu Yin, Wenbo Hu, Xiaoman Wang, Yuntian Tang, et al. Interleaving reasoning for better text-to-image generation. *arXiv preprint arXiv:2509.06945*, 2025. 2, 3
- [27] Ziqi Huang, Ning Yu, Gordon Chen, Haonan Qiu, Paul Debevec, and Ziwei Liu. Vchain: Chain-of-visual-thought for reasoning in video generation. *arXiv preprint arXiv:2510.05094*, 2025. 3

- [28] Mude Hui, Siwei Yang, Bingchen Zhao, Yichun Shi, Heng Wang, Peng Wang, Yuyin Zhou, and Cihang Xie. Hq-edit: A high-quality dataset for instruction-based image editing. *arXiv preprint arXiv:2404.09990*, 2024. 6
- [29] Robert Irvine, Douglas Boubert, Vyas Raina, Adian Liusie, Ziyi Zhu, Vineet Mudupalli, Aliaksei Korshuk, Zongyi Liu, Fritz Cremer, Valentin Assassi, Christie-Carol Beauchamp, Xiaoding Lu, Thomas Rialan, and William Beauchamp. Rewarding chatbots for real-world engagement with millions of users, 2023. 3
- [30] Bohan Jia, Wenxuan Huang, Yuntian Tang, Junbo Qiao, Jincheng Liao, Shaosheng Cao, Fei Zhao, Zhaopeng Feng, Zhouhong Gu, Zhenfei Yin, et al. Compbench: Benchmarking complex instruction-guided image editing. *arXiv preprint arXiv:2505.12200*, 2025. 6
- [31] Dongzhi Jiang, Ziyu Guo, Renrui Zhang, Zhuofan Zong, Hao Li, Le Zhuo, Shilin Yan, Pheng-Ann Heng, and Hongsheng Li. T2i-r1: Reinforcing image generation with collaborative semantic-level and token-level cot. *arXiv preprint arXiv:2505.00703*, 2025. 3
- [32] Black Forest Labs. Flux.1 kontext: Flow matching for in-context image generation and editing in latent space. *arXiv preprint arXiv:2506.15742*, 2025. 4, 6
- [33] Kuang-Huei Lee, Ian Fischer, Yueh-Hua Wu, Dave Marwood, Shumeet Baluja, Dale Schuurmans, and Xinyun Chen. Evolving deeper llm thinking, 2025. 3
- [34] Noam Levi. A simple model of inference scaling laws, 2024. 3
- [35] Hunter Lightman, Vineet Kosaraju, Yura Burda, Harri Edwards, Bowen Baker, Teddy Lee, Jan Leike, John Schulman, Ilya Sutskever, and Karl Cobbe. Let’s verify step by step, 2023. 3
- [36] Fangfu Liu, Hanyang Wang, Yimo Cai, Kaiyan Zhang, Xiaohang Zhan, and Yueqi Duan. Video-t1: Test-time scaling for video generation. *arXiv preprint arXiv:2503.18942*, 2025. 3
- [37] Jiacheng Liu, Andrew Cohen, Ramakanth Pasunuru, Yejin Choi, Hannaneh Hajishirzi, and Asli Celikyilmaz. Don’t throw away your value model! generating more preferable text with value-guided monte-carlo tree search decoding, 2024. 3
- [38] Shiyu Liu, Yucheng Han, Peng Xing, Fukun Yin, Rui Wang, Wei Cheng, Jiaqi Liao, Yingming Wang, Honghao Fu, Chunrui Han, Guopeng Li, Yuang Peng, Quan Sun, Jingwei Wu, Yan Cai, Zheng Ge, Ranchen Ming, Lei Xia, Xianfang Zeng, Yibo Zhu, Binxing Jiao, Xiangyu Zhang, Gang Yu, and Daxin Jiang. Step1x-edit: A practical framework for general image editing. *arXiv preprint arXiv:2504.17761*, 2025. 6
- [39] Yang Liu, Hao Zhang, Xiang Wang, and Jiaming Chen. Mvot: Multi-view thinking for enhanced visual reasoning. *arXiv preprint arXiv:2407.12345*, 2024. 3
- [40] Jiasen Lu, Christopher Clark, Sangho Lee, Zichen Zhang, Savya Khosla, Ryan Marten, Derek Hoiem, and Aniruddha Kembhavi. Unified-io 2: Scaling autoregressive multimodal models with vision language audio and action. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 26439–26455, 2024. 3
- [41] Yiyang Ma, Xingchao Liu, Xiaokang Chen, Wen Liu, Chengyue Wu, Zhiyu Wu, Zizheng Pan, Zhenda Xie, Haowei Zhang, Xingkai yu, Liang Zhao, Yisong Wang, Jiaying Liu, and Chong Ruan. Janusflow: Harmonizing autoregression and rectified flow for unified multimodal understanding and generation, 2024. 3
- [42] Yuhang Ma, Xiaoshi Wu, Keqiang Sun, and Hongsheng Li. Hpsv3: Towards wide-spectrum human preference score. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 15086–15095, 2025. 7
- [43] Aman Madaan, Niket Tandon, Prakhar Gupta, Skyler Hallinan, Luyu Gao, Sarah Wiegrefe, Uri Alon, Nouha Dziri, Shrimai Prabhumoye, Yiming Yang, Sean Welleck, Bodhisattwa Prasad Majumder, Shashank Gupta, Amir Yazdanbakhsh, and Peter Clark. Self-refine: Iterative refinement with self-feedback. *ArXiv*, abs/2303.17651, 2023. 3
- [44] Meta. Llama 4 scout: 17b parameter mixture-of-experts model. <https://huggingface.co/meta-llama/Llama-4-Scout-17B-16E>, 2025. Multimodal mixture-of-experts language model with 17B activated parameters. 4
- [45] Niklas Muennighoff, Zitong Yang, Weijia Shi, Xiang Lisa Li, Li Fei-Fei, Hannaneh Hajishirzi, Luke Zettlemoyer, Percy Liang, Emmanuel Candès, and Tatsunori Hashimoto. s1: Simple test-time scaling. *arXiv preprint arXiv:2501.19393*, 2025. 3, 5, 7
- [46] OpenAI. Openai o1 system card, 2024. 2, 3
- [47] OpenAI. Introducing gpt-5. <https://openai.com/index/introducing-gpt-5/>, 2025. Accessed: 2025-11-08. 8
- [48] Xichen Pan, Yan Cai, Junyi Li, Yuheng Lu, Xicheng Huang, Xu Jia, Juefei Xu, Runxin Xu, Peng Li, Zheng Ge, Zhenguo Ying, Zheng Liu, Shuicheng Huang, Jiazhi Zhao, and Yuanhan Peng. Transfer between modalities with metaqueries. *arXiv preprint arXiv:2504.06256*, 2025. 3
- [49] Luozheng Qin, Jia Gong, Yuqing Sun, Tianjiao Li, Mengping Yang, Xiaomeng Yang, Chao Qu, Zhiyu Tan, and Hao Li. Uni-cot: Towards unified chain-of-thought reasoning across text and vision. *arXiv preprint arXiv:2508.05606*, 2025. 3
- [50] Liao Qu, Huichao Zhang, Yiheng Liu, Xu Wang, Yi Jiang, Yiming Gao, Hu Ye, Daniel K Du, Zehuan Yuan, and Xinglong Wu. Tokenflow: Unified image tokenizer for multimodal understanding and generation. *arXiv preprint arXiv:2412.03069*, 2024. 3
- [51] Qwen Team. Qwen3-vl: Technical report. *Technical Report, Alibaba Cloud*, 2025. Available at <https://github.com/QwenLM/Qwen3-VL>. 4
- [52] Qwen Team. Qwen-image-edit: Ai-powered image editing model. *Technical Report, Alibaba Cloud*, 2025. 20B parameter image editing model. 4, 6
- [53] Hao Shao, Shengju Wang, Liang Li, Xiang Wang, and Liu Yang. Visual-cot: Unleashing chain-of-thought reasoning in multi-modal language models. *arXiv preprint arXiv:2403.16999*, 2024. 2, 3
- [54] Weijia Shi, Xiaochuang Han, Chunting Zhou, Weixin Liang, Xi Victoria Lin, Luke Zettlemoyer, and Lili Yu. Lmfusion: Adapting pretrained language models for multimodal generation. *arXiv preprint arXiv:2412.15188*, 2024. 3

- [55] Charlie Snell, Jaehoon Lee, Kelvin Xu, and Aviral Kumar. Scaling llm test-time compute optimally can be more effective than scaling model parameters. *arXiv preprint arXiv:2408.03314*, 2024. 2
- [56] Charlie Snell, Jaehoon Lee, Kelvin Xu, and Aviral Kumar. Scaling llm test-time compute optimally can be more effective than scaling model parameters, 2024. 3
- [57] Aarohi Srivastava, Abhinav Rastogi, Abhishek Rao, Abu Awal Md Shoeb, Abubakar Abid, Adam Fisch, Adam R Brown, Adam Santoro, et al. Beyond the imitation game: Quantifying and extrapolating the capabilities of language models. *arXiv preprint arXiv:2206.04615*, 2023. 2
- [58] T2i-CoReBench Team. T2i-corebench: A comprehensive benchmark for evaluating composition and reasoning in text-to-image generation. <https://t2i-corebench.github.io/>, 2025. 4
- [59] Chameleon Team. Chameleon: Mixed-modal early-fusion foundation models. *arXiv preprint arXiv:2405.09818*, 2024. 3
- [60] Shengbang Tong, David Fan, Jiachen Zhu, Yunyang Xiong, Xinlei Chen, Koustuv Sinha, Michael Rabbat, Yann LeCun, Saining Xie, and Zhuang Liu. Metamorph: Multimodal understanding and generation via instruction tuning. *arXiv preprint arXiv:2412.14164*, 2024. 3
- [61] Shengbang Tong et al. Metamorph: Multimodal understanding and generation via instruction tuning. *arXiv preprint*, 2024. Placeholder for wu2024next - update with correct citation if different. 3
- [62] Peiyi Wang, Lei Li, Zhihong Shao, R. X. Xu, Damai Dai, Yifei Li, Deli Chen, Y. Wu, and Zhifang Sui. Math-shepherd: Verify and reinforce llms step-by-step without human annotations, 2024. 3
- [63] Xinlong Wang, Xiaosong Zhang, Zhengxiong Luo, Quan Sun, Yufeng Cui, Jinsheng Wang, Fan Zhang, Yueze Wang, Zhen Li, Qiying Yu, et al. Emu3: Next-token prediction is all you need. *arXiv preprint arXiv:2409.18869*, 2024. 3
- [64] Zhilin Wang, Yi Dong, Olivier Delalleau, Jiaqi Zeng, Gerald Shen, Daniel Egert, Jimmy J. Zhang, Makesh Narsimhan Sreedhar, and Oleksii Kuchaiev. Helpsteer2: Open-source dataset for training top-performing reward models, 2024. 3
- [65] Jason Wei, Xuezhi Wang, Dale Schuurmans, Maarten Bosma, Fei Xia, Ed Chi, Quoc V Le, Denny Zhou, et al. Chain-of-thought prompting elicits reasoning in large language models. *Advances in neural information processing systems*, 35:24824–24837, 2022. 3
- [66] Chengyue Wu, Xiaokang Chen, Zhiyu Wu, Yiyang Ma, Xingchao Liu, Zizheng Pan, Wen Liu, Zhenda Xie, Xing kai Yu, Chong Ruan, et al. Janus: Decoupling visual encoding for unified multimodal understanding and generation. *arXiv preprint arXiv:2410.13848*, 2024. 2, 3
- [67] Chenyuan Wu, Pengfei Zheng, Ruiran Yan, Shitao Xiao, Xin Luo, Yueze Wang, Wanli Li, Xiyan Jiang, Yexin Liu, Junjie Zhou, Ze Liu, Ziyi Xia, Chaofan Li, Haoge Deng, Jiahao Wang, Kun Luo, Bo Zhang, Defu Lian, Xinlong Wang, Zhongyuan Wang, Tiejun Huang, and Zheng Liu. Omnigen2: Exploration to advanced multimodal generation. *arXiv preprint arXiv:2506.18871*, 2025. 3
- [68] Yangzhen Wu, Zhiqing Sun, Shanda Li, Sean Welleck, and Yiming Yang. Inference scaling laws: An empirical analysis of compute-optimal inference for problem-solving with language models, 2024. 3
- [69] Yicheng Xiao, Lin Song, Yukang Chen, Yingmin Luo, Yuxin Chen, Yukang Gan, Wei Huang, Xiu Li, Xiaojuan Qi, and Ying Shan. Mindomni: Unleashing reasoning generation in vision language models with rgpo. *arXiv preprint arXiv:2505.13031*, 2025. 2, 3
- [70] Huajian Xin, Daya Guo, Zhihong Shao, Zhizhou Ren, Qihao Zhu, Bo Liu, Chong Ruan, Wenda Li, and Xiaodan Liang. Deepseek-prover: Advancing theorem proving in llms through large-scale synthetic data, 2024. 3
- [71] Yang Ye, Xianyi He, Zongjian Li, Bin Lin, Shenghai Yuan, Zhiyuan Yan, Bohan Hou, and Li Yuan. Imgedit: A unified image editing dataset and benchmark. *arXiv preprint arXiv:2505.20275*, 2025. 6
- [72] Haofei Yu, Zhengyang Qi, Lawrence Jang, Ruslan Salakhutdinov, Louis-Philippe Morency, and Paul Pu Liang. Mmoe: Enhancing multimodal models with mixtures of multimodal interaction experts. In *Proceedings of the 2024 Conference on Empirical Methods in Natural Language Processing*, pages 10006–10030, 2024. 3
- [73] Qifan Yu, Wei Chow, Zhongqi Yue, Kaihang Pan, Yang Wu, Xiaoyang Wan, Juncheng Li, Siliang Tang, Hanwang Zhang, and Yueting Zhuang. Anyedit: Mastering unified high-quality image editing for any idea. *arXiv preprint arXiv:2411.15738*, 2024. 6
- [74] Richard Zhang, Phillip Isola, Alexei A Efros, Eli Shechtman, and Oliver Wang. The unreasonable effectiveness of deep features as a perceptual metric. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 586–595, 2018. 4
- [75] Shun Zhang, Zhenfang Chen, Yikang Shen, Mingyu Ding, Joshua B. Tenenbaum, and Chuang Gan. Planning with large language models for code generation, 2023. 3
- [76] Yiyang Zhang, Wei Chen, Xiang Wang, and Jun Li. Chain-of-image: Advancing visual reasoning via image sequence generation. *arXiv preprint arXiv:2312.06141*, 2024. 3
- [77] Haozhe Zhao, Xiaojian Shawn Ma, Liang Chen, Shuzheng Si, Rujie Wu, Kaikai An, Peiyu Yu, Minjia Zhang, Qing Li, and Baobao Chang. Ultraedit: Instruction-based fine-grained image editing at scale. *Advances in Neural Information Processing Systems*, 37:3058–3093, 2024. 6
- [78] Andy Zhou, Kai Yan, Michal Shlapentokh-Rothman, Hao-han Wang, and Yu-Xiong Wang. Language agent tree search unifies reasoning acting and planning in language models, 2024. 3
- [79] Chunting Zhou, Lili Yu, Arun Babu, Kushal Tirumala, Michihiro Yasunaga, Leonid Shamis, Jacob Kahn, Xuezhe Ma, Luke Zettlemoyer, and Omer Levy. Transfusion: Predict the next token and diffuse images with one multi-modal model. *arXiv preprint arXiv:2408.11039*, 2024. 3
- [80] Chunting Zhou, Lili Yu, Arun Babu, Kushal Tirumala, Michihiro Yasunaga, Leonid Shamis, Jacob Kahn, Xuezhe Ma, Luke Zettlemoyer, and Omer Levy. Transfusion: Predict the next token and diffuse images with one multi-modal model. *arXiv preprint arXiv:2408.11039*, 2024. 2