RePrompt: Reasoning-Augmented Reprompting for Text-to-Image Generation via Reinforcement Learning

Mingrui Wu[1,3] ∗, Lu Wang[2] †, Pu Zhao[2] †, Fangkai Yang[2] †, Jianjin Zhang[2], Jianfeng Liu[2], Yuefeng Zhan[2], Weihao Han[2], Hao Sun[2], Jiayi Ji[1] ‡, Xiaoshuai Sun[1], Qingwei Lin[2], Weiwei Deng[2], Dongmei Zhang[2], Feng Sun[2], Qi Zhang[2], Rongrong Ji[1]

1 Key Laboratory of Multimedia Trusted Perception and Efficient Computing, Ministry of Education of China, Xiamen University, 361005, P.R. China.

2 Microsoft. 3 Zhongguancun Academy, Beijing, China.100094.

Introduction & Method

Challenge: Existing models often struggle to faithfully capture user intentions from short and underspecified prompts.

.

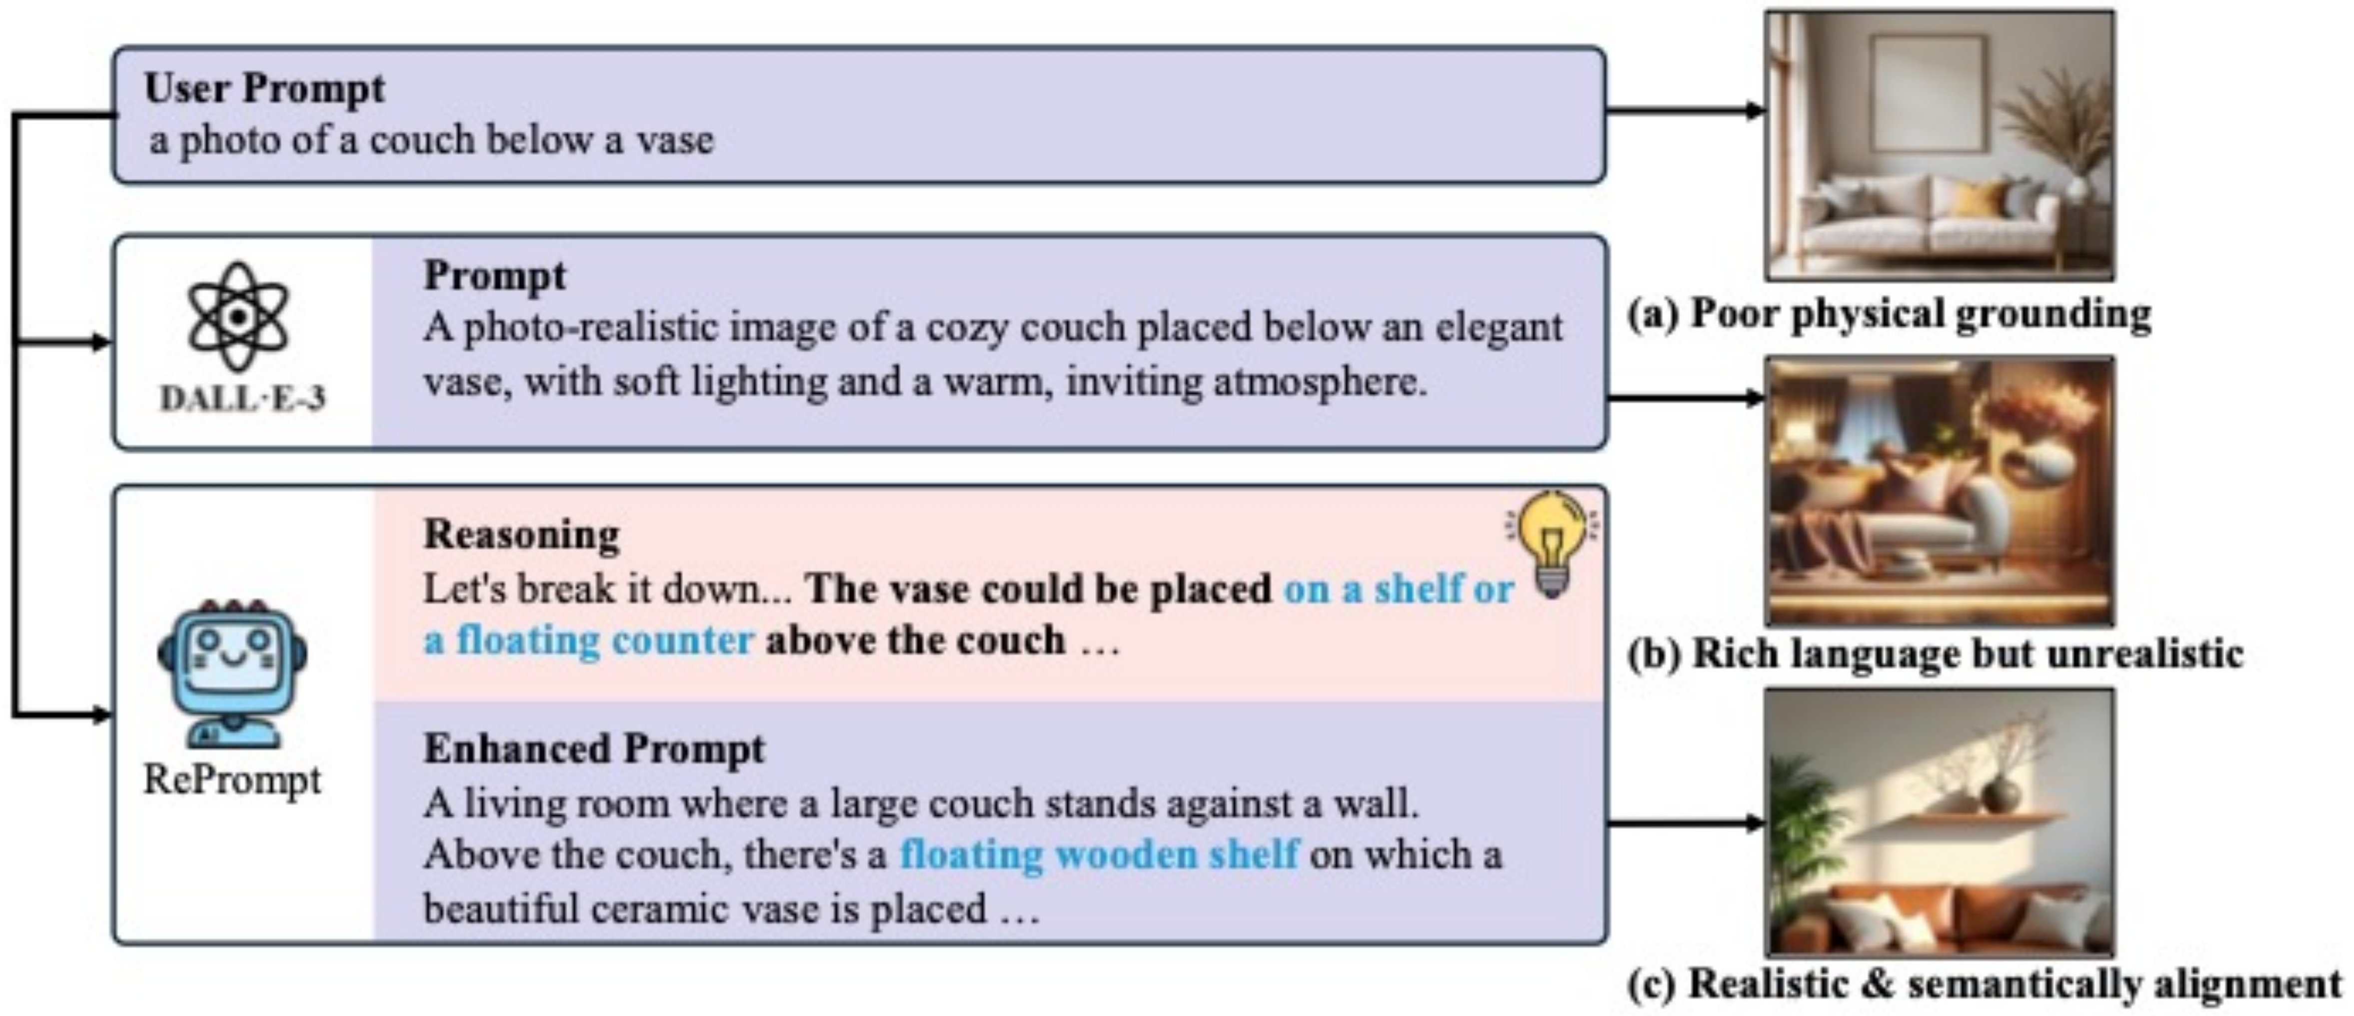

Our Contribution: We propose RePrompt, a novel reprompting framework that introduces explicit reasoning into the prompt enhancement process via reinforcement learning. Instead of relying on handcrafted rules or stylistic rewrites, our method trains a language model to generate structured, self-reflective prompts by optimizing for image-level outcomes.

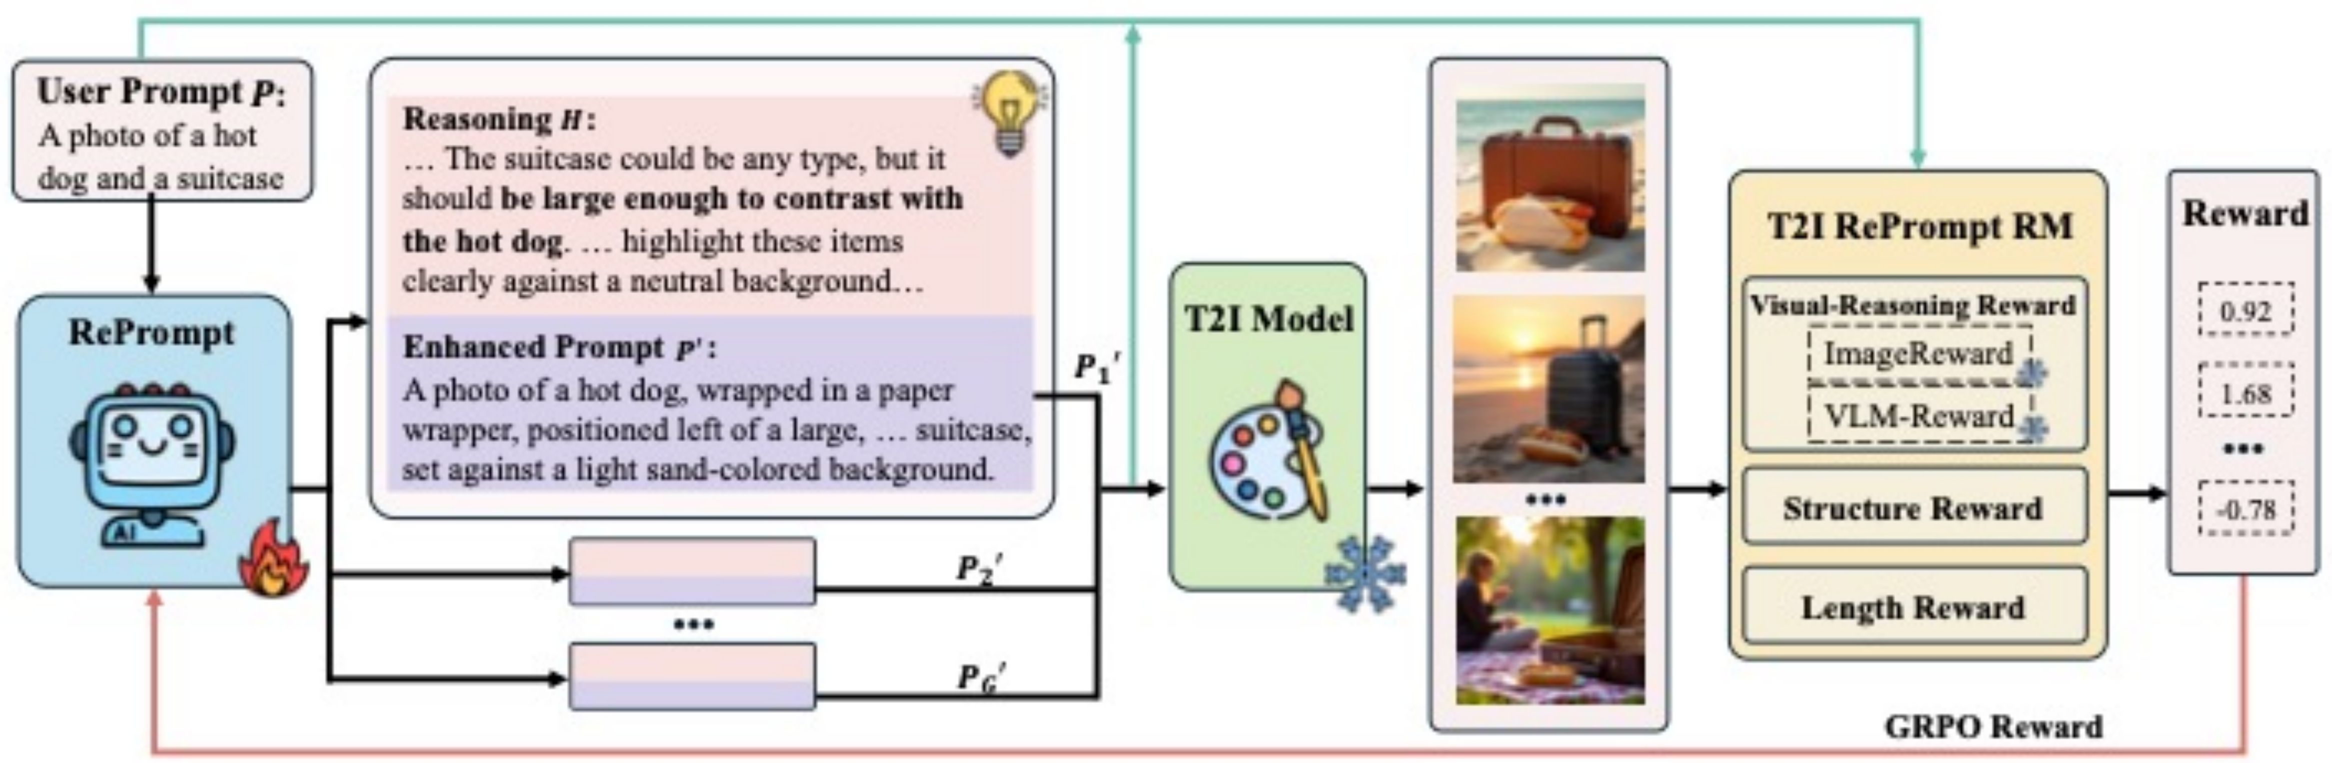

Results

- **Our approach shows substantial gains in spatial position understanding over Qwen2.5 3B-enhanced baselines.**

Method	Single object	Two object	Counting	Colors	Position	Attribute binding	Overall ↑
FLUX Labs (2024)	0.99	0.79	0.75	0.78	0.18	0.45	0.66
+ Promptist Hao et al. (2023)	0.98	0.72	0.70	0.78	0.21	0.44	0.66
+ PAG Yun et al. (2025)	0.97	0.74	0.73	0.80	0.36	0.46	0.69
+ GPT4	0.99	0.79	0.68	0.84	0.51	0.52	0.72
+ GPT5	1.00	0.81	0.70	0.85	0.51	0.51	0.73
+ GPTo3	0.98	0.76	0.65	0.83	0.48	0.50	0.70
+ Deepseek-r1	1.00	0.81	0.56	0.78	0.47	0.43	0.67
+ Qwen2.5 7B	0.99	0.83	0.62	0.84	0.36	0.51	0.69
+ Qwen2.5 3B	0.99	0.84	0.63	0.81	0.35	0.48	0.68
+ Ours (train w/ FLUX)	0.98	**0.87**	**0.77**	**0.85**	**0.62**	**0.49**	**0.76**
Improvement over Qwen2.5 3B	-1.0%	+3.6%	+22.2%	+4.9%	**+77.1%**	+2.1%	+11.8%
SD3 Esser et al. (2024)	1.00	0.85	0.62	0.88	0.22	0.58	0.69
+ Promptist	0.99	0.84	0.66	0.84	0.45	0.52	0.69
+ PAG	0.99	0.85	0.68	0.85	0.49	0.53	0.71
+ GPT4	1.00	0.84	0.51	0.85	0.48	0.54	0.70
+ GPT5	1.00	0.85	0.53	0.85	0.47	0.50	0.70
+ GPTo3	0.99	0.82	0.50	0.84	0.41	0.52	0.68
+ Deepseek-r1	0.99	0.82	0.53	0.80	0.44	0.46	0.67
+ Qwen2.5 7B	1.00	0.82	0.49	0.85	0.34	0.58	0.68
+ Qwen2.5 3B	1.00	0.86	0.53	0.84	0.33	0.55	0.68
+ Ours (train w/ FLUX)	0.99	**0.86**	0.60	0.86	**0.59**	**0.60**	**0.75**
Improvement over Qwen2.5 3B	-1.0%	0.0%	+13.2%	+2.4%	**+78.8%**	+9.1%	+10.3%
Pixart-Σ Chen et al. (2024b)	0.99	0.60	0.47	0.81	0.10	0.26	0.54
+ Promptist	0.98	0.60	0.49	0.80	0.20	0.27	0.55
+ PAG	0.98	0.63	0.52	0.80	0.28	0.29	0.56
+ GPT4	0.96	0.67	0.48	0.84	0.36	0.31	0.60
+ GPT5	0.97	0.68	0.49	0.84	0.37	0.31	0.61
+ GPTo3	0.96	0.67	0.48	0.83	0.35	0.31	0.60
+ Deepseek-r1	0.99	0.63	0.43	0.78	0.24	0.27	0.56
+ Qwen2.5 7B	0.96	0.67	0.43	0.83	0.20	0.32	0.57
+ Qwen2.5 3B	0.99	0.68	0.48	0.82	0.18	0.32	0.58
+ Ours (train w/ FLUX)	0.98	0.64	**0.56**	**0.81**	**0.40**	**0.35**	**0.62**
Improvement over Qwen2.5 3B	-1.0%	-5.9%	+16.7%	-1.2%	**+122.2%**	+9.4%	+6.9%

- **Our RePrompt improves spatial layout and object relations by generating enhanced prompts with explicit reasoning, leading to more faithful compositions.**

- **Our approach consistently improves performance across most aspects, particularly in Spatial compositions.**

Method	Color	Shape	Texture	Spatial	Numeracy	Complex
FLUX	0.7223	0.4796	0.6522	0.2494	0.6101	0.3616
+ Qwen2.5 3B	0.7149	0.5103	0.6012	0.2579	0.5982	0.3325
+ Ours (train w/ FLUX)	**0.7501**	**0.5276**	0.6515	**0.3301**	**0.6499**	**0.3721**
SD3	0.7941	0.5812	0.7224	0.2815	0.5871	0.3714
+ Qwen2.5 3B	0.7227	0.5478	0.6581	0.2549	0.5934	0.3307
+ Ours (train w/ FLUX)	0.7866	**0.5891**	0.7184	**0.3315**	**0.6289**	**0.3744**
Pixart-Σ	0.5682	0.4717	0.5622	0.2497	0.5366	0.3655
+ Qwen2.5 3B	0.6618	0.4814	0.5662	0.2481	0.5443	0.3335
+ Ours (train w/ FLUX)	**0.6665**	**0.5011**	**0.6190**	**0.2913**	**0.5716**	**0.3680**

- **Our method consistently improves the compositional understanding across all base models.**