

000 001 002 003 004 005 006 007 008 009 010 011 012 013 014 015 016 017 018 019 020 021 022 023 024 025 026 027 028 029 030 031 032 033 034 035 036 037 038 039 040 041 042 043 044 045 046 047 048 049 050 051 052 053 FROM BROAD EXPLORATION TO STABLE SYNTHESIS: ENTROPY-GUIDED OPTIMIZATION FOR AUTOREGRESSIVE IMAGE GENERATION

Anonymous authors

Paper under double-blind review

ABSTRACT

Combining Chain-of-Thought (CoT) with Reinforcement Learning (RL) improves text-to-image (T2I) generation, yet the underlying interaction between CoT’s exploration and RL’s optimization remains unclear. We present a systematic entropy-based analysis that yields three key insights: (1) CoT expands the generative exploration space, while RL contracts it toward high-reward regions; (2) final reward is strongly negatively correlated with both the mean and variance of image-token entropy, highlighting the need to reduce uncertainty and instability; and (3) the entropy of the textual CoT directly governs downstream image quality, with lower-entropy CoTs leading to better generations. Motivated by these findings, we propose *Entropy-Guided Group Relative Policy Optimization* (EG-GRPO), a fine-tuning strategy that reallocates optimization budget by uncertainty: low-entropy tokens are excluded from reward-driven updates to preserve stability, while high-entropy tokens receive an entropy bonus that encourages structured exploration without collapse. Experiments on standard T2I benchmarks demonstrate that EG-GRPO achieves state-of-the-art performance.

1 INTRODUCTION

Text-to-image (T2I) generation has progressed rapidly with large-scale pretraining and strong autoregressive and diffusion architectures (Chen et al., 2023; Labs, 2024; Wang et al., 2024), yet two core challenges remain: (i) balancing *exploration* for diversity against *exploitation* for reward-aligned fidelity, and (ii) ensuring *generation stability* under repeated sampling. Chain-of-Thought (CoT) prompting promises richer semantic planning (Wei et al., 2022), while reinforcement learning (RL) directly optimizes task or preference rewards (Jiang et al., 2025). However, how CoT’s exploratory behavior interacts with RL’s optimization in T2I, and how this interaction governs uncertainty and stability, has not been systematically understood.

We analyze *entropy dynamics* in autoregressive T2I models that combine CoT with RL (via Group Relative Policy Optimization, GRPO (Jiang et al., 2025)) and use Shannon entropy to quantify token-level uncertainty in both modalities: textual CoT tokens and image tokens. Three empirical findings emerge. First, CoT *expands* the exploration space, broadening the entropy distribution of generated outputs, whereas RL *contracts* this space toward higher-reward regions. Second, final reward exhibits a strong negative correlation with both the *mean* and the *standard deviation* of image-token entropy, indicating that reducing uncertainty and instability is central to quality. Third, the entropy of the textual CoT *directly* influences downstream image quality: lower-entropy CoTs yield tighter, higher-reward clusters under stable sampling.

Guided by these findings, we propose *Entropy-Guided Group Relative Policy Optimization* (EG-GRPO), a token-level modification of GRPO that reallocates gradient budget by uncertainty. Low-entropy (high-confidence) tokens are excluded from reward-driven updates, retaining only the KL-to-reference term to preserve stability and previously acquired knowledge. High-entropy tokens receive an *entropy bonus* added to the advantage, encouraging structured exploration and accelerating uncertainty reduction where it matters. A batch-level calibration ties the bonus magnitude to the mass saved on low-entropy tokens, keeping update scale close to GRPO, and the bonus vanishes at GRPO equilibrium, preserving the stationary points of the base objective (Wang et al., 2025c).

We evaluate EG-GRPO on T2I-CompBench (Huang et al., 2023) and WISE (Niu et al., 2025) using a Janus-Pro autoregressive backbone in a discrete latent space (Chen et al., 2025) and a standard reward pipeline combining human-preference scoring, object grounding, and VQA signals (Wu et al., 2023; Liu et al., 2024; Wang et al., 2022). EG-GRPO attains state-of-the-art results, with pronounced gains in compositional generalization (e.g., attribute binding and object relations). Ablations that apply entropy guidance to only CoT tokens or only image tokens underperform the full model, confirming the need to control uncertainty in both semantic planning and visual decoding.

The main contributions of our paper can be summarized as follows:

- We provide a quantitative account of the CoT-RL interaction through entropy dynamics: CoT expands exploration, RL contracts toward high-reward regions; reward is strongly negatively correlated with entropy mean and std; and textual CoT entropy governs downstream image quality.
- We introduce an entropy-guided, token-level refinement of GRPO that protects confident tokens via KL-only updates and focuses optimization on uncertain tokens via an entropy bonus, with calibrated budget and equilibrium-vanishing properties.
- On T2I-CompBench and WISE, EG-GRPO achieves state-of-the-art performance and reduces uncertainty and instability in line with the analysis.

2 RELATED WORK

2.1 TEXT-TO-IMAGE GENERATION MODEL

Text-to-image generation has been advanced along two major paradigms: autoregressive modeling and diffusion-based approaches. On the autoregressive side, Parti (Yu et al., 2022) demonstrates that large-scale transformer models can achieve impressive compositional generation by predicting image tokens sequentially. Fluid (Fan et al., 2024) further explores continuous tokens and generation order, showing improved efficiency and quality. STAR (Ma et al., 2024) introduces a scale-wise autoregressive framework that progressively generates images from coarse to fine scales, while JetFormer (Tschannen et al., 2024) directly models raw images and text in a unified autoregressive manner without discrete tokenization. More recently, NextStep-1 (Team et al., 2025) scales continuous-token autoregressive generation to 14B parameters, achieving strong performance in both image synthesis and editing. In parallel, diffusion models have become dominant for high-quality synthesis. VQ-Diffusion (Gu et al., 2022) combines vector-quantized representations with diffusion for discrete latent modeling, and ERNIE-ViLG 2.0 (Feng et al., 2023) extends this to large-scale multilingual text-to-image generation. UPainting (Li et al., 2022) introduces cross-modal guidance to unify simple and complex scenarios, while RPG (Yang et al., 2024) enhances controllability via recaptioning, planning, and region-based diffusion. These works illustrate the complementary strengths of autoregressive and diffusion frameworks in advancing the controllability, fidelity, and scalability of text-to-image generation.

2.2 CHAIN OF THINKING AND REINFORCEMENT LEARNING

Recent works have explored integrating Chain-of-Thought (CoT) reasoning with reinforcement learning (RL) to improve text-to-image generation. Visual-CoG (Li et al., 2025) introduces stage-aware RL with intermediate rewards across semantic planning, refinement, and evaluation, while ReasonGen-R1 (Zhang et al., 2025b) and T2I-R1 (Jiang et al., 2025) incorporate rationale-augmented data and bi-level reasoning chains optimized via GRPO. PromptEnhancer (Wang et al., 2025b) further demonstrates that CoT-based prompt rewriting with RL can enhance image quality without modifying the generator, and verification-based methods also inject preference alignment during synthesis (Zhang et al., 2025a). Beyond explicit CoT, RL has been applied for alignment in diffusion models, such as comparing DPO and GRPO (Tong et al., 2025), subject-driven preference optimization (Miao et al., 2024), and DPOK fine-tuning with KL regularization (Fan et al., 2023). Related multimodal reasoning approaches, including ImageGen-CoT (Liao et al., 2025) and reflective CoT for retrieval (Wu et al., 2024), highlight the broader potential of structured reasoning. Combining CoT and RL provides a promising avenue for enhancing controllability, interpretability, and human alignment in text-to-image generation. While SimpleAR (Wang et al., 2025a) and Gallici & Borde (2025) demonstrate the efficacy of standard GRPO for high-fidelity and style-aligned

108 autoregressive generation, our approach introduces an entropy-guided mechanism to reallocate the
 109 optimization budget at the token level dynamically.
 110

111 3 PRELIMINARIES

113 We begin by formalizing autoregressive text-to-image generation, introducing Shannon entropy as a
 114 measure of generative uncertainty, and reviewing Group Relative Policy Optimization (GRPO) for
 115 model refinement.
 116

117 3.1 AUTOREGRESSIVE GENERATION IN DISCRETE LATENT SPACE

119 Autoregressive text-to-image models operate in a discrete latent space. An image I is first encoded
 120 into tokens $z = (z_1, z_2, \dots, z_L)$ via a pre-trained tokenizer such as VQ-VAE. The conditional like-
 121 lihood given a text prompt c factorizes autoregressively:

$$122 \quad p(z | c) = \prod_{i=1}^L p(z_i | z_{<i}, c). \quad (1)$$

125 A policy π_θ parameterized by θ models these conditionals. At step i , the policy outputs a categorical
 126 distribution $\pi_\theta(\cdot | z_{<i}, c)$ over the vocabulary \mathcal{V} , from which z_i is sampled. The final image is
 127 reconstructed by decoding the full sequence z .

128 3.2 ENTROPY AS PREDICTIVE UNCERTAINTY

130 For each step i , the policy distribution admits a Shannon entropy:

$$132 \quad H(\pi_\theta(\cdot | z_{<i}, c)) = - \sum_{j \in \mathcal{V}} p_j \log p_j, \quad (2)$$

134 where p_j denotes the probability of token j . Low entropy reflects confident, often high-fidelity pre-
 135 dictions with reduced diversity, whereas high entropy reflects uncertainty, encouraging exploration
 136 and diverse generations at the cost of possible incoherence.

137 3.3 GROUP RELATIVE POLICY OPTIMIZATION

139 Group Relative Policy Optimization (GRPO) refines generative policies using relative rewards with-
 140 out requiring a value function. For a prompt c , the policy π_θ samples G candidate sequences

$$142 \quad \{o^{(i)}\}_{i=1}^G \sim \pi_\theta(\cdot | c), \quad o^{(i)} = (o_1^{(i)}, \dots, o_{T^{(i)}}^{(i)}). \quad (3)$$

143 Each sequence receives a reward $r^{(i)}$, which is normalized within the group:

$$145 \quad \mu = \frac{1}{G} \sum_{i=1}^G r^{(i)}, \quad \sigma = \sqrt{\frac{1}{G} \sum_{i=1}^G (r^{(i)} - \mu)^2}, \quad A^{(i)} = \frac{r^{(i)} - \mu}{\max\{\sigma, \varepsilon\}}. \quad (4)$$

148 Broadcasting $A^{(i)}$ to all tokens yields the training objective

$$149 \quad \mathcal{L}_{\text{GRPO}}(\theta) = -\frac{1}{G} \sum_{i=1}^G \frac{1}{T^{(i)}} \sum_{t=1}^{T^{(i)}} A^{(i)} \log \pi_\theta(o_t^{(i)} | c, o_{<t}^{(i)}) + \beta D_{\text{KL}}(\pi_\theta \| \pi_{\text{ref}}), \quad (5)$$

152 where $\beta \geq 0$ controls a KL regularizer toward a reference policy. This group-based normaliza-
 153 tion yields scale-invariant advantages and stable updates, making GRPO well-suited for aligning
 154 autoregressive generators with task-specific rewards.

155 4 ANALYSIS OF ENTROPY DYNAMICS IN TEXT-TO-IMAGE GENERATION

158 To investigate the role of uncertainty in text-to-image generation, we analyze *generative entropy* as
 159 a quantitative indicator of model behavior. Our study focuses on disentangling how CoT and RL
 160 fine-tuning affect entropy within the pipeline. By examining their individual and combined effects,
 161 we clarify how these techniques balance exploration and exploitation, and how the resulting entropy
 dynamics define the optimization objective for high-quality visual synthesis.

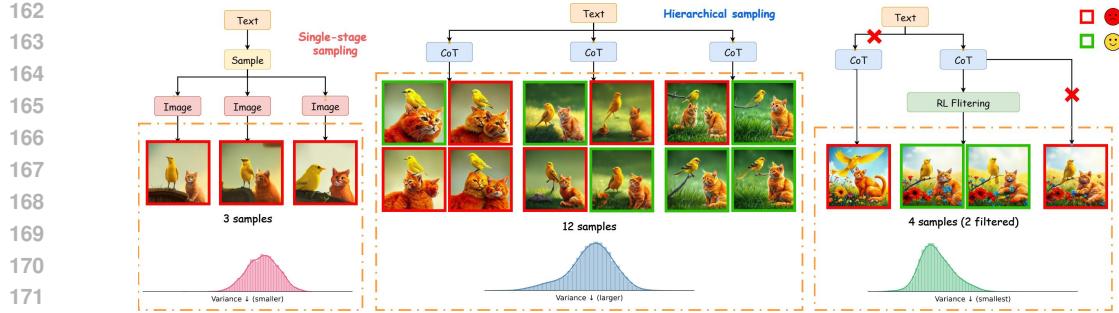


Figure 1: Comparison of different text-to-image generation methods: (a) autoregressive text-to-image generation, (b) CoT, and (c) with CoT and GRPO optimization.

4.1 THE DICHOTOMY OF EXPLORATION AND EXPLOITATION: CoT vs. RL

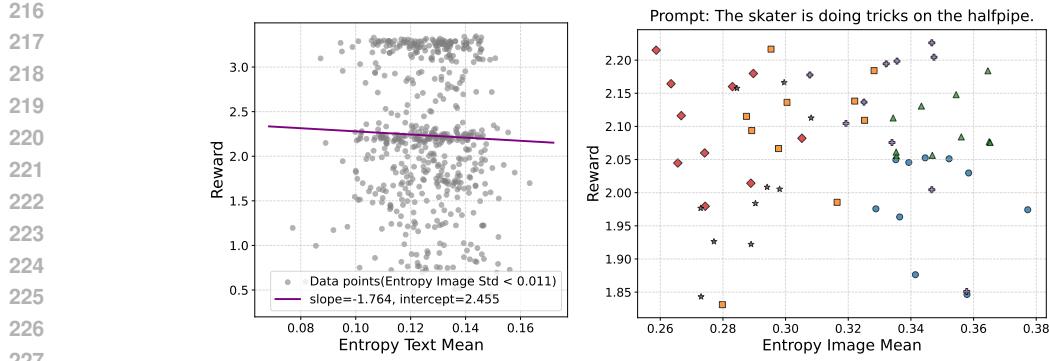
We begin our analysis by examining the distinct yet complementary roles of Chain-of-Thought (CoT) prompting and reinforcement learning (RL) fine-tuning. For each textual prompt, we generate multiple image candidates under three settings: the baseline model (*Janus-Pro*), the baseline augmented with CoT reasoning (*Janus-Pro+CoT*), and a GRPO-finetuned variant built upon *Janus-Pro+CoT* (*T2I-RL*). To investigate their differences, we assess each generated image by jointly measuring its output mean entropy and reward score, and visualize the resulting distributions in a two-dimensional space.

As illustrated in Figure 2, the introduction of CoT significantly broadens the entropy distribution of the generated images. It shows that CoT expands the model’s exploratory space, enabling it to generate a more diverse set of outputs. This expanded space contains both low- and high-reward samples, indicating that CoT itself does not guarantee quality but rather increases the range of generative possibilities. Conversely, the application of Group Relative Policy Optimization (GRPO), which results in the T2I-RL, leads to a notable contraction and leftward shift of the entropy distribution. This demonstrates that the model learns to exploit the vast space unlocked by CoT, converging towards a much narrower, lower-entropy region that consistently yields higher rewards. This reveals a complementary relationship: CoT serves to **expand the exploratory landscape**, while RL acts as a refinement mechanism to **exploit this landscape and guide the model towards stable, high-quality regions**.

4.2 UPSTREAM INFLUENCE: HOW CoT’S TEXTUAL ENTROPY GOVERNS IMAGE QUALITY

Since Chain-of-Thought serves as the entry point of exploration, we further examine how the intrinsic uncertainty of the textual CoT affects downstream image generation. To isolate this effect, we focus on a subset of samples where image generation remains stable, defined as those with entropy variance below a small threshold of 0.011 across multiple runs. This filtering minimizes confounding factors due to unstable sampling, allowing us to probe the impact of CoT entropy directly.

As shown in Figure 3 (Left), we observe a clear negative correlation between the mean entropy of the textual CoT and the average reward of the resulting images. Intuitively, high-entropy CoTs correspond to less coherent or more uncertain reasoning traces, which tend to degrade visual quality.



228
229
230
231
232
233
234
235
236
237
238
239
240
241
242
243
244
245
246
247
248
249
250
251
252
253
254
255
256
257
258
259
260
261
262
263
264
265
266
267
268
269

Figure 3: **Left:** Reward vs. CoT entropy (stable cases, Image Entropy Std < 0.011). Higher CoT entropy correlates with lower image reward. **Right:** Reward distributions across different CoTs for the same prompt. Images from the same CoT cluster together, with certain CoTs consistently yielding lower rewards.

In contrast, low-entropy CoTs provide more confident and consistent reasoning, ultimately leading to higher-reward generations. This analysis is grounded on the stable subset, ensuring that the observed trend is not an artifact of sampling instability.

To further dissect this phenomenon, we visualize the reward distributions of images conditioned on different CoTs for the same prompt (Figure 3, Right). Each distinct CoT forms a compact cluster in the reward space, highlighting its consistent influence on generation outcomes. Notably, certain CoTs repeatedly produce clusters with lower average rewards, suggesting that the quality bottleneck is determined upstream, at the textual reasoning stage. These findings demonstrate a direct transmission of uncertainty from text to image: the entropy of CoT reasoning acts as a critical upstream factor that governs the attainable quality of visual outputs.

4.3 THE LEARNED OBJECTIVE: MINIMIZING UNCERTAINTY AND INSTABILITY FOR HIGHER REWARDS

We next investigate how the fully RL-trained model internalizes entropy control to optimize for reward. For each prompt–CoT pair, multiple images are generated, and their entropy statistics are aggregated. Specifically, the mean entropy characterizes the global level of uncertainty in the generative process, while the standard deviation (std) captures the degree of instability across different runs.

Empirical results reveal a consistent negative correlation between reward and standard deviation of entropy (Figure 4, Left). This indicates that generative instability is inherently detrimental: mod-

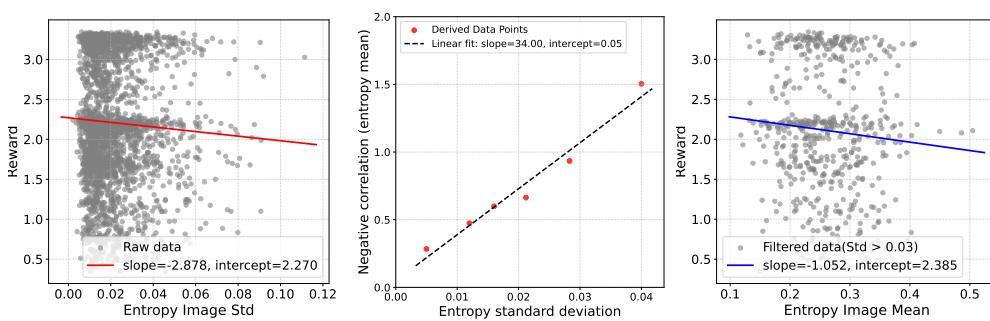


Figure 4: **Left:** Reward vs. entropy std. Higher instability (larger std) consistently lowers reward. **Middle:** Relation between entropy std (x-axis) and the negative correlation of reward–entropy mean (y-axis). Greater instability strengthens the negative correlation. **Right:** Reward vs. entropy mean under high-variance cases (std > 0.03). Large std implies exploratory generation where RL has not converged; in this regime, reducing mean entropy is especially beneficial.

270 els that produce more consistent entropy trajectories across samples yield higher-quality outputs.
 271 Beyond this first-order observation, we further examine how instability modulates the role of un-
 272 certainty. Figure 4 (Middle) demonstrates that the negative correlation between mean entropy and
 273 reward becomes progressively stronger as std increases. In other words, instability amplifies the
 274 adverse impact of uncertainty, rendering mean entropy a more decisive factor under high-variance
 275 conditions.

276 To isolate this effect, we analyze the high-variance regime ($\text{std} > 0.03$), where the image token
 277 generation remains exploratory and the RL policy has not yet converged. As shown in Figure 4
 278 (Right), in such cases, lowering mean entropy proves particularly effective in improving reward.
 279 This finding suggests that, when the model operates in an unstable exploratory state, suppressing
 280 overall uncertainty is a critical pathway toward higher-quality generations.

281 These analyses indicate that the RL agent implicitly optimizes a compound objective: **simulta-**
 282 **neously minimizing overall uncertainty (mean entropy) and reducing instability across runs**
 283 **(entropy std).**

285 5 ENTROPY-GUIDED GROUP RELATIVE POLICY OPTIMIZATION

288 Building on Section 4, we introduce a token-level optimization scheme that preserves GRPO’s
 289 group-relative structure while reallocating updates toward uncertain parts of the generation pro-
 290 cess. *Entropy-Guided GRPO (EG-GRPO)* applies to both textual CoT and image tokens and adds
 291 an entropy-based bonus only where uncertainty is high.

292 5.1 DESIGN PRINCIPLES FROM ENTROPY ANALYSIS

294 Section 4 showed that: (i) CoT broadens exploration while RL contracts it toward high-reward re-
 295 gions; (ii) rewards are negatively correlated with both the mean and the variance of token entropies;
 296 and (iii) textual (CoT) entropy causally influences downstream image quality. We adopt three prin-
 297 ciples:

- 299 1. **Focus on uncertainty.** Allocate more update mass to high-entropy tokens to reduce insta-
 300 bility where it matters.
- 301 2. **Protect confidence.** On the lowest-entropy tokens, set the reward-driven advantage to zero
 302 so that only the KL-to-reference acts, preventing drift on confident regions.
- 303 3. **Stay reward-driven.** Retain the GRPO group-relative objective; entropy contributes an
 304 *additive* per-token bonus at high entropy without replacing advantage.

306 5.2 REVISITING GRPO AND TOKEN BROADCAST

308 For prompt c , policy π_θ samples G sequences $\{o^{(i)}\}_{i=1}^G$ with rewards $\{r^{(i)}\}$ normalized to group-
 309 relative advantages $A^{(i)}$. GRPO optimizes

$$310 \quad \mathcal{L}_{\text{GRPO}}(\theta) = -\frac{1}{G} \sum_{i=1}^G \frac{1}{T^{(i)}} \sum_{t=1}^{T^{(i)}} A^{(i)} \log \pi_\theta(o_t^{(i)} \mid c, o_{<t}^{(i)}) + \beta D_{\text{KL}}(\pi_\theta \parallel \pi_{\text{ref}}). \quad (6)$$

313 This *broadcasts* the same coefficient $A^{(i)}$ to all tokens of sequence i , ignoring per-token uncertainty
 314 and potentially wasting gradient budget on already-confident tokens.

316 5.3 ENTROPY-GUIDED TOKEN SELECTION

318 Let $H_t^{(i)} \triangleq H(\pi_\theta(\cdot \mid c, o_{<t}^{(i)}))$ be the Shannon entropy at token t of sequence i , and define the
 319 normalized entropy $\bar{H}_t^{(i)} \triangleq H_t^{(i)} / \log |\mathcal{V}| \in [0, 1]$. For each sequence i and each modality $m \in$
 320 $\{\text{textual, image}\}$ independently,¹ compute per-sequence percentiles and define

$$322 \quad \mathcal{S}_{\text{hi}}^{(i,m)} = \text{top-50\% by } \bar{H}_t^{(i)}, \quad \mathcal{S}_{\text{lo}}^{(i,m)} = \text{bottom-20\%}, \quad \mathcal{S}_{\text{mid}}^{(i,m)} = \text{remaining}.$$

323 ¹We rank and threshold entropies on textual CoT and image tokens separately, consistent with Section 4.

324 Introduce masks $M_t^{(i)}, U_t^{(i)} \in \{0, 1\}$:

$$326 \quad M_t^{(i)} = \mathbb{I}[t \notin \mathcal{S}_{\text{lo}}^{(i,m)}], \quad U_t^{(i)} = \mathbb{I}[t \in \mathcal{S}_{\text{hi}}^{(i,m)}],$$

327 so $M_t^{(i)} = 0$ removes reward-driven updates on low-entropy tokens, and $U_t^{(i)} = 1$ marks high-
328 entropy tokens to receive a bonus.
329

330 **Budget view.** Let $p_{\text{lo}}, p_{\text{mid}}, p_{\text{hi}}$ be the fractions of tokens in $\mathcal{S}_{\text{lo}}^{(i,m)}, \mathcal{S}_{\text{mid}}^{(i,m)}, \mathcal{S}_{\text{hi}}^{(i,m)}$ ($p_{\text{hi}} = 0.5$,
331 $p_{\text{lo}} = 0.2, p_{\text{mid}} = 0.3$). GRPO’s per-sequence coefficient budget is $B_{\text{GRPO}}^{(i)} \triangleq \frac{1}{T^{(i)}} \sum_{t=1}^{T^{(i)}} |A^{(i)}| =$
332 $|A^{(i)}|$. Under EG-GRPO,
333

$$335 \quad B_{\text{EG}}^{(i)} \triangleq \frac{1}{T^{(i)}} \sum_{t=1}^{T^{(i)}} |M_t^{(i)} A^{(i)} + U_t^{(i)} \lambda \text{sg}[\bar{H}_t^{(i)}]|, \quad (7)$$

336 with $\lambda \geq 0$ and stop-gradient $\text{sg}[\cdot]$. Zeroing low-entropy updates saves roughly $p_{\text{lo}} |A^{(i)}|$ of mass
337 and reinvests it on high-entropy tokens through the additive bonus $\lambda \text{sg}[\bar{H}_t^{(i)}]$.
338

339 **Proposition 1** (Per-batch budget balance). *For a batch \mathcal{B} , choose*

$$340 \quad \lambda^* \triangleq \kappa \cdot \frac{\sum_{i \in \mathcal{B}} |A^{(i)}| \cdot \frac{1}{T^{(i)}} \sum_{t \in \mathcal{S}_{\text{lo}}^{(i,m)}} 1}{\sum_{i \in \mathcal{B}} \frac{1}{T^{(i)}} \sum_{t \in \mathcal{S}_{\text{hi}}^{(i,m)}} \text{sg}[\bar{H}_t^{(i)}]} \quad \text{with } \kappa \in (0, 1]. \quad (8)$$

341 Then $\mathbb{E}_{\mathcal{B}}[B_{\text{EG}}^{(i)}] \approx \kappa \cdot \mathbb{E}_{\mathcal{B}}[B_{\text{GRPO}}^{(i)}]$. A detailed derivation and calibration discussion are deferred
342 to Appendix B.1. Setting $\kappa = 1$ yields batch-level budget neutrality in the calibrated upper-bound
343 sense.
344

345 **Fixed-point neutrality.** Because λ^* scales with $\sum_i |A^{(i)}|$, the bonus vanishes at GRPO equilibrium
346 where group-relative advantages cancel. See Appendix B.2 for a formal proof.
347

348 **Corollary 5.1** (Preserving GRPO stationary points). *If $A^{(i)} \equiv 0$ for all i and $\lambda = \lambda^*$ with $\kappa = 1$,
349 then $\lambda^* = 0$ and $\tilde{A}_t^{(i)} \equiv 0$ for all t ; EG-GRPO reduces to the KL regularizer and preserves the
350 stationary point.*
351

352 5.4 ENTROPY-BIASED ADVANTAGE

353 We modify the broadcasted coefficient at token t of sequence i by

$$354 \quad \tilde{A}_t^{(i)} = M_t^{(i)} A^{(i)} + U_t^{(i)} \lambda \text{sg}[\bar{H}_t^{(i)}], \quad (9)$$

355 where $M_t^{(i)}$ removes reward-driven updates on the lowest-entropy 20% tokens and $U_t^{(i)}$ adds an
356 entropy bonus on the highest-entropy 50% tokens. The EG-GRPO loss is
357

$$358 \quad \mathcal{L}_{\text{EG-GRPO}}(\theta) = -\frac{1}{G} \sum_{i=1}^G \frac{1}{T^{(i)}} \sum_{t=1}^{T^{(i)}} \tilde{A}_t^{(i)} \log \pi_{\theta}(o_t^{(i)} | c, o_{<t}^{(i)}) + \beta D_{\text{KL}}(\pi_{\theta} \| \pi_{\text{ref}}). \quad (10)$$

359 β and the reference-policy KL are unchanged; low-entropy tokens are therefore governed solely by
360 KL when $M_t^{(i)} = 0$.
361

362 **Reward-shaping view.** Define a token-level pseudo-reward $\tilde{r}_t^{(i)} \triangleq r_{\text{grp}}^{(i)} \cdot M_t^{(i)} + \lambda \text{sg}[\bar{H}_t^{(i)}] \cdot U_t^{(i)}$,
363 where $r_{\text{grp}}^{(i)}$ induces $A^{(i)}$. Then $\mathbb{E}[\nabla_{\theta} \log \pi_{\theta}(o_t^{(i)} | \cdot) \tilde{r}_t^{(i)}]$ is an unbiased policy-gradient estimator
364 for an augmented objective whose baseline component is GRPO.
365

366 5.5 WHY EG-GRPO REDUCES UNCERTAINTY AND PRESERVES KNOWLEDGE

367 **(A) Targeted entropy reduction.** For small steps, the update at token t is proportional to
368 $\tilde{A}_t^{(i)} \nabla_{\theta} \log \pi_{\theta}(o_t^{(i)} | \cdot)$. On high-entropy tokens, $\tilde{A}_t^{(i)} = A^{(i)} + \lambda \text{sg}[\bar{H}_t^{(i)}]$ strengthens positive
369 updates and attenuates negative ones, lowering entropy where it is largest under softmax parameter-
370izations.
371

(B) **Stability on confident tokens.** When $M_t^{(i)} = 0$, reward-driven gradients vanish and only KL-to-reference remains, protecting confident regions from drift and preserving learned knowledge.

(C) **Proximity to GRPO equilibrium.** With $\lambda = \lambda^*$ and $\kappa = 1$, the batch-wise coefficient budget matches GRPO (Appendix B.1), and the bonus disappears at equilibrium (Appendix B.2), reallocating update mass without altering stationary points.

6 EXPERIMENTS

6.1 EXPERIMENTAL SETTINGS

Training Details. Following (Jiang et al., 2025), we train our policy on the same 6,786 text prompts drawn from T2I-CompBench (Huang et al., 2023), which contain only textual descriptions without paired images. The prompts are accompanied by structured object–attribute annotations that were automatically extracted using GPT-4o mini in prior work (Jiang et al., 2025). The policy backbone is initialized from Janus-Pro-7B (Chen et al., 2025). Optimization uses a learning rate of 1×10^{-6} and a KL coefficient $\beta = 0.01$. For the reward pipeline, we combine HPS (Wu et al., 2023) as the human-preference estimator, GroundingDINO (Liu et al., 2024) as the object detector, and GIT (Wang et al., 2022) as the VQA model. The object–relation module (ORM) is implemented by finetuning LLaVA-OneVision-7B following the procedure of Guo et al. (2025).

Benchmark. To assess the effectiveness of our approach, we rely on two established evaluation suites: T2I-CompBench (Huang et al., 2023) and WISE (Niu et al., 2025). T2I-CompBench provides 6,000 prompts designed to test compositional generalization. The benchmark covers three broad categories: attribute binding, object relations, and complex compositions, which are further split into six sub-categories, such as color/shape/texture bindings, spatial and non-spatial relations, and multi-object compositions. In contrast, WISE focuses on knowledge-intensive reasoning. It contains 1,000 prompts spanning cultural commonsense, spatial–temporal reasoning, and natural science, requiring the model to infer what specific entity or situation should appear in the image. For WISE, since the corresponding reasoning instructions from prior work (Jiang et al., 2025) were not released, we reimplemented them ourselves and provide the exact templates in the Appendix C for transparency. For both benchmarks, we otherwise follow the official evaluation protocols.

6.2 MAIN RESULTS

| Model | T2I-CompBench | | | WISE | | |
|--------------------------------------|---------------|-------|---------|---------|-----------------|---------|
| | Color | Shape | Texture | Culture | Spatio-temporal | Science |
| PixArt- α (Chen et al., 2023) | 66.90 | 49.27 | 64.77 | 45.00 | 50.00 | 46.30 |
| SD-v1.5 (Rombach et al., 2022) | 37.58 | 37.13 | 41.86 | 34.00 | 33.50 | 26.00 |
| FLUX.1-dev (Labs, 2024) | 74.07 | 57.18 | 69.22 | 48.00 | 59.00 | 45.00 |
| Emu3 (Wang et al., 2024) | 75.44 | 57.06 | 71.64 | 34.00 | 45.00 | 37.70 |
| Show-o (Xie et al., 2024) | 56.00 | 41.00 | 46.00 | 28.00 | 45.00 | 37.30 |
| Janus-Pro-7B (Chen et al., 2025) | 63.59 | 35.28 | 49.36 | 30.00 | 37.00 | 34.70 |
| T2I-R1* (Jiang et al., 2025) | 81.86 | 57.35 | 75.40 | 47.00 | 55.50 | 43.67 |
| EG-GRPO (Ours) | 84.85 | 62.61 | 78.13 | 48.00 | 55.00 | 44.00 |

Table 1: Comparison of models across **T2I-CompBench** and **WISE**. Spatio-temporal is the average of Time and Space; Science is the average of Biology, Physics, and Chemistry. *Results for T2I-R1 are reproduced under the same experimental settings for fair comparison.

As shown in Table 1, EG-GRPO achieves the strongest results on T2I-CompBench, surpassing all baselines in Color (84.85), Shape (62.61), and Texture (78.13), with particularly notable gains on Shape binding. On WISE, EG-GRPO improves over T2I-R1 in Culture (48.00 vs. 47.00) while maintaining comparable performance in Spatio-temporal (55.00 vs. 55.50) and Science (44.00 vs. 43.67). These consistent improvements demonstrate that entropy-guided updates effectively enhance compositional reasoning and robustness while preserving the stability of knowledge learned by the base model.

432

6.3 ABLATION STUDY

434 Table 2 summarizes the effect of applying entropy
 435 guidance on different token types. The full EG-
 436 GRPO model, which adds entropy bonuses to both
 437 textual CoT tokens and image tokens, achieves the
 438 best overall performance. Restricting entropy guid-
 439 ance to only CoT tokens (*w/ only sem*) or only image
 440 tokens (*w/ only tok*) leads to weaker results, indi-
 441 cating that both modalities benefit from uncertainty
 442 reduction. The baseline without entropy guidance
 443 (*w/o All*) performs worst, confirming that the pro-
 444 posed entropy-aware updates are essential for im-
 445 proved compositional generalization.

446

6.4 ANALYSIS OF ENTROPY

448 Figure 5 shows the entropy distributions of
 449 EG-GRPO and T2I-R1 on textual CoT tokens
 450 (left) and image tokens (right). EG-GRPO re-
 451 duces both the mean and variance of entropy,
 452 with a stronger effect on image tokens, indi-
 453 cating that our method concentrates updates
 454 on uncertain regions and yields more confi-
 455 dent yet stable predictions without sacrificing
 456 diversity.

457

6.5 CASE STUDY

459 As shown in Figure 6, our method con-
 460 sistently produces high-quality genera-
 461 tions across a wide range of prompts. It captures
 462 fine-grained attributes such as colors and
 463 textures with higher fidelity, preserves coherent
 464 spatial layouts in complex scenes, and main-
 465 tains stability when composing multiple ob-
 466 jects. These results confirm that entropy-
 467 guided optimization enhances both the accu-
 468 racy and consistency of text-to-image gen-
 469 eration.

470

7 CONCLUSION

473 We studied entropy dynamics in text-to-
 474 image generation, showing that CoT ex-
 475 pands exploration while reinforcement learn-
 476 ing contracts it toward stable, high-reward
 477 regions. Both the mean and variance of
 478 entropy strongly predict image quality, motivat-
 479 ing our Entropy-Guided GRPO (EG-GRPO).
 480 By protecting low-entropy tokens and focus-
 481 ing updates on high-entropy ones, EG-GRPO
 482 balances stability with structured explora-
 483 tion. Experiments on T2I-CompBench and WISE
 484 confirm its state-of-the-art performance and
 485 reduced instability, underscoring uncertainty
 486 control as a key principle for advancing text-
 487 to-image generation.

| Model | Color | Shape | Texture |
|-------------|-------|-------|---------|
| EG-GRPO | 84.85 | 62.61 | 78.13 |
| w/ only sem | 81.87 | 56.79 | 75.02 |
| w/ only tok | 84.08 | 61.80 | 77.88 |
| w/o All | 81.86 | 57.35 | 75.40 |

Table 2: Ablation results of EG-GRPO. We com-
 pare the full method with variants that apply en-
 tropy guidance only to textual CoT tokens (*w/*
 only sem), only to image tokens (*w/ only tok*), or
 not at all (*w/o All*).

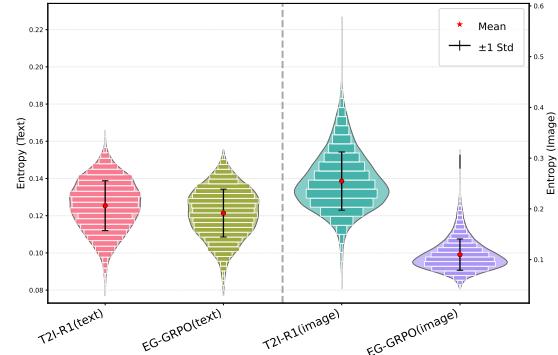


Figure 5: Entropy distributions of EG-GRPO vs. T2I-
 R1: left for textual CoT tokens, right for image tokens.

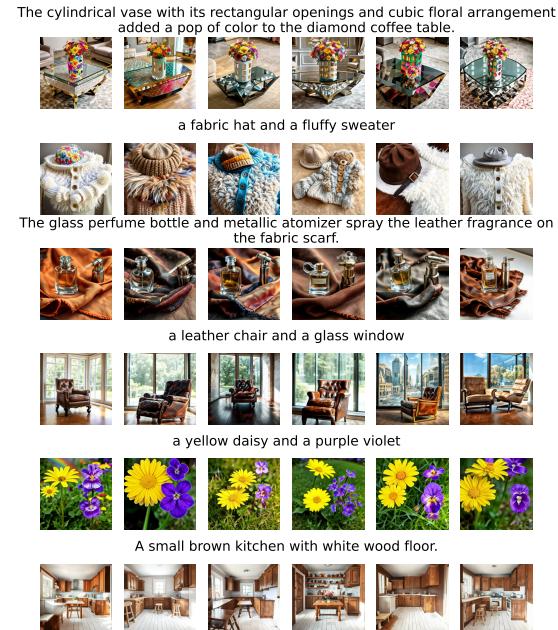


Figure 6: Qualitative case study of our method on diverse
 prompts. These results are randomly sampled.

486 REFERENCES
487

- 488 Romain Beaumont, Christoph Schuhmann, and contributors. Laion-ai/aesthetic-predictor: A linear
489 estimator on top of clip to predict the aesthetic quality of images. <https://github.com/LAION-AI/aesthetic-predictor>, 2022. Commit “main”, MIT license.
490
- 491 Junsong Chen, Jincheng Yu, Chongjian Ge, Lewei Yao, Enze Xie, Yue Wu, Zhongdao Wang, James
492 Kwok, Ping Luo, Huchuan Lu, et al. Pixart-alpha: Fast training of diffusion transformer for
493 photorealistic text-to-image synthesis. *arXiv preprint arXiv:2310.00426*, 2023.
494
- 495 Xiaokang Chen, Zhiyu Wu, Xingchao Liu, Zizheng Pan, Wen Liu, Zhenda Xie, Xingkai Yu, and
496 Chong Ruan. Janus-pro: Unified multimodal understanding and generation with data and model
497 scaling. *arXiv preprint arXiv:2501.17811*, 2025.
498
- 499 Lijie Fan, Tianhong Li, Siyang Qin, Yuanzhen Li, Chen Sun, Michael Rubinstein, Deqing Sun,
500 Kaiming He, and Yonglong Tian. Fluid: Scaling autoregressive text-to-image generative models
501 with continuous tokens. *arXiv preprint arXiv:2410.13863*, 2024.
502
- 503 Ying Fan, Olivia Watkins, Yuqing Du, Hao Liu, Moonkyung Ryu, Craig Boutilier, Pieter Abbeel,
504 Mohammad Ghavamzadeh, Kangwook Lee, and Kimin Lee. Dpok: Reinforcement learning for
505 fine-tuning text-to-image diffusion models. *Advances in Neural Information Processing Systems*,
506 36:79858–79885, 2023.
507
- 508 Zhida Feng, Zhenyu Zhang, Xintong Yu, Yewei Fang, Lanxin Li, Xuyi Chen, Yuxiang Lu, Jiaxi-
509 ang Liu, Weichong Yin, Shikun Feng, et al. Ernie-vilg 2.0: Improving text-to-image diffusion
510 model with knowledge-enhanced mixture-of-denoising-experts. In *Proceedings of the IEEE/CVF
Conference on Computer Vision and Pattern Recognition*, pp. 10135–10145, 2023.
511
- 512 Dan Friedman and Adji Bousso Dieng. The vendi score: A diversity evaluation metric for machine
513 learning. *Trans. Mach. Learn. Res.*, 2023, 2023. URL <https://openreview.net/forum?id=g970HbQyk1>.
514
- 515 Matteo Gallici and Haitz Sáez de Ocáriz Borde. Fine-tuning next-scale visual autoregressive models
516 with group relative policy optimization. *arXiv preprint arXiv:2505.23331*, 2025.
517
- 518 Shuyang Gu, Dong Chen, Jianmin Bao, Fang Wen, Bo Zhang, Dongdong Chen, Lu Yuan, and
519 Baining Guo. Vector quantized diffusion model for text-to-image synthesis. In *Proceedings of
520 the IEEE/CVF conference on computer vision and pattern recognition*, pp. 10696–10706, 2022.
521
- 522 Ziyu Guo, Renrui Zhang, Chengzhuo Tong, Zhizheng Zhao, Rui Huang, Haoquan Zhang, Manyuan
523 Zhang, Jiaming Liu, Shanghang Zhang, Peng Gao, et al. Can we generate images with cot? let’s
524 verify and reinforce image generation step by step. *arXiv preprint arXiv:2501.13926*, 2025.
525
- 526 Kaiyi Huang, Kaiyue Sun, Enze Xie, Zhenguo Li, and Xihui Liu. T2i-compbench: A com-
527 prehensive benchmark for open-world compositional text-to-image generation. *Advances in Neural
528 Information Processing Systems*, 36:78723–78747, 2023.
529
- 530 Dongzhi Jiang, Ziyu Guo, Renrui Zhang, Zhuofan Zong, Hao Li, Le Zhuo, Shilin Yan, Pheng-Ann
531 Heng, and Hongsheng Li. T2i-r1: Reinforcing image generation with collaborative semantic-level
and token-level cot. *arXiv preprint arXiv:2505.00703*, 2025.
532
- 533 Yuval Kirstain, Adam Polyak, Uriel Singer, Shahbuland Matiana, Joe Penna, and Omer Levy.
534 Pick-a-pic: An open dataset of user preferences for text-to-image generation. In Alice Oh,
535 Tristan Naumann, Amir Globerson, Kate Saenko, Moritz Hardt, and Sergey Levine (eds.),
536 *Advances in Neural Information Processing Systems 36: Annual Conference on Neural In-
537 formation Processing Systems 2023, NeurIPS 2023, New Orleans, LA, USA, December 10 -
538 16, 2023*, 2023. URL http://papers.nips.cc/paper_files/paper/2023/hash/73aacd8b3b05b4b503d58310b523553c-Abstract-Conference.html.
539
- Black Forest Labs. Flux. <https://github.com/black-forest-labs/flux>, 2024.

- 540 Junnan Li, Dongxu Li, Silvio Savarese, and Steven C. H. Hoi. BLIP-2: bootstrapping language-
 541 image pre-training with frozen image encoders and large language models. In Andreas Krause,
 542 Emma Brunskill, Kyunghyun Cho, Barbara Engelhardt, Sivan Sabato, and Jonathan Scarlett
 543 (eds.), *International Conference on Machine Learning, ICML 2023, 23-29 July 2023, Honolulu,
 544 Hawaii, USA*, volume 202 of *Proceedings of Machine Learning Research*, pp. 19730–19742.
 545 PMLR, 2023. URL <https://proceedings.mlr.press/v202/li23q.html>.
- 546 Wei Li, Xue Xu, Xinyan Xiao, Jiachen Liu, Hu Yang, Guohao Li, Zhanpeng Wang, Zhifan Feng,
 547 Qiaoqiao She, Yajuan Lyu, et al. Upainting: Unified text-to-image diffusion generation with
 548 cross-modal guidance. *arXiv preprint arXiv:2210.16031*, 2022.
- 549
- 550 Yaqi Li, Peng Chen, Mingyang Han, Bu Pi, Haoxiang Shi, Runzhou Zhao, Yang Yao, Xuan Zhang,
 551 and Jun Song. Visual-cog: Stage-aware reinforcement learning with chain of guidance for text-
 552 to-image generation. *arXiv preprint arXiv:2508.18032*, 2025.
- 553 Jiaqi Liao, Zhengyuan Yang, Linjie Li, Dianqi Li, Kevin Lin, Yu Cheng, and Lijuan Wang.
 554 Imagegen-cot: Enhancing text-to-image in-context learning with chain-of-thought reasoning.
 555 *arXiv preprint arXiv:2503.19312*, 2025.
- 556
- 557 Shilong Liu, Zhaoyang Zeng, Tianhe Ren, Feng Li, Hao Zhang, Jie Yang, Qing Jiang, Chunyuan
 558 Li, Jianwei Yang, Hang Su, et al. Grounding dino: Marrying dino with grounded pre-training
 559 for open-set object detection. In *European conference on computer vision*, pp. 38–55. Springer,
 560 2024.
- 561 Xiaoxiao Ma, Mohan Zhou, Tao Liang, Yalong Bai, Tiejun Zhao, Huaiyan Chen, and Yi Jin.
 562 Star: Scale-wise text-to-image generation via auto-regressive representations. *arXiv e-prints*,
 563 pp. arXiv–2406, 2024.
- 564
- 565 Yanting Miao, William Loh, Suraj Kothawade, Pascal Poupart, Abdullah Rashwan, and Yeqing Li.
 566 Subject-driven text-to-image generation via preference-based reinforcement learning. *Advances
 567 in Neural Information Processing Systems*, 37:123563–123591, 2024.
- 568
- 569 Yuwei Niu, Munan Ning, Mengren Zheng, Weiyang Jin, Bin Lin, Peng Jin, Jiaqi Liao, Chaoran
 570 Feng, Kunpeng Ning, Bin Zhu, et al. Wise: A world knowledge-informed semantic evaluation
 571 for text-to-image generation. *arXiv preprint arXiv:2503.07265*, 2025.
- 572
- 573 Robin Rombach, Andreas Blattmann, Dominik Lorenz, Patrick Esser, and Björn Ommer. High-
 574 resolution image synthesis with latent diffusion models. In *Proceedings of the IEEE/CVF confer-
 575 ence on computer vision and pattern recognition*, pp. 10684–10695, 2022.
- 576
- 577 NextStep Team, Chunrui Han, Guopeng Li, Jingwei Wu, Quan Sun, Yan Cai, Yuang Peng, Zheng
 578 Ge, Deyu Zhou, Haomiao Tang, et al. Nextstep-1: Toward autoregressive image generation with
 579 continuous tokens at scale. *arXiv preprint arXiv:2508.10711*, 2025.
- 580
- 581 Chengzhuo Tong, Ziyu Guo, Renrui Zhang, Wenyu Shan, Xinyu Wei, Zhenghao Xing, Hongsheng
 582 Li, and Pheng-Ann Heng. Delving into rl for image generation with cot: A study on dpo vs. grpo.
 583 *arXiv preprint arXiv:2505.17017*, 2025.
- 584
- 585 Michael Tschannen, André Susano Pinto, and Alexander Kolesnikov. Jetformer: An autoregressive
 586 generative model of raw images and text. *arXiv preprint arXiv:2411.19722*, 2024.
- 587
- 588 Jianfeng Wang, Zhengyuan Yang, Xiaowei Hu, Linjie Li, Kevin Lin, Zhe Gan, Zicheng Liu, Ce Liu,
 589 and Lijuan Wang. Git: A generative image-to-text transformer for vision and language. *arXiv
 590 preprint arXiv:2205.14100*, 2022.
- 591
- 592 Junke Wang, Zhi Tian, Xun Wang, Xinyu Zhang, Weilin Huang, Zuxuan Wu, and Yu-Gang Jiang.
 593 Simplear: Pushing the frontier of autoregressive visual generation through pretraining, sft, and rl.
 594 *arXiv preprint arXiv:2504.11455*, 2025a.
- 595
- 596 Linqing Wang, Ximing Xing, Yiji Cheng, Zhiyuan Zhao, Jiale Tao, Qixun Wang, Ruihuang Li, Xin
 597 Li, Mingrui Wu, Xinchi Deng, et al. Promptenhancer: A simple approach to enhance text-to-
 598 image models via chain-of-thought prompt rewriting. *arXiv preprint arXiv:2509.04545*, 2025b.

- 594 Shenzhi Wang, Le Yu, Chang Gao, Chujie Zheng, Shixuan Liu, Rui Lu, Kai Dang, Xionghui Chen,
 595 Jianxin Yang, Zhenru Zhang, et al. Beyond the 80/20 rule: High-entropy minority tokens drive
 596 effective reinforcement learning for llm reasoning. *arXiv preprint arXiv:2506.01939*, 2025c.
 597
- 598 Xinlong Wang, Xiaosong Zhang, Zhengxiong Luo, Quan Sun, Yufeng Cui, Jinsheng Wang, Fan
 599 Zhang, Yueze Wang, Zhen Li, Qiying Yu, et al. Emu3: Next-token prediction is all you need.
 600 *arXiv preprint arXiv:2409.18869*, 2024.
- 601 Jason Wei, Xuezhi Wang, Dale Schuurmans, Maarten Bosma, Fei Xia, Ed Chi, Quoc V Le, Denny
 602 Zhou, et al. Chain-of-thought prompting elicits reasoning in large language models. *Advances in
 603 neural information processing systems*, 35:24824–24837, 2022.
- 604 Ren-Di Wu, Yu-Yen Lin, and Huei-Fang Yang. Training-free zero-shot composed image retrieval
 605 via weighted modality fusion and similarity. In *International Conference on Technologies and
 606 Applications of Artificial Intelligence*, pp. 77–90. Springer, 2024.
 607
- 608 Xiaoshi Wu, Yiming Hao, Keqiang Sun, Yixiong Chen, Feng Zhu, Rui Zhao, and Hongsheng Li.
 609 Human preference score v2: A solid benchmark for evaluating human preferences of text-to-
 610 image synthesis. *arXiv preprint arXiv:2306.09341*, 2023.
- 611 Jinheng Xie, Weijia Mao, Zechen Bai, David Junhao Zhang, Weihao Wang, Kevin Qinghong Lin,
 612 Yuchao Gu, Zhijie Chen, Zhenheng Yang, and Mike Zheng Shou. Show-o: One single transformer
 613 to unify multimodal understanding and generation. *arXiv preprint arXiv:2408.12528*, 2024.
 614
- 615 Ling Yang, Zhaochen Yu, Chenlin Meng, Minkai Xu, Stefano Ermon, and Bin Cui. Mastering text-
 616 to-image diffusion: Recaptioning, planning, and generating with multimodal llms. In *Forty-first
 617 International Conference on Machine Learning*, 2024.
- 618 Jiahui Yu, Yuanzhong Xu, Jing Yu Koh, Thang Luong, Gunjan Baid, Zirui Wang, Vijay Vasudevan,
 619 Alexander Ku, Yafei Yang, Burcu Karagol Ayan, et al. Scaling autoregressive models for content-
 620 rich text-to-image generation. *arXiv preprint arXiv:2206.10789*, 2(3):5, 2022.
- 621 Renrui Zhang, Chengzhuo Tong, Zhizheng Zhao, Ziyu Guo, Haoquan Zhang, Manyuan Zhang,
 622 Jiaming Liu, Peng Gao, and Hongsheng Li. Let’s verify and reinforce image generation step by
 623 step. In *Proceedings of the Computer Vision and Pattern Recognition Conference*, pp. 28662–
 624 28672, 2025a.
 625
- 626 Yu Zhang, Yunqi Li, Yifan Yang, Rui Wang, Yuqing Yang, Dai Qi, Jianmin Bao, Dongdong Chen,
 627 Chong Luo, and Lili Qiu. Reasongen-r1: Cot for autoregressive image generation models through
 628 sft and rl. *arXiv preprint arXiv:2505.24875*, 2025b.
 629
- 630
- 631
- 632
- 633
- 634
- 635
- 636
- 637
- 638
- 639
- 640
- 641
- 642
- 643
- 644
- 645
- 646
- 647

648 A THE USE OF LARGE LANGUAGE MODELS
649650 In this work, we employed GPT-5 to assist with the polishing of English text. The model was used
651 primarily for improving readability, clarity, and fluency of the written content, ensuring that the final
652 manuscript meets high academic standards.
653654 B PROOFS AND DERIVATIONS FOR EG-GRPO
655656 B.1 DETAILED PROOF OF PROPOSITION 1 (PER-BATCH BUDGET BALANCE)
657658 **Setup and notation.** For sequence i , write $a^{(i)} \triangleq |A^{(i)}| \geq 0$, $s^{(i)} \triangleq \text{sign}(A^{(i)}) \in \{\pm 1\}$, and
659 $h_t^{(i)} \triangleq \bar{H}_t^{(i)} \in [0, 1]$. With the low/mid/high partitions from Section 5.3, EG-GRPO’s per-sequence
660 budget can be written as
661

662
$$B_{\text{EG}}^{(i)} = (1 - p_{\text{lo}}) a^{(i)} + \frac{1}{T^{(i)}} \sum_{t \in \mathcal{S}_{\text{hi}}^{(i,m)}} (|s^{(i)} a^{(i)} + \lambda h_t^{(i)}| - a^{(i)}). \quad (11)$$

663
664

665 **Lemma 1 (Symmetric sign averaging).** For any $a \geq 0$ and $b \geq 0$,
666

667
$$\frac{1}{2}(|a + b| + |-a + b|) = \max(a, b).$$

668

669 *Proof.* If $b \geq a$, then $|a + b| = a + b$ and $|-a + b| = b - a$, so the average is b . If $b < a$, then
670 $|a + b| = a + b$ and $|-a + b| = a - b$, so the average is a . \square
671672 **Exact decomposition.** Taking expectation over $s^{(i)}$ (assumed approximately symmetric due to
673 group-normalization), Lemma 1 yields

674
$$\mathbb{E}_{s^{(i)}}[|s^{(i)} a^{(i)} + \lambda h_t^{(i)}|] = \max(a^{(i)}, \lambda h_t^{(i)}). \quad (12)$$

675

676 Plugging into equation 11 and using $(x - a)_+ \triangleq \max(0, x - a)$,

677
$$\mathbb{E}_{s^{(i)}}[B_{\text{EG}}^{(i)}] = (1 - p_{\text{lo}}) a^{(i)} + \underbrace{\frac{1}{T^{(i)}} \sum_{t \in \mathcal{S}_{\text{hi}}^{(i,m)}} (\lambda h_t^{(i)} - a^{(i)})_+}_{\delta^{(i)}(\lambda)}. \quad (13)$$

678
679
680
681

682 Averaging over the batch \mathcal{B} gives
683

684
$$\mathbb{E}_{\mathcal{B}}[B_{\text{EG}}^{(i)}] = (1 - p_{\text{lo}}) \mathbb{E}_{\mathcal{B}}[a^{(i)}] + \mathbb{E}_{\mathcal{B}}[\delta^{(i)}(\lambda)]. \quad (14)$$

685

686 Here $\delta^{(i)}(\lambda)$ is nondecreasing in λ and equals zero at $\lambda = 0$.
687688 **Target equation for exact κ -scaling.** If one desires $\mathbb{E}_{\mathcal{B}}[B_{\text{EG}}^{(i)}] = \kappa \mathbb{E}_{\mathcal{B}}[a^{(i)}]$, then equation 14 is
689 equivalent to

690
$$\sum_{i \in \mathcal{B}} \frac{1}{T^{(i)}} \sum_{t \in \mathcal{S}_{\text{hi}}^{(i,m)}} (\lambda h_t^{(i)} - a^{(i)})_+ = (\kappa - 1 + p_{\text{lo}}) \sum_{i \in \mathcal{B}} a^{(i)}. \quad (15)$$

691
692

693 The left-hand side is continuous, nondecreasing in λ , hence admits a (numerically) unique solution
694 for any $\kappa \in (0, 1]$.
695696 **Closed-form calibration (upper-bound match).** For an implementable closed form, use $(x -
697 a)_+ \leq x$ with $x = \lambda h_t^{(i)}$:

698
$$\delta^{(i)}(\lambda) \leq \lambda \cdot \underbrace{\frac{1}{T^{(i)}} \sum_{t \in \mathcal{S}_{\text{hi}}^{(i,m)}} h_t^{(i)}}_{H_{\text{hi}}^{(i)}}.$$

699
700
701

702 Substituting this into equation 14 yields the upper bound
 703

$$705 \quad \mathbb{E}_{\mathcal{B}}[B_{\text{EG}}^{(i)}] \leq (1 - p_{\text{lo}}) \mathbb{E}_{\mathcal{B}}[a^{(i)}] + \lambda \mathbb{E}_{\mathcal{B}}[H_{\text{hi}}^{(i)}]. \quad (16)$$

708 Matching the *upper bound* to $\kappa \mathbb{E}_{\mathcal{B}}[a^{(i)}]$ leads to the batch-calibrated choice
 709

$$712 \quad \lambda^* \triangleq \kappa \cdot \frac{\sum_{i \in \mathcal{B}} a^{(i)} \cdot \frac{1}{T^{(i)}} \sum_{t \in \mathcal{S}_{\text{lo}}^{(i,m)}} 1}{\sum_{i \in \mathcal{B}} \frac{1}{T^{(i)}} \sum_{t \in \mathcal{S}_{\text{hi}}^{(i,m)}} h_t^{(i)}} = \kappa \cdot \frac{\sum_{i \in \mathcal{B}} |A^{(i)}| \cdot \frac{|\mathcal{S}_{\text{lo}}^{(i,m)}|}{T^{(i)}}}{\sum_{i \in \mathcal{B}} \frac{1}{T^{(i)}} \sum_{t \in \mathcal{S}_{\text{hi}}^{(i,m)}} \text{sg}[\bar{H}_t^{(i)}]}, \quad (17)$$

716 which is exactly equation 8. This calibration equates “saved” low-entropy budget with “reinvested”
 717 high-entropy mass in an upper-bound sense; see Remarks below.
 718

721 **Remarks on calibration accuracy.** (i) Define the nonnegative discrepancy
 722

$$724 \quad \varepsilon^{(i)}(\lambda) \triangleq \lambda H_{\text{hi}}^{(i)} - \delta^{(i)}(\lambda) = \frac{1}{T^{(i)}} \sum_{t \in \mathcal{S}_{\text{hi}}^{(i,m)}} \min(\lambda h_t^{(i)}, a^{(i)}).$$

728 It vanishes as $\lambda h_t^{(i)} \gg a^{(i)}$ for most high-entropy tokens, and is statistically damped by batch aver-
 729 aging. (ii) For *exact* κ -scaling, one may solve equation 15 via a 1D root finder; we keep equation 8
 730 for simplicity and stability. (iii) With $\kappa = 1$, equation 8 delivers batch-level budget neutrality in the
 731 calibrated upper-bound sense; empirically it closely tracks neutrality while avoiding per-step root
 732 solving.
 733

735 **Conclusion.** Combining equation 14–equation 16 with the calibration above yields $\mathbb{E}_{\mathcal{B}}[B_{\text{EG}}^{(i)}] \approx$
 736 $\kappa \mathbb{E}_{\mathcal{B}}[B_{\text{GRPO}}^{(i)}]$, as stated in Proposition 1.
 737

740 B.2 PROOF OF COROLLARY 5.1 (FIXED-POINT NEUTRALITY)

744 Suppose $A^{(i)} \equiv 0$ for all i in a batch. Then $a^{(i)} = 0$, and the numerator of equation 8 vanishes,
 745 yielding $\lambda^* = 0$ for any $\kappa \in (0, 1]$. By equation 9, $\tilde{A}_t^{(i)} \equiv 0$ for all tokens t , so the loss equation 10
 746 reduces to the reference-policy KL term. Hence any GRPO stationary point remains stationary under
 747 EG-GRPO, proving the corollary.
 748

750 B.3 OPTIONAL EXACT PER-BATCH SCALING (IMPLEMENTATION NOTE)

753 If precise κ -scaling is required, solve the monotone equation equation 15 for λ by bisection or
 754 Newton’s method. This guarantees $\mathbb{E}_{\mathcal{B}}[B_{\text{EG}}^{(i)}] = \kappa \mathbb{E}_{\mathcal{B}}[B_{\text{GRPO}}^{(i)}]$ *exactly* per batch, at the cost of a 1D
 755 search.

756 C INSTRUCTION FOR WISE
757758 WISE Instruction
759

760 You are asked to write a concise text description to guide the generation of an image based
761 on this prompt: “{}”. Provide a brief, precise visualization of all elements in the prompt.
762 Your description should:
763 1. Include every object mentioned.
764 2. Specify visual attributes (color, number, shape, texture) if given.
765 3. Clarify spatial or relational positioning if specified.
766 4. Be concise (≤ 50 words) but include the most common features and states inferred from
767 real-world knowledge.
768 5. Apply real-world knowledge (cultural, religious, temporal, spatial, biological, physical, or
769 chemical reasoning) and select only the single most relevant aspect that naturally enhances
770 the original prompt to infer context (e.g., season, appearance, identity, cultural usage, or
771 natural state) and reflect it in the objects. Use direct, widely accepted interpretations; include
772 cultural or religious meanings only when they are common real-world associations. Avoid
773 abstract or purely metaphorical interpretations.
774 6. Emphasize the current state of each object individually, as inferred from its environment
775 or context. Reason separately for each object, considering temporal, cultural, or physical
776 factors, and prioritize states logically implied by the prompt.
777 7. If multiple objects are present, reason from each object’s inherent physical or chemical
778 properties and their interactions with the environment and with all other objects. Ensure that
779 the inferred state, behavior, and interaction of every single object is logically correct and
780 consistent with real-world rules, and clearly describe all differences and interactions relative
781 to each other.
782 8. Ensure realism and aesthetic quality: all objects and interactions must follow real-world
783 rules and appear visually consistent and appealing.
784 9. Do not omit objects explicitly mentioned, or add ones not specified or logically inferred.
785 10. Always preserve and emphasize the original objects and scene as the primary focus.
786 11. Always output a complete natural language description, never an image or symbolic
787 shorthand.

788 D COMPUTATIONAL EFFICIENCY AND CONVERGENCE
789 ANALYSIS
790

791 To address concerns regarding the computational overhead of the proposed EG-GRPO, specifically
792 the token-level entropy computation, percentile thresholding, and batch-level bonus recalibration,
793 we conducted a rigorous benchmarking comparison of our method against the baseline T2I-R1.
794

795 D.1 WALL-CLOCK TIME AND MEMORY OVERHEAD
796

797 We measured the training step time and peak GPU memory usage on NVIDIA A100 GPUs under
798 identical experimental settings. As shown in Table 3, EG-GRPO introduces negligible overhead.
799

| Metric | T2I-R1 (Baseline) | EG-GRPO (Ours) | Overhead |
|-----------------|-------------------|----------------|----------|
| Step Time (s) | 50.20 | 50.88 | +1.35% |
| GPU Memory (GB) | 30.79 | 31.17 | +0.38 GB |

800
801 Table 3: Computational overhead comparison between T2I-R1 and EG-GRPO. The overhead intro-
802 duced by our entropy-guided mechanism is marginal in both time and memory.
803

804
805 The minimal increase in wall-clock time ($\sim 1.35\%$) suggests that the complexity of entropy opera-
806 tions ($O(L \cdot V)$ for sequence length L and vocabulary size V) is trivial compared to the computational
807 cost of the model’s forward and backward passes. Importantly, this overhead ratio is expected to de-

crease further as model scale increases (e.g., larger hidden dimensions), as the entropy calculation depends only on the output logits and not the model depth or parameter count.

D.2 CONVERGENCE EFFICIENCY

While the per-step cost is marginally higher, EG-GRPO exhibits significantly better sample efficiency. We compared the reward progression of both models over the same number of training steps. As detailed in Table 4, EG-GRPO consistently achieves higher rewards earlier in the training process.

| Step | 100 | 200 | 400 | 600 | 800 |
|------------------|---------------|---------------|---------------|---------------|---------------|
| Reward (T2I-R1) | 2.1105 | 2.1130 | 2.1208 | 2.1339 | 2.1340 |
| Reward (EG-GRPO) | 2.1411 | 2.1836 | 2.1968 | 2.2075 | 2.2117 |

Table 4: Convergence comparison: Reward scores at different training steps. EG-GRPO achieves higher performance consistently, indicating a superior time-to-convergence ratio.

This superior convergence efficiency outweighs the slight computational overhead, making EG-GRPO a more practical choice for large-scale training.

E ANALYSIS ON DIVERSITY

In this section, we address the trade-off between entropy reduction and generative diversity. While entropy is a measure of uncertainty, reducing it raises the question of whether the model's creative expressivity and output diversity are negatively impacted. To investigate this, we employ the **Vendi Score** (Friedman & Dieng, 2023), a reference-free metric designed to quantify diversity in machine learning models.

Dynamics of Diversity during RL. The RL naturally contracts the exploration space to focus on high-reward regions. As shown in Table 5, our analysis of the GRPO training dynamics confirms this expected behavior: as the model optimizes for reward-aligned fidelity from Step 100 to 800, the diversity naturally shows a slight decrease.

| Step | 100 | 200 | 400 | 600 | 800 |
|--------------------------------|--------|--------|--------|--------|--------|
| Diversity (Vendi Score) | 2.7305 | 2.7233 | 2.7212 | 2.7151 | 2.7159 |

Table 5: Evolution of diversity during the GRPO training process. As the model converges towards high-reward regions, a slight reduction in diversity is observed, reflecting the exploration-exploitation trade-off.

Preserving Diversity at Similar Quality. Crucially, our method targets *instability* rather than *semantic diversity*. To verify this, we compared EG-GRPO against the baseline (T2I-R1) on a filtered subset of generated samples where both models achieved similar quality scores (defined as $|\Delta\text{Quality}| < 0.1$). The quality score is an aggregate of BLIP-2 (Li et al., 2023), LAION-Aesthetics (Beaumont et al., 2022), and PickScore (Kirstain et al., 2023).

| Metric | EG-GRPO (Ours) | T2I-R1 (Baseline) |
|--------------------------------|----------------|-------------------|
| Quality Mean | 13.86 | 13.85 |
| Diversity (Vendi Score) | 2.593 | 2.592 |

Table 6: Comparison of Diversity (Vendi Score) between EG-GRPO and T2I-R1 on a subset of samples with similar quality scores. Our method maintains diversity comparable to the baseline.

As shown in Table 6, EG-GRPO maintains a Vendi Score (2.593) virtually identical to the baseline (2.592) under the same quality constraints. This indicates that the Entropy Bonus (Section 5.3)

864 successfully preserves valid exploration pathways while suppressing “bad” uncertainty that leads to
 865 instability, rather than collapsing the model into a single mode.
 866

867 F QUALITATIVE DIVERSITY CASES

870 To qualitatively compare the diversity differences between T2I-R1 and EG-GRPO, we generate 20
 871 samples for each of three identical prompts and concatenate them for side-by-side visualization, as
 872 shown in Figure 7.

874 The soft, plush texture of the teddy bear was a comforting companion
 875 for the children at bedtime.



887 The triangular shelf with its crescent curves and circular brackets held
 888 items in the pentagonal hallway.



900 a chicken hidden by a couch



912 Figure 7: Comparison of generation diversity for **T2I-R1** (left) and **EG-GRPO** (right).
 913

914
 915
 916
 917