

ReVersion: Diffusion-Based Relation Inversion from Images

ZIQI HUANG*, S-Lab, Nanyang Technological University, Singapore
 TIANXING WU*, S-Lab, Nanyang Technological University, Singapore
 YUMING JIANG, S-Lab, Nanyang Technological University, Singapore
 KELVIN C.K. CHAN, S-Lab, Nanyang Technological University, Singapore
 ZIWEI LIU†, S-Lab, Nanyang Technological University, Singapore

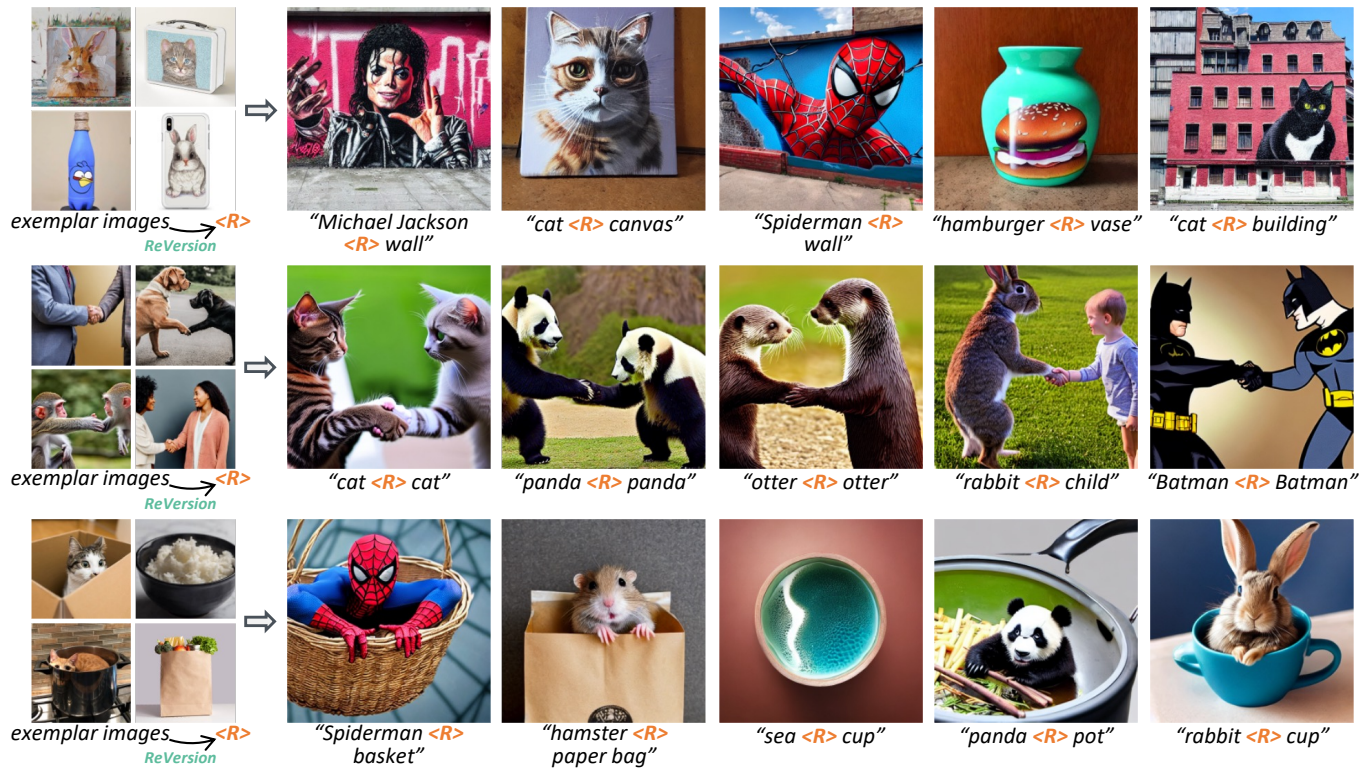


Fig. 1. We propose a new task, **Relation Inversion**: Given a few exemplar images, where a relation co-exists in every image, we aim to find a relation prompt $\langle R \rangle$ to capture this interaction, and apply the relation to new entities to synthesize new scenes. The above images are generated by our **ReVersion Framework**.

Diffusion models gain increasing popularity for their generative capabilities. Recently, there have been surging needs to generate customized images by inverting diffusion models from exemplar images, and existing inversion methods mainly focus on capturing object **appearances** (i.e., the “look”). However, how to invert object **relations**, another important pillar in the visual world, remains unexplored. In this work, we propose the **Relation**

Inversion task, which aims to learn a specific relation (represented as “relation prompt”) from exemplar images. Specifically, we learn a relation prompt with a frozen pre-trained text-to-image diffusion model. The learned relation prompt can then be applied to generate relation-specific images with new objects, backgrounds, and styles. To tackle the *Relation Inversion* task, we propose the **ReVersion Framework**. Specifically, we propose a novel “relation-steering contrastive learning” scheme to steer the relation prompt towards relation-dense regions, and disentangle it away from object appearances. We further devise “relation-focal importance sampling” to emphasize high-level interactions over low-level appearances (e.g., texture, color). To comprehensively evaluate this new task, we contribute the **ReVersion Benchmark**, which provides various exemplar images with diverse relations. Extensive experiments validate the superiority of our approach over existing methods across a wide range of visual relations. Our proposed task and method could be good inspirations for future research in various domains like generative inversion, few-shot learning, and visual relation detection.

*Equal contributions

†Corresponding author

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than the author(s) must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.

SA Conference Papers '24, December 3–6, 2024, Tokyo, Japan

© 2024 Copyright held by the owner/author(s). Publication rights licensed to ACM.

ACM ISBN 979-8-4007-1131-2/24/12...\$15.00

<https://doi.org/10.1145/3680528.3687658>

CCS Concepts: • **Computing methodologies** → **Computer vision**.

Additional Key Words and Phrases: Image generation, relation modeling, diffusion model

ACM Reference Format:

Ziqi Huang, Tianxing Wu, Yuming Jiang, Kelvin C.K. Chan, and Ziwei Liu. 2024. ReVersion: Diffusion-Based Relation Inversion from Images. In *SIGGRAPH Asia 2024 Conference Papers (SA Conference Papers '24)*, December 3–6, 2024, Tokyo, Japan. ACM, New York, NY, USA, 24 pages. <https://doi.org/10.1145/3680528.3687658>

1 INTRODUCTION

Recently, text-to-image (T2I) diffusion models [Ramesh et al. 2022; Rombach et al. 2022; Saharia et al. 2022a] have shown promising results and enabled subsequent explorations of various generative tasks. There have been several explorations [Chen et al. 2023a; Gal et al. 2022; Jia et al. 2023; Kumari et al. 2022; Li et al. 2023a; Ruiz et al. 2022; Wei et al. 2023] on the *appearance inversion* task. Specifically, given a few images of a specific object (e.g., a cat statue), *appearance inversion* learns to map a “new word” to this concept via the text-to-image model. The “new word” can then be used in prompts to generate new images that contain this concept. While existing methods have made substantial progress in capturing object *appearances*, such exploration for *relations* between objects is rare.

In this paper, we study the **Relation Inversion** task, whose objective is to learn a *relation* that co-exists in the given exemplar images. Specifically, with objects in each exemplar image following a specific relation, we aim to obtain a relation prompt in the text embedding space of the pre-trained text-to-image diffusion model. By composing the relation prompt with user-devised text prompts, users are able to synthesize images using the corresponding relation, with new objects, styles, and backgrounds, etc. Studying *Relation Inversion* not only addresses a critical gap in text-to-image model inversion tasks but also paves the way for deeper understanding and generation of relation-rich visual content.

The *Relation Inversion* task is intrinsically different from existing appearance inversion tasks, and thus poses unique challenges. Appearance inversion [Chen et al. 2023a; Gal et al. 2022; Jia et al. 2023; Kumari et al. 2022; Li et al. 2023a; Ruiz et al. 2022; Wei et al. 2023] focuses on capturing the look of a specific entity, thus the commonly used pixel-level reconstruction loss is typically adequate to learn a prompt that encapsulates the shared information among exemplar images. In contrast, *relation* is a more abstract visual concept, and a pixel-wise loss alone is insufficient for precise extraction of the intended relation. Consequently, linguistic and visual priors are needed to accurately represent these high-level relation concepts.

To this end, we propose the **ReVersion Framework** to tackle the *Relation Inversion* problem. First, we exploit linguistic priors to steer the relation prompt in the text embedding space. Specifically, we found that in the text embedding space of Stable Diffusion, embeddings are generally clustered according to their Part-of-Speech (POS), as shown in Figure 2. Also, the concept of “relation” is related to prepositional words. For example, the relation “rides on” is semantically related to the prepositions “atop”, “above”, and “below”;

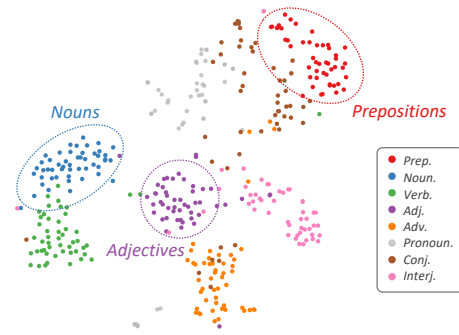


Fig. 2. **Part-of-Speech (POS) Clustering**. We use t-SNE [Van der Maaten and Hinton 2008] to visualize word distribution in CLIP’s input text embedding space, where $\langle R \rangle$ is optimized in our ReVersion framework. We observe that words of the same Part-of-Speech (POS) are closely clustered together, while words of different POS are generally at a distance from each other.

the relation “being contained within” is semantically related to “inside”, “around”, “in”, and “including”. This together with the POS clustering observation motivate us to steer the relation prompts towards the prepositional word cluster. Notably, we design a novel *relation-steering contrastive learning scheme* to steer the relation prompt towards a relatively relation-dense region in the text embedding space. A set of basis prepositions are used as positive samples to pull the relation prompt, while words of other POS (e.g., nouns, adjectives) in text descriptions are regarded as negatives so that the semantics related to object appearances are disentangled away.

Second, to encourage attention on object interactions, we devise a *relation-focal importance sampling* strategy. During the optimization process, we emphasize high-level interactions over relatively lower-level details (e.g., color and texture of objects), effectively leading to better *Relation Inversion* results.

As the first attempt in this direction, we further contribute the **ReVersion Benchmark**, which provides various exemplar images with diverse relations, from simple spatial arrangements to complex interactive behaviours. The benchmark serves as an evaluation tool for future research in *Relation Inversion*. Results on a variety of relations demonstrate the effectiveness of our ReVersion Framework.

Our contributions are summarized as follows:

- We study a new problem, **Relation Inversion**, which requires learning a relation prompt for a relation that co-exists in several exemplar images. While existing T2I inversion works mainly focus on capturing appearances, we take the initiative to explore relation, an under-explored yet important pillar in the visual world.
- We propose the **ReVersion Framework**, where the *relation-steering contrastive learning scheme* steers relation prompt using linguistic priors, and effectively disentangles the learned relation away from object appearances. The *relation-focal importance sampling* further emphasizes high-level relations over low-level details.
- We contribute the **ReVersion Benchmark**, which serves as a diagnostic and benchmarking tool for the new task of *Relation Inversion*.

2 RELATED WORK

Diffusion Models. Diffusion models [Gu et al. 2022; Ho et al. 2020; Rombach et al. 2022; Sohl-Dickstein et al. 2015; Song et al. 2021a,b] have become a mainstream approach for image synthesis [Dhariwal and Nichol 2021; Esser et al. 2021; Meng et al. 2022] apart from GANs [Goodfellow et al. 2014], and have shown success in various domains such as video generation [Blattmann et al. 2023; Harvey et al. 2022; He et al. 2022; Ho et al. 2022b; Singer et al. 2022; Villegas et al. 2022; Wu et al. 2022], image restoration [Ho et al. 2022a; Saharia et al. 2022b], and many more [Amit et al. 2021; Austin et al. 2021; Baranchuk et al. 2022; Graikos et al. 2022]. Diffusion models are usually trained using score-matching objectives [Hyvärinen and Dayan 2005; Vincent 2011] at various noise levels, and sampling is done via iterative denoising. Text-to-Image (T2I) diffusion models [Esser et al. 2021; Gu et al. 2022; Jiang et al. 2022; Nichol et al. 2021; Ramesh et al. 2022; Rombach et al. 2022; Saharia et al. 2022a] demonstrated impressive results in converting user-provided text descriptions into images. Motivated by their success, we build our framework on a state-of-the-art T2I diffusion model, Stable Diffusion [Rombach et al. 2022].

Relation Modeling. Relation modeling has been explored in discriminative tasks such as scene graph generation [Ji et al. 2020; Krishna et al. 2017; Shang et al. 2017; Xu et al. 2017; Yang et al. 2022, 2023] and visual relationship detection [Lu et al. 2016; Yu et al. 2017; Zhuang et al. 2017]. These works aim to detect visual relations between objects in given images and classify them into a predefined, closed-set of relations. However, the finite relation category set intrinsically limits the diversity of captured relations. In contrast, *Relation Inversion* regards relation modeling as a generative task, aiming to capture arbitrary, open-world relations from exemplar images and apply the resulting relation for content creation.

Diffusion-Based Inversion. Given a pre-trained T2I diffusion model, *inversion* aims to find a text embedding vector to express the concepts in the given exemplar images, via optimization-based [Alaluf et al. 2023; Choi et al. 2023; Gal et al. 2022; Han et al. 2023; Hu et al. 2022; Kawar et al. 2022; Kumari et al. 2022; Li et al. 2023b; Ruiz et al. 2022; Voynov et al. 2023], encoder-based [Jia et al. 2023; Ma et al. 2023; Wei et al. 2023; Xu et al. 2023; Ye et al. 2023; Zhou et al. 2023], or hybrid [Arar et al. 2023; Chen et al. 2023b; Gal et al. 2023; Gong et al. 2023; Li et al. 2023a; Ruiz et al. 2024] approaches. For example, given several images of a particular “*cat statue*”, Textual Inversion [Gal et al. 2022] learns a new word to describe its appearance - finding a vector in Latent Diffusion Model (LDM) [Rombach et al. 2022]’s text embedding space, so that the new word can be composed into new sentences to achieve personalized creation. Rather than inverting the appearance information (e.g., color, texture), our proposed *Relation Inversion* task extracts high-level object *relations* from exemplar images, which can be harder as it requires comprehending image compositions and object relationships.

3 THE RELATION INVERSION TASK

Relation Inversion aims to extract the common relation $\langle R \rangle$ from several exemplar images. Let $\mathcal{I} = \{I_1, I_2, \dots, I_n\}$ be a set of exemplar images, and $E_{i,A}$ and $E_{i,B}$ be two dominant entities in image I_i . In *Relation Inversion*, we assume that the entities in each exemplar

image interacts with each other through a common relation R . A set of coarse descriptions $C = \{c_1, c_2, \dots, c_n\}$ is associated to the exemplar images, where “ $c_i = E_{i,A} \langle R \rangle E_{i,B}$ ” denotes the caption corresponding to image I_i . Our objective is to optimize the relation prompt $\langle R \rangle$ such that the co-existing relation can be accurately represented by the optimized prompt.

An immediate application of *Relation Inversion* is relation-specific text-to-image synthesis. Once the prompt is acquired, one can generate images with novel objects interacting with each other following the specified relation. More generally, this task reveals a new direction of inferring relations from a set of exemplar images. This could potentially inspire future research in representation learning, few-shot learning, visual relation detection, scene graph generation, and many more.

4 THE REVERSION FRAMEWORK

4.1 Preliminaries

Stable Diffusion. Diffusion models are a class of generative models that gradually denoise the Gaussian prior \mathbf{x}_T to the data \mathbf{x}_0 (e.g., a natural image). The commonly used training objective L_{DM} [Ho et al. 2020] is:

$$L_{DM}(\theta) := \mathbb{E}_{t, \mathbf{x}_0, \epsilon} [\|\epsilon - \epsilon_\theta(\mathbf{x}_t, t)\|^2], \quad (1)$$

where \mathbf{x}_t is a noisy image constructed by adding noise $\epsilon \sim \mathcal{N}(\mathbf{0}, \mathbf{I})$ to the natural image \mathbf{x}_0 , and the network $\epsilon_\theta(\cdot)$ is trained to predict the added noise. To sample data \mathbf{x}_0 from a trained diffusion model $\epsilon_\theta(\cdot)$, we iteratively denoise \mathbf{x}_t from $t = T$ to $t = 0$ using the predicted noise $\epsilon_\theta(\mathbf{x}_t, t)$ at each timestep t .

LDM [Rombach et al. 2022], the predecessor of Stable Diffusion, mainly introduced two changes to the vanilla diffusion model [Ho et al. 2020]. First, instead of directly modeling the natural image distribution, LDM models images’ projections in autoencoder’s compressed latent space. Second, LDM enables text-to-image generation by feeding encoded text input to the UNet [Ronneberger et al. 2015] $\epsilon_\theta(\cdot)$. The LDM loss is:

$$L_{LDM}(\theta) := \mathbb{E}_{t, \mathbf{x}_0, \epsilon} [\|\epsilon - \epsilon_\theta(\mathbf{x}_t, t, \tau_\theta(c))\|^2], \quad (2)$$

where \mathbf{x} is the autoencoder latents for images, and $\tau_\theta(\cdot)$ is the text encoder that encodes the text descriptions c into the text embedding space. Stable Diffusion extends LDM by training on the larger LAION dataset [Schuhmann et al. 2022], with some architectural and training changes.

Inversion on Text-to-Image Diffusion Models. Existing inversion methods focus on appearance inversion. Given several images that all contain a specific entity, they [Gal et al. 2022; Kumari et al. 2022; Ruiz et al. 2022] find a text embedding V^* for the pre-trained T2I model. The obtained V^* can then be used as a new word to generate this entity in different scenarios.

In this work, we aim to capture object relations instead. Given several exemplar images which share a common relation R , we aim to find a relation prompt $\langle R \rangle$ to capture this relation, such that “ $E_A \langle R \rangle E_B$ ” can be used to generate an image where E_A and E_B interact via relation R .

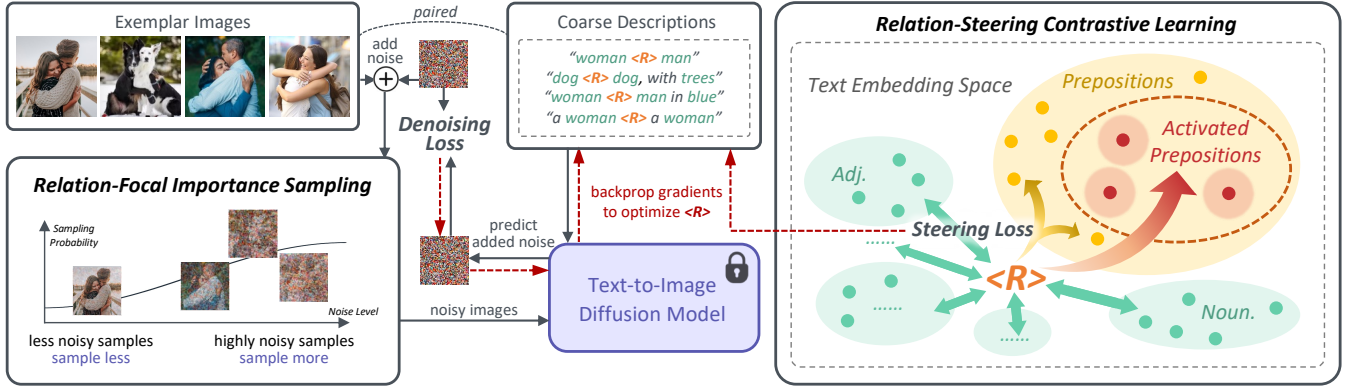


Fig. 3. **ReVersion Framework.** Given exemplar images and their entities’ coarse descriptions, our ReVersion framework optimizes the relation prompt ($\langle R \rangle$) to capture the relation that co-exists in all the exemplar images. During optimization, the *relation-focal importance sampling* strategy encourages $\langle R \rangle$ to focus on high-level relations, and the *relation-steering contrastive learning* scheme induces the relation prompt ($\langle R \rangle$) towards relation-dense regions and away from entities or appearances. Upon optimization, $\langle R \rangle$ can be used as a word in new sentences to make novel entities interact via the relation in exemplar images.

4.2 Relation-Steering Contrastive Learning

Recall that our goal is to acquire a relation prompt $\langle R \rangle$ that accurately captures the co-existing relation in the exemplar images. A basic objective is to reconstruct the exemplar images using $\langle R \rangle$:

$$\langle R \rangle = \arg \min_{\langle r \rangle} \mathbb{E}_{t, x_0, \epsilon} [\|\epsilon - \epsilon_{\theta}(x_t, t, \tau_{\theta}(c))\|^2], c \text{ contains } \langle r \rangle \quad (3)$$

where $\epsilon \sim \mathcal{N}(0, I)$, $\langle R \rangle$ is the optimized text embedding, and $\epsilon_{\theta}(\cdot)$ is a pre-trained text-to-image diffusion model whose weights are frozen throughout optimization. $\langle r \rangle$ is the relation prompt being optimized, and is fed into the pre-trained T2I model as part of the text description c .

However, this pixel-level reconstruction loss mainly focus on reconstructing visual details, without emphasis on object relations. Consequently, we find that directly optimizing with Equation 3 could lead the relation prompt $\langle R \rangle$ to be more associated with the look of objects rather than the relation between them, undesirably leaking entity appearance from exemplar images into the generated images, and also causing wrong object relations.

To mitigate this problem, we propose the “*relation-steering contrastive learning*” scheme, leveraging linguistic priors discussed in Section 1 to emphasis more on object relation during the optimization of $\langle R \rangle$. Specifically, we sample a set of prepositions as positives and use other Part-of-Speech (POS)’ words (e.g., nouns, adjectives) in the text descriptions as negatives to steer the relation prompt towards a relation-dense text embedding subspace, and push it away from appearance-related semantics. Following InfoNCE [Miech et al. 2020; Oord et al. 2018], we formulate the Steering Loss by:

$$L_{\text{steer}} = -\log \frac{\sum_{l=1}^L e^{\langle r \rangle^{\top} \cdot P_l^+ / \gamma}}{\sum_{l=1}^L e^{\langle r \rangle^{\top} \cdot P_l^+ / \gamma} + \sum_{m=1}^M e^{\langle r \rangle^{\top} \cdot N_m^+ / \gamma}}, \quad (4)$$

where $\langle r \rangle$ is the relation embedding, and γ is the temperature parameter. $P_l = \{P_l^1, \dots, P_l^L\}$ (i.e., positive samples) refers to a set of a randomly sampled preposition embeddings from basis prepositions

(more details provided in Supplementary File) at the i -th optimization iteration, and $N_i = \{N_i^1, \dots, N_i^M\}$ (i.e., negative samples) are the embeddings of all other POS’ words (e.g., nouns, adjectives) in the exemplars’ text descriptions in the current batch. All embeddings are normalized to unit length. We find that our relation-steering contrastive learning scheme can effectively help $\langle r \rangle$ to focus on relation and mitigate the appearance leakage problem (see Figure 7 and Section 6.5).

4.3 Relation-Focal Importance Sampling

In the sampling process of diffusion models, high-level semantics usually appear first, and fine details emerge at later stages [Huang et al. 2023; Liew et al. 2022; Patashnik et al. 2023; Wang and Vas-tola 2023]. As our objective is to capture the relation (a high-level concept) in exemplar images, it is undesirable to focus too much on fine-grained visual details (e.g., color, texture) during optimization. Therefore, we further conduct an importance sampling strategy to encourage the learning of high-level relations. Specifically, unlike previous reconstruction objectives, which samples the timestep t from a uniform distribution, we skew the sampling distribution so that a higher probability is assigned to a larger t . The Denoising Loss for “*relation-focal importance sampling*” becomes:

$$L_{\text{denoise}} = \mathbb{E}_{t \sim f(t), x_0, \epsilon} [\|\epsilon - \epsilon_{\theta}(x_t, t, \tau_{\theta}(c))\|^2], \quad (5)$$

$$f(t) = \frac{1}{T} (1 - \alpha \cos \frac{\pi t}{T}),$$

where $f(t)$ is the importance sampling function, which characterizes the probability density function to sample t from. The skewness of $f(t)$ increases with $\alpha \in (0, 1]$. We set $\alpha = 0.5$ throughout our experiments. The overall optimization objective of the **ReVersion Framework** is:

$$\langle R \rangle = \arg \min_{\langle r \rangle} (\lambda_{\text{steer}} L_{\text{steer}} + \lambda_{\text{denoise}} L_{\text{denoise}}), \quad (6)$$

where λ_{steer} and λ_{denoise} are the weighting factors.

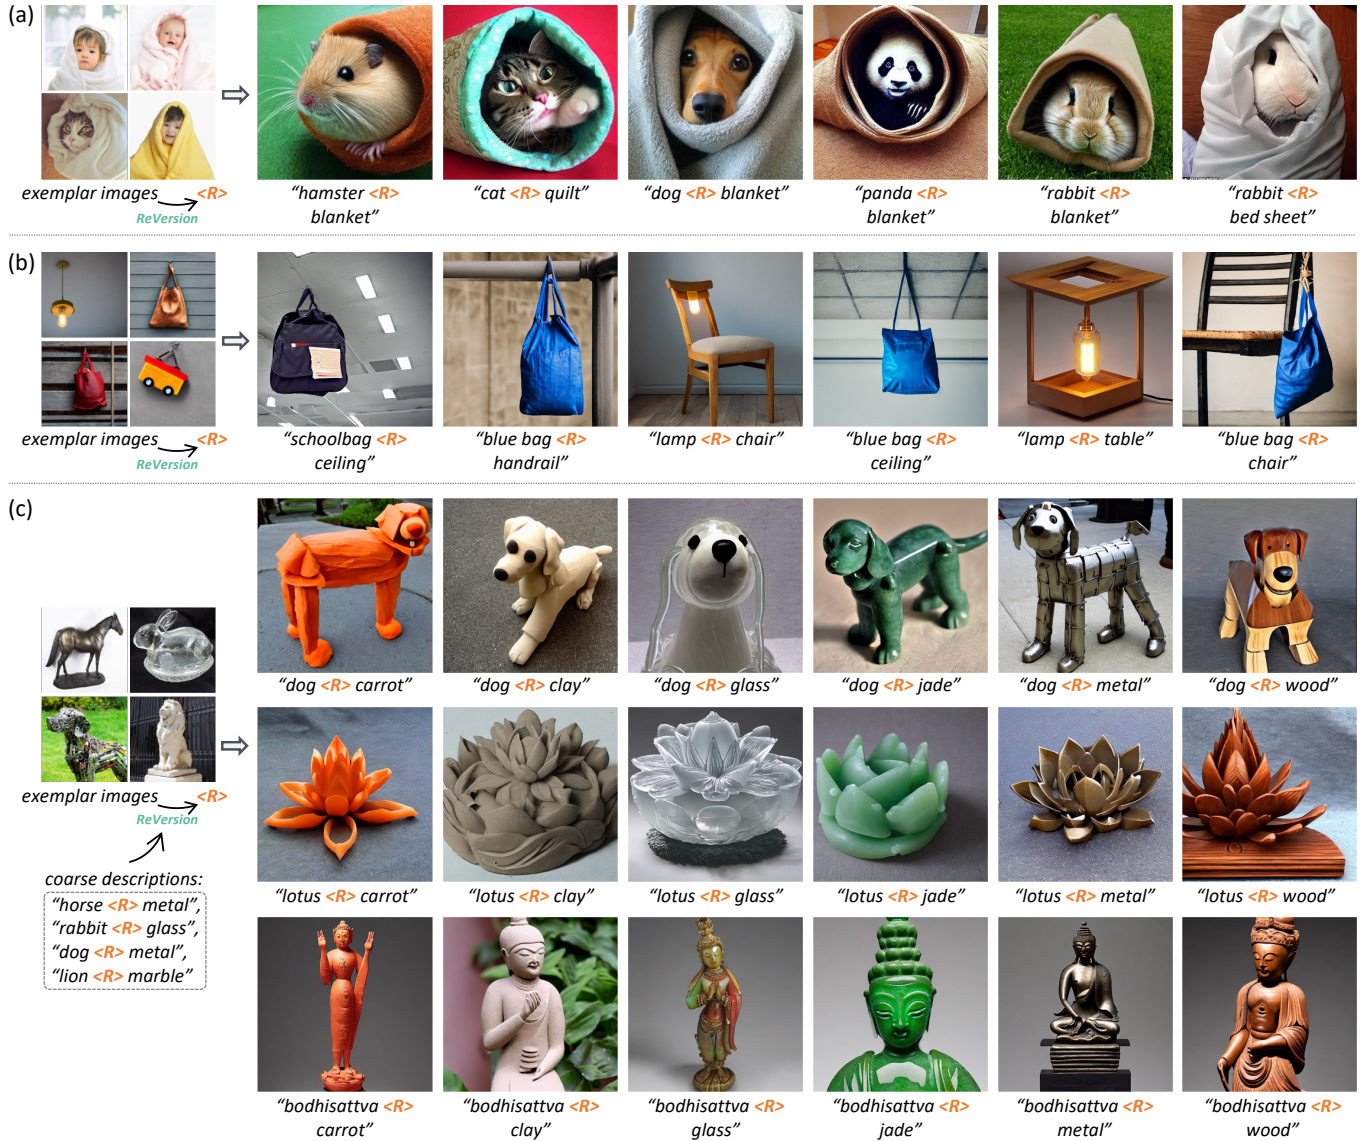


Fig. 4. **Qualitative Results.** Our ReVersion Framework successfully captures the relation that co-exists in the exemplar images, and applies the extracted relation prompt $\langle R \rangle$ to compose novel entities.

5 THE REVERSION BENCHMARK

To facilitate fair comparison for *Relation Inversion*, we present the **ReVersion Benchmark**. It consists of diverse *relations* and *entities*, along with a set of well-defined text *descriptions*. This benchmark can be used for conducting qualitative and quantitative evaluations. Additional details are in Supplementary File.

Relations and Entities. We define ten representative object relations with different abstraction levels, ranging from basic spatial relations (e.g., “on top of”), entity interactions (e.g., “shakes hands with”), to abstract concepts (e.g., “is carved by”). A wide range of entities, such as animals, human, household items, are involved to further increase the diversity of the benchmark.

Exemplar Images and Text Descriptions. For each relation, we collect four to ten exemplar images containing different entities. We further annotate several text templates for each exemplar image to describe them with different levels of details¹. These training templates can be used for the optimization of the relation prompt. **Benchmark Scenarios.** To validate the robustness of the *Relation Inversion* methods, we design 100 inference templates composing of different object entities for each of the ten relations. This provides a total of 1,000 inference templates for performance evaluation.

¹For example, a photo of a cat sitting on a box could be annotated as 1) “cat $\langle R \rangle$ box”, 2) “an orange cat $\langle R \rangle$ a black box” and 3) “an orange cat $\langle R \rangle$ a black box, with trees in the background”. Detailed examples will be in the Supplementary File.

Table 1. **Comparisons via Objective Metrics.** We compare our performance against existing methods and ablation variants using objective evaluation metrics.

(a) **Baseline Comparison.** Performance against several existing methods.

Method	Relation Score \uparrow	Entity Score \uparrow
Text-to-Image [Rombach et al. 2022]	0.3516	0.2896
Textual Inversion [Gal et al. 2022]	0.3785	0.2679
DreamBooth [Ruiz et al. 2022]	0.3576	0.2902
Ours	0.3817	0.2820

(b) **Ablation Study.** Steering or importance sampling is removed.

Method	Relation Score \uparrow	Entity Score \uparrow
Ours w/o Relation-Steering	0.3748	0.2766
Ours w/o Importance Sampling	0.3464	0.2790
Ours	0.3817	0.2820

6 EXPERIMENTS

We present qualitative and quantitative results in this section, and more experiments and analysis are in the Supplementary File. We adopt Stable Diffusion [Rombach et al. 2022] for all experiments since it achieves a good balance between quality and speed. We generate images at 512×512 resolution.

6.1 Comparison Methods

Text-to-Image Generation using Stable Diffusion [Rombach et al. 2022]. We use the original Stable Diffusion 1.5 as the text-to-image generation baseline. Since there is no ground-truth textual description for the relation in each set of exemplar images, we use natural language that can best describe the relation to replace the $\langle R \rangle$ token. For example, in Figure 5 (a), the co-existing relation in the reference images can be roughly described as “*is painted on*”. Thus, we use it to replace the $\langle R \rangle$ token in the inference template “*Spiderman $\langle R \rangle$ building*”, resulting in a sentence “*Spiderman is painted on building*”, which is then used as the text prompt for generation.

Textual Inversion [Gal et al. 2022]. For fair comparison with our method developed on Stable Diffusion 1.5, we use the *diffusers* [Face [n. d.]] implementation of Textual Inversion [Gal et al. 2022] on Stable Diffusion 1.5. Based on the default hyper-parameter settings, we tuned the learning rate and batch size for its optimal performance on our *Relation Inversion* task. We use Textual Inversion’s LDM objective to optimize $\langle R \rangle$ for 3000 iterations, and generate images using the obtained $\langle R \rangle$.

DreamBooth [Ruiz et al. 2022]. We use *diffusers* implementation of DreamBooth on Stable Diffusion 1.5. To adapt DreamBooth to our *Relation Inversion* task for fair comparison, we made three modifications to the original implementation. First, instead of using the original training template like “A photo of V^* dog”, we explicitly inject the word “relation” into the text template to help DreamBooth focus on relation instead of entity, thereby using “A photo of $\langle R \rangle$ relation” to fine-tune the model. Second, the class-specific prior preservation loss is implemented with a text prompt “A photo of *relation*” to avoid overfitting or language drift. Third, to align with fine-tuning stage’s template, the template “Entity A is in $\langle R \rangle$ relation with Entity B” is used during inference.

Table 2. **Comparison with Existing Methods (Human Preference).** Percentage of votes where users favor our results vs. comparison methods. Our method outperforms the baselines under all metrics.

Method	Relation Accuracy	Entity Accuracy	Overall Quality
Text-to-Image Generation [Rombach et al. 2022]	6.45%	10.32%	9.68%
Textual Inversion [Gal et al. 2022]	6.13%	5.81%	5.16%
DreamBooth [Ruiz et al. 2022]	18.39%	18.39%	19.03%
Ours	69.03%	65.48%	66.13%

Table 3. **Ablation Study (Human Preference).** Suppressing relation-steering or importance sampling introduces performance drops, which shows the necessity of both relation-steering and importance sampling.

Method	Relation Accuracy	Entity Accuracy	Overall Quality
w/o Relation-Steering	11.20%	10.90%	13.31%
w/o Importance Sampling	11.20%	13.62%	7.14%
Ours	77.60%	75.48%	79.55%

6.2 Qualitative Comparisons

Our Results. In Figure 4, we provide the generation results using $\langle R \rangle$ inverted by ReVersion. We observe that our framework is capable of 1) synthesizing the entities in the inference template and 2) ensuring that entities follow the relation co-existing in the exemplar images. We provide additional qualitative results in the Supplementary File due to space constraint.

Comparison of Relation Accuracy. Figure 5 shows qualitative comparisons with existing methods. We compare our method with 1) Text-to-Image Generation via Stable Diffusion [Rombach et al. 2022], 2) Textual Inversion [Gal et al. 2022], and 3) DreamBooth [Ruiz et al. 2022]. In Figure 5 (a), although “Text-to-Image Generation” and “DreamBooth” successfully generate both entities (Spiderman and building), they fail to *paint* Spiderman on the building as the exemplar images do. They severely rely on the bias between two entities: Spiderman usually *climbs/jumps* on the buildings, instead of being *painted* onto the buildings. Similarly, in Figure 5 (b), although all methods in comparison can generate at least one monkey, the relation between generated monkeys does not follow the “*back to back*” relation in the exemplar images. In contrast, Our ReVersion Framework does not have this problem.

Entity Leakage in Existing Methods. In Textual Inversion, entities in the exemplar images like canvas are leaked to $\langle R \rangle$, such that the generated image shows a Spiderman on the canvas even when the word “*canvas*” is not in the inference prompt (see Figure 5 (a)). In DreamBooth, the “*basket*” in exemplar images sometimes leak to the generated images (see Figure 9). In Figure 6, we include comparisons with NeTI [Alaluf et al. 2023] and also discuss its entity leakage problem.

6.3 Quantitative Comparisons via Human Evaluation

We conduct user studies with 68 human evaluators to assess the performance of our ReVersion Framework on the *Relation Inversion* task. We sampled 20 groups of images. Each group has images generated by different methods or ablation variants. For each group, apart from the generated images, the following information is presented: 1) exemplar images of a particular relation, 2) text description of the exemplar images. We then ask the evaluators to vote for the best generated image with respect to the following metrics.

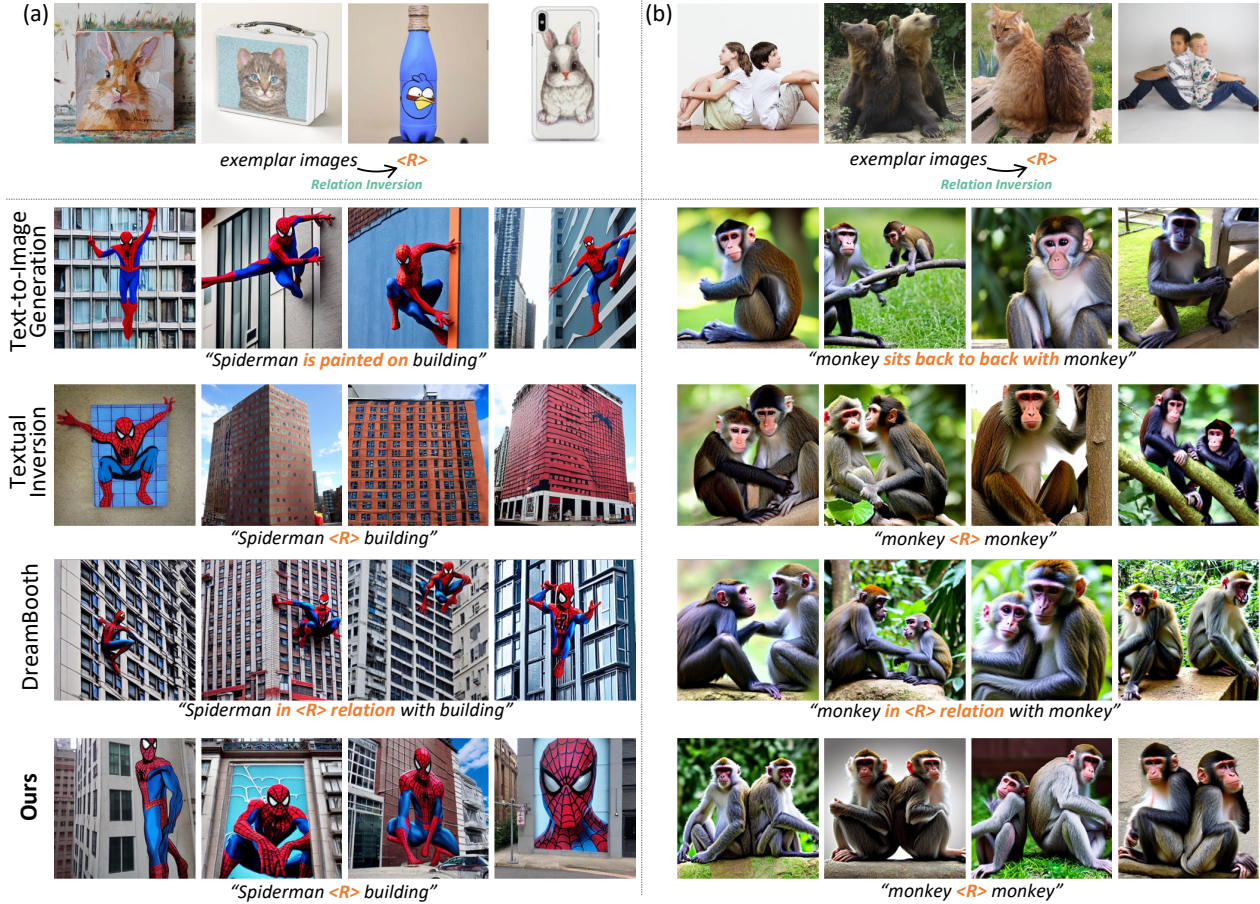


Fig. 5. **Qualitative Comparisons with Existing Methods.** Our method can generate entity and relation accurately. “Text-to-Image Generation” and “DreamBooth” can correctly generate entities described in text prompt, but fail to compose them following the desired relation. “Textual Inversion” suffers from appearance leakage (e.g., $\langle R \rangle$ unexpectedly capturing the canvas in exemplar images), thus resulting in low entity accuracy (e.g., cannot generate spiderman and building simultaneously).

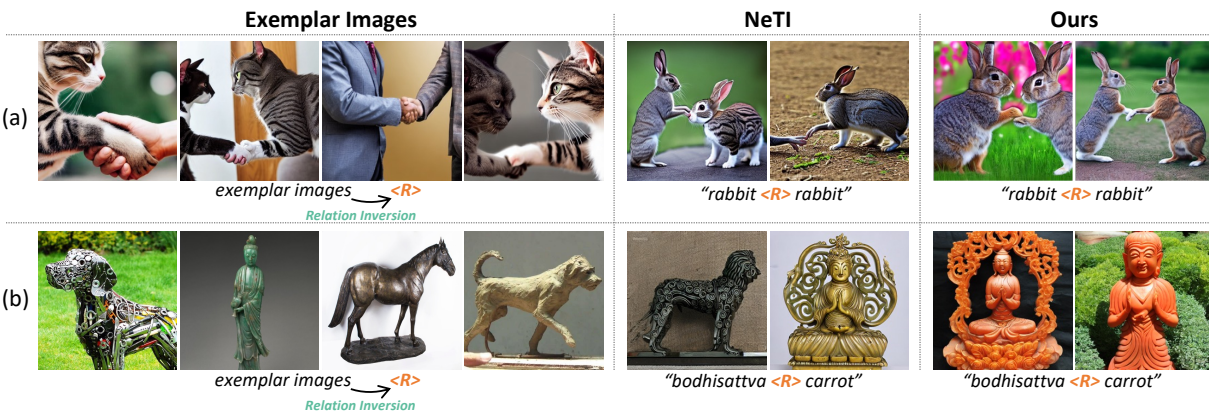


Fig. 6. **Comparisons with Newer Method.** NeTI [Alaluf et al. 2023] demonstrates *some degree of effectiveness* for relation inversion, attributed to its adaptive adjustment at different network layers and denoising timesteps. For example, in (a) where $\langle R \rangle$ denotes “shaking hands”, NeTI successfully rendered rabbits extending their hands, trying to engage in the “shake hands” behaviour. However, NeTI is still prone to *texture leakage*. For instance: (a) The striped patterns of cat fur from the exemplar images are unintentionally transferred to the rabbit fur in NeTI’s outputs. (b) With the “carved by” relation, the metal dog appearance in the exemplar images is unintentionally captured by NeTI, resulting in images resembling a metal animal even when the text prompt is “bodhisattva $\langle R \rangle$ carrot”. Our relation steering is essential to help $\langle R \rangle$ focus on the relation rather than the appearance, thereby producing results without texture leakage.



Fig. 7. **Ablation Study (Qualitative)**. Without relation-steering, $\langle R \rangle$ suffers from appearance leak (e.g., white puppy in (a), gray background in (b)) and inaccurate relation capture (e.g., dog not being on top of plate in (b)). Without importance sampling, $\langle R \rangle$ focuses on lower-level visual details (e.g., rattan around puppy in (a)) and misses high-level relations.

Relation Accuracy. Human evaluators are asked to evaluate whether the relations of the two entities in the generated image are consistent with the relation co-existing in the exemplar images.

Entity Accuracy. Given an inference template in the form of “ $\langle \text{Entity A} \rangle \langle R \rangle \langle \text{Entity B} \rangle$ ”, we ask evaluators to determine whether $\langle \text{Entity A} \rangle$ and $\langle \text{Entity B} \rangle$ are both authentically generated in each image.

Overall Quality. Human evaluators are asked to assess the overall performance on the ReVersion task, considering both the alignment of relation and entity, and the image quality.

Table 2 shows our method clearly obtains better results under all three metrics.

6.4 Quantitative Comparisons via Objective Metrics

We devise automatic metrics to objectively evaluate “relation accuracy” and “entity accuracy”, which are briefly introduced below. More implementation details of the objective metrics will be detailed in the Supplementary File. For comparison experiments, we use the 1,000 inference templates in the ReVersion Benchmark for all relations, and generate 10 images using each template.

Relation Score. We use PSGFormer [Yang et al. 2022], a pre-trained scene-graph generation network, to extract the relation features

for relation accuracy evaluation. Table 1a shows that our method outperforms all existing methods in comparison.

Entity Score. We use CLIP [Radford et al. 2021] score to calculate the alignment between the entity types in the text prompt versus the generated entities. Table 1a shows that our method outperforms Textual Inversion in terms of entity accuracy. This is because the $\langle R \rangle$ learned by Textual Inversion contains leaked entity information, which distracts the model from generating the desired “ E_A ” and “ E_B ”. Our steering loss effectively prevents entity information from leaking into $\langle R \rangle$, allowing for accurate entity synthesis. Furthermore, our approach achieves comparable entity score with “Text-to-Image Generation” and “DreamBooth”, and significantly surpasses them in terms of relation score. It is worth mentioning that the CLIP-based metrics mainly focus on whether the correct class of object is generated, and does not fully take the pixel-level object quality into account. For example, as shown in Figure 9, the stripe textures of cat fur in exemplar images often leak to $\langle R \rangle$, resulting in unrealistic textures in generated rabbits.

6.5 Ablation Study

From both Table 3 (human evaluation) and Table 1b (objective metrics), we observe that removing steering or importance sampling results in deterioration in both relation accuracy and entity accuracy.



Fig. 8. **ReVersion for Complicated Relation.** (a) **Exemplar images.** In each exemplar image, people exhibit the similar relation of “holding hands, leaning backwards”. (b) **Ours.** ReVersion effectively captures this relation by $\langle R \rangle$ and successfully applies it to new entities. (c) **Describe and T2I.** The “first describe the relation, then use text-to-image” approach struggles to accurately represent such complex relation in newly synthesized images.

This corroborates our observations that 1) relation-steering effectively guides $\langle R \rangle$ towards the relation-dense regions and disentangles $\langle R \rangle$ away from exemplar entities, and 2) importance sampling emphasizes high-level relations over low-level details, aiding $\langle R \rangle$ to be relation-focal. We further show qualitatively the necessity of both modules in Figure 7.

Effectiveness of Relation-Steering. In “w/o Relation-Steering”, we remove the Steering Loss L_{steer} in the optimization process. As shown in Figure 7 (a), the appearance of the white puppy in the lower-left exemplar image is leaked into $\langle R \rangle$, resulting in similar puppies in the generated images. In Figure 7 (b), many appearance elements are leaked into $\langle R \rangle$, such as the gray background, the black cube, and the husky dog. The dog and the plate also do not follow the relation of “being on top of” as shown in exemplar images. Consequently, the images generated via $\langle R \rangle$ do not present the correct relation and introduced unwanted leaked imagery.

Effectiveness of Importance Sampling. We replace our relation-focal importance sampling with uniform sampling, and observe that $\langle R \rangle$ pays too much attention to low-level details rather than high-level relations. For instance, in Figure 7 (a) “w/o Importance Sampling”, the basket rattan wraps around puppy’s head in the same way as the exemplar image, instead of containing the puppy inside.

6.6 Further Analysis

Diverse Styles and Backgrounds. As shown in Figure 10, the $\langle R \rangle$ inverted by ReVersion can be applied robustly to relate entities in scenes with diverse backgrounds or styles.

More Comparisons on Complicated Relation. Some relations are hard to accurately express by text, or the description of such



Fig. 9. **Appearance Leakage of DreamBooth.** (a) **Entity Leakage (Red Boxes):** The basket from the exemplar images significantly leaks into images generated by DreamBooth. In contrast, our approach avoids this issue of entity leakage. (b) **Texture Leakage (Green Boxes):** While DreamBooth accurately generates the entity “rabbit”, it encounters texture leakage from the exemplar images. That is, stripe patterns of cat fur texture (marked with green boxes) unintentionally transfer to the rabbit’s fur in DreamBooth’s outputs. Our method, in contrast, is free from such texture leakage.

relation may be complex and difficult for the text-to-image generation model to effectively comprehend. For the relation shown in Figure 8 (a), our method (Figure 8 (b)) effectively captures these relations using $\langle R \rangle$ and applies them to new entities. In Figure 8 (c), we engage four human subjects to observe the exemplar images in (a) and describe scenes where these relations are applied to new entities (detailed process in Supplementary File). Subsequently, we utilize text-to-image (T2I) to synthesize images based on these human descriptions. The results demonstrate that this “describe and T2I” approach struggles to accurately represent such complex relations in the newly synthesized images.

6.7 Limitations and Potential Societal Impacts

Limitations. Our performance is dependent on the generative capabilities of Stable Diffusion. For instance, it might produce sub-optimal synthesis results for entities that Stable Diffusion struggles at, such as human body and human face. We discuss limitations of “human synthesis” and “concept blending” in detail in the Supplementary File with qualitative examples.

Potential Negative Societal Impacts. The entity relational composition capabilities of ReVersion could be applied maliciously on real human figures. Additional potential impacts are discussed in the Supplementary File in depth.



Fig. 10. **ReVersion for Diverse Styles and Backgrounds.** The $\langle R \rangle$ inverted by ReVersion can be applied robustly to relate entities under diverse backgrounds or styles.

7 CONCLUSION

In this work, we take the first step forward and propose the **Relation Inversion** task, which aims to learn a relation prompt to capture the relation that co-exists in multiple exemplar images. In our **ReVersion Framework**, we use *relation-steering contrastive learning* scheme to effectively guide the relation prompt towards relation-dense regions in the text embedding space, and our *relation-focal importance sampling* scheme shift the focus from visual details to high-level relations. We also contribute the **ReVersion Benchmark** for performance evaluation. Our proposed **Relation Inversion** task would be a good inspiration for future works in various domains

such as generative model inversion, representation learning, few-shot learning, visual relation detection, and scene graph generation.

ACKNOWLEDGMENTS

This study is supported by the Ministry of Education, Singapore, under its MOE AcRF Tier 2 (MOET2EP20221- 0012), NTU NAP, and under the RIE2020 Industry Alignment Fund – Industry Collaboration Projects (IAF-ICP) Funding Initiative, as well as cash and in-kind contribution from the industry partner(s).

REFERENCES

- Yuval Alaluf, Elad Richardson, Gal Metzger, and Daniel Cohen-Or. 2023. A neural space-time representation for text-to-image personalization. *ACM TOG* 42, 6 (2023), 1–10.
- Tomer Amit, Eliya Nachmani, Tal Shaharabany, and Lior Wolf. 2021. SegDiff: Image Segmentation with Diffusion Probabilistic Models. *arXiv preprint arXiv:2112.00390* (2021).
- Moab Arar, Rinon Gal, Yuval Atzmon, Gal Chechik, Daniel Cohen-Or, Ariel Shamir, and Amit H Bermano. 2023. Domain-agnostic tuning-encoder for fast personalization of text-to-image models. *arXiv preprint arXiv:2307.06925* (2023).
- Jacob Austin, Daniel D Johnson, Jonathan Ho, Daniel Tarlow, and Rianne van den Berg. 2021. Structured denoising diffusion models in discrete state-spaces. In *NeurIPS*.
- Dmitry Baranchuk, Ivan Rubachev, Andrey Voynov, Valentin Khulkov, and Artem Babenko. 2022. Label-efficient semantic segmentation with diffusion models. In *ICLR*.
- Andreas Blattmann, Robin Rombach, Huan Ling, Tim Dockhorn, Seung Wook Kim, Sanja Fidler, and Karsten Kreis. 2023. Align your Latents: High-Resolution Video Synthesis with Latent Diffusion Models. In *CVPR*.
- Hong Chen, Yipeng Zhang, Xin Wang, Xuguang Duan, Yuwei Zhou, and Wenwu Zhu. 2023b. DisenBooth: Disentangled Parameter-Efficient Tuning for Subject-Driven Text-to-Image Generation. *arXiv preprint arXiv:2305.03374* (2023).
- Wenhu Chen, Hexiang Hu, Yandong Li, Nataniel Ruiz, Xuhui Jia, Ming-Wei Chang, and William W Cohen. 2023a. Subject-driven Text-to-Image Generation via Apprenticeship Learning. *arXiv preprint arXiv:2304.00186* (2023).
- Jooyoung Choi, Yunjeong Choi, Yunji Kim, Junho Kim, and Sungroh Yoon. 2023. Custom-Edit: Text-Guided Image Editing with Customized Diffusion Models. *arXiv preprint arXiv:2305.15779* (2023).
- Prafulla Dhariwal and Alexander Nichol. 2021. Diffusion Models Beat GANs on Image Synthesis. In *NeurIPS*.
- Patrick Esser, Robin Rombach, Andreas Blattmann, and Bjorn Ommer. 2021. Image-BART: Bidirectional context with multinomial diffusion for autoregressive image synthesis. In *NeurIPS*.
- Patrick Esser, Robin Rombach, and Björn Ommer. 2020. A note on data biases in generative models. In *NeurIPS Workshop*.
- Hugging Face. [n. d.]. *Diffusers*.
- Rinon Gal, Yuval Alaluf, Yuval Atzmon, Or Patashnik, Amit H. Bermano, Gal Chechik, and Daniel Cohen-Or. 2022. An Inverse is Worth One Word: Personalizing Text-to-Image Generation using Textual Inversion. *arXiv preprint arXiv:2208.01618* (2022).
- Rinon Gal, Moab Arar, Yuval Atzmon, Amit H Bermano, Gal Chechik, and Daniel Cohen-Or. 2023. Encoder-based domain tuning for fast personalization of text-to-image models. *ACM TOG* (2023).
- Yuan Gong, Youxin Pang, Xiaodong Cun, Menghan Xia, Haoxin Chen, Longyue Wang, Yong Zhang, Xintao Wang, Ying Shan, and Yujiu Yang. 2023. TaleCrafter: Interactive Story Visualization with Multiple Characters. *arXiv preprint arXiv:2305.18247* (2023).
- Ian J Goodfellow, Jean Pouget-Abadie, Mehdi Mirza, Bing Xu, David Warde-Farley, Sherjil Ozair, Aaron C Courville, and Yoshua Bengio. 2014. Generative Adversarial Nets. In *NeurIPS*.
- Alexandros Graikos, Nikolay Malkin, Nebojsa Jojic, and Dimitris Samaras. 2022. Diffusion Models as Plug-and-Play Priors. In *NeurIPS*.
- Shuyang Gu, Dong Chen, Jianmin Bao, Fang Wen, Bo Zhang, Dongdong Chen, Lu Yuan, and Baining Guo. 2022. Vector quantized diffusion model for text-to-image synthesis. In *CVPR*.
- Ligong Han, Yinxiao Li, Han Zhang, Peyman Milanfar, Dimitris Metaxas, and Feng Yang. 2023. SVDiff: Compact Parameter Space for Diffusion Fine-Tuning. *arXiv preprint arXiv:2303.11305* (2023).
- William Harvey, Saeid Naderiparizi, Vaden Masrani, Christian Weillbach, and Frank Wood. 2022. Flexible Diffusion Modeling of Long Videos. *arXiv preprint arXiv:2205.11495* (2022).
- Kaiming He, Haoqi Fan, Yuxin Wu, Saining Xie, and Ross Girshick. 2020. Momentum contrast for unsupervised visual representation learning. In *CVPR*.
- Yingqing He, Tianyu Yang, Yong Zhang, Ying Shan, and Qifeng Chen. 2022. Latent video diffusion models for high-fidelity video generation with arbitrary lengths. *arXiv preprint arXiv:2211.13221* (2022).
- Jonathan Ho, Ajay Jain, and Pieter Abbeel. 2020. Denoising diffusion probabilistic models. In *NeurIPS*.
- Jonathan Ho, Chitwan Saharia, William Chan, David J Fleet, Mohammad Norouzi, and Tim Salimans. 2022a. Cascaded Diffusion Models for High Fidelity Image Generation. *JMLR* (2022).
- Jonathan Ho, Tim Salimans, Alexey Gritsenko, William Chan, Mohammad Norouzi, and David J Fleet. 2022b. Video diffusion models. *arXiv preprint arXiv:2204.03458* (2022).
- Edward J Hu, Yelong Shen, Phillip Wallis, Zeyuan Allen-Zhu, Yuanzhi Li, Shean Wang, Lu Wang, and Weizhu Chen. 2022. LoRA: Low-Rank Adaptation of Large Language Models. In *ICLR*.
- Ziqi Huang, Kelvin C.K. Chan, Yuming Jiang, and Ziwei Liu. 2023. Collaborative Diffusion for Multi-Modal Face Generation and Editing. In *CVPR*.
- Aapo Hyvärinen and Peter Dayan. 2005. Estimation of non-normalized statistical models by score matching. *JMLR* (2005).
- Jingwei Ji, Ranjay Krishna, Fei-Fei Li, and Juan Carlos Niebles. 2020. Action genome: Actions as compositions of spatio-temporal scene graphs. In *CVPR*. 10236–10247.
- Xuhui Jia, Yang Zhao, Kelvin C.K. Chan, Yandong Li, Han Zhang, Boqing Gong, Tingbo Hou, Huisheng Wang, and Yu-Chuan Su. 2023. Taming Encoder for Zero Fine-tuning Image Customization with Text-to-Image Diffusion. (2023).
- Yuming Jiang, Shuai Yang, Haonan Qiu, Wayne Wu, Chen Change Loy, and Ziwei Liu. 2022. Text2human: Text-driven controllable human image generation. *ACM TOG* (2022).
- Bahjat Kawar, Shiran Zada, Oran Lang, Omer Tov, Huiwen Chang, Tali Dekel, Inbar Mosseri, and Michal Irani. 2022. Imagic: Text-based real image editing with diffusion models. *arXiv preprint arXiv:2210.09276* (2022).
- Ranjay Krishna, Yuke Zhu, Oliver Groth, Justin Johnson, Kenji Hata, Joshua Kravitz, Stephanie Chen, Yannis Kalantidis, Li-Jia Li, David A Shamma, Michael Bernstein, and Fei-Fei Li. 2017. Visual Genome: Connecting Language and Vision Using Crowdsourced Dense Image Annotations. *IJCV* (2017).
- Nupur Kumari, Bingliang Zhang, Richard Zhang, Eli Shechtman, and Jun-Yan Zhu. 2022. Multi-Concept Customization of Text-to-Image Diffusion. *arXiv preprint arXiv:2212.04488* (2022).
- Dongxu Li, Junnan Li, and Steven CH Hoi. 2023a. BLIP-Diffusion: Pre-trained Subject Representation for Controllable Text-to-Image Generation and Editing. *arXiv preprint arXiv:2305.14720* (2023).
- Yuheng Li, Haotian Liu, Yangming Wen, and Yong Jae Lee. 2023b. Generate Anything Anywhere in Any Scene. *arXiv preprint arXiv:2306.17154* (2023).
- Jun Hao Liew, Hanshu Yan, Daquan Zhou, and Jiashi Feng. 2022. MagicMix: Semantic Mixing with Diffusion Models. *arXiv preprint arXiv:2210.16056* (2022).
- Ilya Loshchilov and Frank Hutter. 2019. Decoupled weight decay regularization. In *ICLR*.
- Cewu Lu, Ranjay Krishna, Michael Bernstein, and Fei-Fei Li. 2016. Visual Relationship Detection with Language Priors. In *ECCV*.
- Jian Ma, Junhao Liang, Chen Chen, and Haonan Lu. 2023. Subject-diffusion: Open domain personalized text-to-image generation without test-time fine-tuning. *arXiv preprint arXiv:2307.11410* (2023).
- Chenlin Meng, Yutong He, Yang Song, Jiaming Song, Jiajun Wu, Jun-Yan Zhu, and Stefano Ermon. 2022. SDEdit: Guided image synthesis and editing with stochastic differential equations. In *ICLR*.
- Antoine Miech, Jean-Baptiste Alayrac, Lucas Smaira, Ivan Laptev, Josef Sivic, and Andrew Zisserman. 2020. End-to-end learning of visual representations from uncurated instructional videos. In *CVPR*. 9879–9889.
- Alex Nichol, Prafulla Dhariwal, Aditya Ramesh, Pranav Shyam, Pamela Mishkin, Bob McGrew, Ilya Sutskever, and Mark Chen. 2021. GLIDE: Towards photorealistic image generation and editing with text-guided diffusion models. *arXiv preprint arXiv:2112.10741* (2021).
- Aaron van den Oord, Yazhe Li, and Oriol Vinyals. 2018. Representation learning with contrastive predictive coding. *arXiv preprint arXiv:1807.03748* (2018).
- Or Patashnik, Daniel Garibi, Idan Azuri, Hadar Averbuch-Elor, and Daniel Cohen-Or. 2023. Localizing object-level shape variations with text-to-image diffusion models. In *ICCV*.
- Alec Radford, Jong Wook Kim, Chris Hallacy, Aditya Ramesh, Gabriel Goh, Sandhini Agarwal, Girish Sastry, Amanda Askell, Pamela Mishkin, Jack Clark, et al. 2021. Learning transferable visual models from natural language supervision. In *ICML*.
- Aditya Ramesh, Prafulla Dhariwal, Alex Nichol, Casey Chu, and Mark Chen. 2022. Hierarchical text-conditional image generation with CLIP latents. *arXiv preprint arXiv:2204.06125* (2022).
- Robin Rombach, Andreas Blattmann, Dominik Lorenz, Patrick Esser, and Björn Ommer. 2022. High-resolution image synthesis with latent diffusion models. In *CVPR*.
- Olaf Ronneberger, Philipp Fischer, and Thomas Brox. 2015. U-Net: Convolutional networks for biomedical image segmentation. In *MICCAI*.
- Nataniel Ruiz, Yuanzhen Li, Varun Jampani, Yael Pritch, Michael Rubinstein, and Kfir Aberman. 2022. DreamBooth: Fine Tuning Text-to-image Diffusion Models for Subject-Driven Generation. *arXiv preprint arXiv:2208.12242* (2022).
- Nataniel Ruiz, Yuanzhen Li, Varun Jampani, Wei Wei, Tingbo Hou, Yael Pritch, Neal Wadhwa, Michael Rubinstein, and Kfir Aberman. 2024. Hyperdreambooth: Hypernetworks for fast personalization of text-to-image models. In *CVPR*.
- Chitwan Saharia, William Chan, Saurabh Saxena, Lala Li, Jay Whang, Emily Denton, Seyed Kamyar Seyed Ghasemipour, Burcu Karagol Ayan, S Sara Mahdavi, Rapha Gontijo Lopes, et al. 2022a. Photorealistic Text-to-Image Diffusion Models with Deep Language Understanding. *arXiv preprint arXiv:2205.11487* (2022).
- Chitwan Saharia, Jonathan Ho, William Chan, Tim Salimans, David J Fleet, and Mohammad Norouzi. 2022b. Image super-resolution via iterative refinement. *IEEE TPAMI* (2022).
- Christoph Schuhmann, Romain Beaumont, Richard Vencu, Cade Gordon, Ross Wightman, Mehdi Cherti, Theo Coombes, Aarush Katta, Clayton Mullis, Mitchell Wortsman, et al. 2022. Laion-5b: An open large-scale dataset for training next generation image-text models. *arXiv preprint arXiv:2210.08402* (2022).

- Xindi Shang, Tongwei Ren, Jingfan Guo, Hanwang Zhang, and Tat-Seng Chua. 2017. Video Visual Relation Detection. In *ACM MM*.
- Uriel Singer, Adam Polyak, Thomas Hayes, Xi Yin, Jie An, Songyang Zhang, Qiyuan Hu, Harry Yang, Oron Ashual, Oran Gafni, et al. 2022. Make-a-video: Text-to-video generation without text-video data. *arXiv preprint arXiv:2209.14792* (2022).
- Jascha Sohl-Dickstein, Eric Weiss, Niru Maheswaranathan, and Surya Ganguli. 2015. Deep unsupervised learning using nonequilibrium thermodynamics. In *ICML*.
- Jiaming Song, Chenlin Meng, and Stefano Ermon. 2021a. Denoising diffusion implicit models. In *ICLR*.
- Yang Song, Jascha Sohl-Dickstein, Diederik P Kingma, Abhishek Kumar, Stefano Ermon, and Ben Poole. 2021b. Score-based generative modeling through stochastic differential equations. In *ICLR*.
- Angus Stevenson. 2010. *Oxford dictionary of English*. Oxford University Press, USA.
- Patrick Tinsley, Adam Czajka, and Patrick Flynn. 2021. This face does not exist... but it might be yours! identity leakage in generative models. In *WACV*.
- Laurens Van der Maaten and Geoffrey Hinton. 2008. Visualizing data using t-SNE. *JMLR* 9, 11 (2008).
- Ruben Villegas, Mohammad Babaeizadeh, Pieter-Jan Kindermans, Hernan Moraldo, Han Zhang, Mohammad Taghi Saffar, Santiago Castro, Julius Kunze, and Dumitru Erhan. 2022. Phenaki: Variable length video generation from open domain textual description. *arXiv preprint arXiv:2210.02399* (2022).
- Pascal Vincent. 2011. A connection between score matching and denoising autoencoders. *Neural Computation* (2011).
- Andrey Voynov, Qinghao Chu, Daniel Cohen-Or, and Kfir Aberman. 2023. p+: Extended textual conditioning in text-to-image generation. *arXiv preprint arXiv:2303.09522* (2023).
- Binxu Wang and John J. Vastola. 2023. Diffusion Models Generate Images Like Painters: an Analytical Theory of Outline First, Details Later. *arXiv preprint arXiv:2303.02490* (2023).
- Yuxiang Wei, Yabo Zhang, Zhilong Ji, Jinfeng Bai, Lei Zhang, and Wangmeng Zuo. 2023. ELITE: Encoding Visual Concepts into Textual Embeddings for Customized Text-to-Image Generation. *arXiv preprint arXiv:2302.13848* (2023).
- Jay Zhangjie Wu, Yixiao Ge, Xintao Wang, Stan Weixian Lei, Yuchao Gu, Wynne Hsu, Ying Shan, Xiaohu Qie, and Mike Zheng Shou. 2022. Tune-A-Video: One-Shot Tuning of Image Diffusion Models for Text-to-Video Generation. *arXiv preprint arXiv:2212.11565* (2022).
- Danfei Xu, Yuke Zhu, Christopher B Choy, and Fei-Fei Li. 2017. Scene graph generation by iterative message passing. In *CVPR*.
- Xingqian Xu, Jiayi Guo, Zhangyang Wang, Gao Huang, Irfan Essa, and Humphrey Shi. 2023. Prompt-Free Diffusion: Taking "Text" out of Text-to-Image Diffusion Models. *arXiv preprint arXiv:2305.16223* (2023).
- Jingkang Yang, Yi Zhe Ang, Zujin Guo, Kaiyang Zhou, Wayne Zhang, and Ziwei Liu. 2022. Panoptic Scene Graph Generation. In *ECCV*. Springer, 178–196.
- Jingkang Yang, Wenxuan Peng, Xiangtai Li, Zujin Guo, Liangyu Chen, Bo Li, Zheng Ma, Kaiyang Zhou, Wayne Zhang, Chen Change Loy, and Ziwei Liu. 2023. Panoptic Video Scene Graph Generation. In *CVPR*.
- Hu Ye, Jun Zhang, Sibao Liu, Xiao Han, and Wei Yang. 2023. Ip-adapter: Text compatible image prompt adapter for text-to-image diffusion models. *arXiv preprint arXiv:2308.06721* (2023).
- Ruichi Yu, Ang Li, Vlad I Morariu, and Larry S Davis. 2017. Visual relationship detection with internal and external linguistic knowledge distillation. In *ICCV*.
- Yufan Zhou, Ruiyi Zhang, Tong Sun, and Jinhui Xu. 2023. Enhancing Detail Preservation for Customized Text-to-Image Generation: A Regularization-Free Approach. *arXiv preprint arXiv:2305.13579* (2023).
- Bohan Zhuang, Lingqiao Liu, Chunhua Shen, and Ian Reid. 2017. Towards context-aware interaction recognition for visual relationship detection. In *ICCV*.

SUPPLEMENTARY

In this *supplementary file*, we provide more experimental details in Section A, and elaborate on the ReVersion Benchmark details in Section B. We then provide further explanations on basis prepositions in Section C. We also discuss our limitations in Section D, and the potential societal impacts of our work in Section E. At the end of the supplementary file, we show various qualitative results of ReVersion in Section F.

A MORE EXPERIMENTAL DETAILS

In this section, we provide more experimental details.

A.1 Implementation Details of ReVersion

We introduce the implementation details of the ReVersion Framework. Our framework is built on top of the *diffusers* [Face [n. d.]] implementation of Stable Diffusion [Rombach et al. 2022] 1.5. All experiments are conducted on 512×512 image resolution. In Equation 4, the temperature parameter γ in the steering loss L_{steer} is set as 0.07, following [He et al. 2020]. In each iteration, 8 positive samples are randomly selected from the basis preposition set (see Table A4). In Equation 6, to ensure that the numerical values $\lambda_{\text{denoise}}L_{\text{denoise}}$ and $\lambda_{\text{steer}}L_{\text{steer}}$ are in comparable order of magnitude, we set $\lambda_{\text{denoise}} = 1.0$ and $\lambda_{\text{steer}} = 0.01$. During the optimization process, we first initialize our relation prompt $\langle R \rangle$ using the word “and”, then optimize the prompt using the AdamW [Loshchilov and Hutter 2019] optimizer for 3,000 steps, with learning rate 2.5×10^{-4} and batch size 2. During the inference process, we use classifier-free guidance for all experiments including the baselines and ablation variants, with a constant guidance weight 7.5.

A.2 Human Evaluation

We introduce the implementation details of the user studies in the main paper’s Section 6.3.

Figure A11 is a screenshot of the user study form we distributed for main paper’s Table 1, namely “Comparing with Existing Methods”. We employ preference voting to differentiate the performance of different methods. To ensure unbiased responses, the order of different methods’ results is randomized. That is, the orders of generated images A, B, C , and D are random and different in each question. For main paper’s Table 1, “Comparison with Existing Methods”, four methods are in comparison, so there are four choices: A, B, C , and D . For main paper’s Table 2, “Ablation Study”, three methods are in comparison, so there are three choices: A, B , and C .

A.3 Objective Evaluation Metrics

We introduce the implementation details of the objective metrics used in the main paper’s Section 6.4.

Relation Score. We devise an objective evaluation metric to measure the quality and accuracy of the inverted relation. To do this, we train relation classifiers that categorize the ten relations in our ReVersion benchmark. We then use these classifiers to determine whether the entities in the generated images follow the specified

relation. We employ PSGFormer [Yang et al. 2022], a pre-trained scene-graph generation network, to extract the relation feature vectors from a given image. The feature vectors are averaged-pooled and fed into linear SVMs for classification. We calculate the *Relation Score* as the percentage of generated images that follow the relation class in the exemplar images.

Entity Score. To evaluate whether the generated image contains the entities specified by the text prompt, we compute the CLIP score [Radford et al. 2021] between a revised text prompt and the generated image, which we refer to as the *Entity Score*. CLIP [Radford et al. 2021] is a vision-language model that has been trained on large-scale datasets. It uses an image encoder and a text encoder to project images and text into a common feature space. The CLIP score is calculated as the cosine similarity between the normalized image and text embeddings. A higher score usually indicates greater consistency between the output image and the text prompt. In our approach, we calculate the CLIP score between the generated image and the revised text prompt “ E_A, E_B ”, which only includes the entity information.

A.4 Implementation of “Describe and Text-to-Image”

In main paper’s Section 6.6 and Figure 6, we compared our method against the “Describe and Text-to-Image (T2I)” approach. We provide detailed process in Figure A12.

B REVERSION BENCHMARK DETAILS

In this section, we provide the details of our ReVersion Benchmark. The full benchmark will be publicly available.

B.1 Relations

To benchmark the Relation Inversion task, we define ten diverse and representative object relations as follows:

- E_A **is painted on (the surface of)** E_B
- E_A **is carved by / is made of the material of** E_B
- E_A **shakes hands with** E_B
- E_A **hugs** E_B
- E_A **sits back to back with** E_B
- E_A **is contained inside** E_B
- E_A **on / is on top of** E_B
- E_A **is hanging from** E_B
- E_A **is wrapped in** E_B
- E_A **rides (on)** E_B

where E_A and E_B are the two entities that follow the specified relation. It is worth mentioning that the relations can be best described by the exemplar images, and the text descriptions provided above are simply approximated summarizations of the true relations.

B.2 Exemplar Images


A wide range of entities, such as animals, human, household items, are involved to further increase the diversity of the benchmark. In Figure A13, we show the *exemplar images* and *text descriptions* for the relation “ E_A sits back to back with E_B ”. The exemplar images

Q3*

Metrics:


- (1) Relation Accuracy: whether the generated relation is consistent with that in the exemplar images
- (2) Entity Accuracy: whether the generated entities are consistent with text prompt, and their realism
- (3) Overall Quality: overall performance, consider both *relation* and *entity*, and *image quality*


Exemplar images:





Text prompt: panda <R> panda

Generated images:

A


B


C


D


	A	B	C	D
Best Relation Accuracy	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>
Best Entity Accuracy	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>
Best Overall Quality	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>

Fig. A11. **Example of Human Evaluation.** This is a screenshot of a user study question distributed to human evaluators. The order of different methods (*i.e.*, *A*, *B*, *C*, and *D*) is randomized. Human evaluators are provided with the exemplar images, text prompt, and generated images. They are asked to vote for the best generated image among *A*, *B*, *C*, and *D*, for the three metrics (*i.e.*, *Relation Accuracy* / *Entity Accuracy* / *Overall Quality*) respectively.

There are several images where the entities (*i.e.*, people) in each image follow a common relation. Now, imagine some other entities (*e.g.* two rabbits) also follow this relation. Use English to best describe the new scenario.

You can say something like "One rabbit <TO BE FILLED BY YOU> one rabbit", or "Two rabbits <TO BE FILLED BY YOU>", depending on how you think can best describe it.



Your answer

Fig. A12. **Human Description of Relation.** This is a screenshot of a user study question distributed to human subjects. The human subjects are asked to observe the exemplar images and identify the co-existing relation in the exemplar images. They are then asked to use natural language to describe the relation. The description will then be used for the "Describe and T2I" baseline.

Table A4. **Basis Preposition Set.** We list the set of 56 basis prepositions.

aboard	astride	in	regarding
about	at	including	round
above	atop	inside	through
across	before	into	throughout
after	behind	near	to
against	below	of	toward
along	beneath	off	towards
alongside	beside	on	under
amid	between	onto	underneath
amidst	beyond	opposite	up
among	by	out	upon
amongst	down	outside	versus
anti	following	over	with
around	from	past	within

contain both human figures and animals to emphasize the invariant “back to back” relation in different scenarios.

B.3 Text Descriptions

As shown in Figure A13, the *text descriptions* for each image contains several levels, from short sentences which only mention the class names, to complex and comprehensive sentences that describe each entity and the scene backgrounds. The $\langle R \rangle$ in each description will be replaced by the learnable relation prompt during optimization.

B.4 Inference Templates

To evaluate the performance of relation inversion methods, we devise 100 inference templates for each relation. The inference templates contains diverse entity combinations to test the robustness and generalizability of the inverted relation $\langle R \rangle$. To quantitatively evaluate relation inversion performance, we use each inference template to synthesize 10 images, resulting in a total of 1,000 synthesized images for each inverted $\langle R \rangle$.

Below, we show the 100 inference templates for the relation “ E_A sits back to back with E_B ”:

- *man* $\langle R \rangle$ *man*, *man* $\langle R \rangle$ *woman*, *man* $\langle R \rangle$ *child*, *man* $\langle R \rangle$ *cat*, *man* $\langle R \rangle$ *rabbit*, *man* $\langle R \rangle$ *monkey*, *man* $\langle R \rangle$ *dog*, *man* $\langle R \rangle$ *hamster*, *man* $\langle R \rangle$ *kangaroo*, *man* $\langle R \rangle$ *panda*,
- *woman* $\langle R \rangle$ *man*, *woman* $\langle R \rangle$ *woman*, *woman* $\langle R \rangle$ *child*, *woman* $\langle R \rangle$ *cat*, *woman* $\langle R \rangle$ *rabbit*, *woman* $\langle R \rangle$ *monkey*, *woman* $\langle R \rangle$ *dog*, *woman* $\langle R \rangle$ *hamster*, *woman* $\langle R \rangle$ *kangaroo*, *woman* $\langle R \rangle$ *panda*,
- *child* $\langle R \rangle$ *man*, *child* $\langle R \rangle$ *woman*, *child* $\langle R \rangle$ *child*, *child* $\langle R \rangle$ *cat*, *child* $\langle R \rangle$ *rabbit*, *child* $\langle R \rangle$ *monkey*, *child* $\langle R \rangle$ *dog*, *child* $\langle R \rangle$ *hamster*, *child* $\langle R \rangle$ *kangaroo*, *child* $\langle R \rangle$ *panda*,
- *cat* $\langle R \rangle$ *man*, *cat* $\langle R \rangle$ *woman*, *cat* $\langle R \rangle$ *child*, *cat* $\langle R \rangle$ *cat*, *cat* $\langle R \rangle$ *rabbit*, *cat* $\langle R \rangle$ *monkey*, *cat* $\langle R \rangle$ *dog*, *cat* $\langle R \rangle$ *hamster*, *cat* $\langle R \rangle$ *kangaroo*, *cat* $\langle R \rangle$ *panda*,
- *rabbit* $\langle R \rangle$ *man*, *rabbit* $\langle R \rangle$ *woman*, *rabbit* $\langle R \rangle$ *child*, *rabbit* $\langle R \rangle$ *cat*, *rabbit* $\langle R \rangle$ *rabbit*, *rabbit* $\langle R \rangle$ *monkey*, *rabbit* $\langle R \rangle$ *dog*, *rabbit* $\langle R \rangle$ *hamster*, *rabbit* $\langle R \rangle$ *kangaroo*, *rabbit* $\langle R \rangle$ *panda*,
- *monkey* $\langle R \rangle$ *man*, *monkey* $\langle R \rangle$ *woman*, *monkey* $\langle R \rangle$ *child*, *monkey* $\langle R \rangle$ *cat*, *monkey* $\langle R \rangle$ *rabbit*, *monkey* $\langle R \rangle$ *monkey*, *monkey*

$\langle R \rangle$ *dog*, *monkey* $\langle R \rangle$ *hamster*, *monkey* $\langle R \rangle$ *kangaroo*, *monkey* $\langle R \rangle$ *panda*,

- *dog* $\langle R \rangle$ *man*, *dog* $\langle R \rangle$ *woman*, *dog* $\langle R \rangle$ *child*, *dog* $\langle R \rangle$ *cat*, *dog* $\langle R \rangle$ *rabbit*, *dog* $\langle R \rangle$ *monkey*, *dog* $\langle R \rangle$ *dog*, *dog* $\langle R \rangle$ *hamster*, *dog* $\langle R \rangle$ *kangaroo*, *dog* $\langle R \rangle$ *panda*,
- *hamster* $\langle R \rangle$ *man*, *hamster* $\langle R \rangle$ *woman*, *hamster* $\langle R \rangle$ *child*, *hamster* $\langle R \rangle$ *cat*, *hamster* $\langle R \rangle$ *rabbit*, *hamster* $\langle R \rangle$ *monkey*, *hamster* $\langle R \rangle$ *dog*, *hamster* $\langle R \rangle$ *hamster*, *hamster* $\langle R \rangle$ *kangaroo*, *hamster* $\langle R \rangle$ *panda*,
- *kangaroo* $\langle R \rangle$ *man*, *kangaroo* $\langle R \rangle$ *woman*, *kangaroo* $\langle R \rangle$ *child*, *kangaroo* $\langle R \rangle$ *cat*, *kangaroo* $\langle R \rangle$ *rabbit*, *kangaroo* $\langle R \rangle$ *monkey*, *kangaroo* $\langle R \rangle$ *dog*, *kangaroo* $\langle R \rangle$ *hamster*, *kangaroo* $\langle R \rangle$ *kangaroo*, *kangaroo* $\langle R \rangle$ *panda*,
- *panda* $\langle R \rangle$ *man*, *panda* $\langle R \rangle$ *woman*, *panda* $\langle R \rangle$ *child*, *panda* $\langle R \rangle$ *cat*, *panda* $\langle R \rangle$ *rabbit*, *panda* $\langle R \rangle$ *monkey*, *panda* $\langle R \rangle$ *dog*, *panda* $\langle R \rangle$ *hamster*, *panda* $\langle R \rangle$ *kangaroo*, *panda* $\langle R \rangle$ *panda*

C FURTHER EXPLANATIONS ON BASIS PREPOSITIONS

As stated in the manuscript, we devise a set of basis prepositions to steer the learning process of the relation prompt. Specifically, we collect a comprehensive list of ~ 100 prepositions from [Stevenson 2010], and drop the prepositions that describes non-visual relations (i.e., temporal relations, causal relations, etc.), while keep the ones that are related to visual relations. For example, the prepositional word “until” is discarded as a temporal preposition, while words like “above”, “beneath”, “toward” will be kept as plausible basis prepositions.

The basis preposition set contains a total of 56 words, listed in Table A4.

D LIMITATIONS

Our performance is capped by the generative capabilities of the pre-trained text-to-image model, Stable Diffusion (SD). This dependency might lead to suboptimal synthesis in scenarios where SD faces challenges, as shown in Figure A14.

Concept Blending. SD suffers from the concept blending problem. This issue arises when the model generates multiple entities within a single scene, leading to a fusion of characteristics from different classes. For example, when tasked with depicting a “rabbit” and a “cat” together, SD creates entities that blend features of both - such as rabbit ears and cat-like fur color and texture. Consequently, when ReVersion applies the learned $\langle R \rangle$ on two entities of different classes, the same issue might occur.

Human. When SD attempts to render human faces and bodies, the outcomes are often less than ideal. Consequently, even though ReVersion effectively captures the relation, the quality of the faces and bodies of the human subjects might remain suboptimal.

Given that these limitations are inherent to the pre-trained text-to-image model, exploring and developing better text-to-image diffusion models is an orthogonal direction for performance improvements.

E POTENTIAL SOCIETAL IMPACTS

Although *ReVersion* can generate diverse entity combinations through inverted relations, this capability can also be exploited to synthesize real human figures interacting in ways they never did. As a result, we strongly advise users to only use *ReVersion* for proper recreational purposes.

The rapid advancement of generative models has unlocked new levels of creativity but has also introduced various societal concerns. First, it is easier to create false imagery or manipulate data maliciously, leading to the spread of misinformation. Second, data used to train these models might be revealed during the sampling process without explicit consent from the data owner [Tinsley et al. 2021]. Third, generative models can suffer from the biases present in the training data [Esser et al. 2020]. We used the pre-trained Stable Diffusion [Rombach et al. 2022] for *ReVersion*, which has been shown to suffer from data bias in certain scenarios. For example, when prompted with the phrase “a professor”, Stable Diffusion tends to generate human figures that are white-passing and male-passing. We hope that more research will be conducted to address the risks and biases associated with generative models, and we advise everyone to use these models with discretion.

F MORE QUALITATIVE RESULTS

We show various qualitative results in Figure A15-A21, which are located at the end of this Supplementary File.

F.1 ReVersion with Diverse Styles and Backgrounds

As shown in Figure A15, we apply the $\langle R \rangle$ inverted by *ReVersion* in scenarios with diverse backgrounds and styles, and show that $\langle R \rangle$ robustly adapt these environments with impressive results.

F.2 ReVersion with Arbitrary Entity Combinations

In Figure A16 and A17, we show that the $\langle R \rangle$ inverted by *ReVersion* can be applied to robustly relate arbitrary entity combinations. For example, in Figure A16, for the $\langle R \rangle$ extracted from the exemplar images where one entity is “*painted on*” the other entity, we enumerate over all combinations among “*{cat / flower / guitar / hamburger / Michael Jackson / Spiderman} $\langle R \rangle$ {building / canvas / paper / vase / wall}*”, and observe that $\langle R \rangle$ successfully links these entities together via exactly the same relation in the exemplar images.

F.3 Additional Qualitative Results

We show additional qualitative results of *ReVersion* in Figure A18, A19, A20, and A21.



"girl <R> boy",
 "a girl with in white and green <R> a boy in white and light grey",
 "a girl wearing white T-shirt and green skirt <R> a boy in white T-shirt and grey shorts, white background"



"cat <R> cat",
 "a long haired cat <R> a long haired cat",
 "a dark long haired cat <R> a grey long haired cat, white background"



"woman <R> man",
 "a woman wearing in white trousers and blue shirt <R> a man in grey",
 "a woman wearing in white trousers and blue shirt <R> a man in khaki trousers and light grey shirt, white background"



"girl <R> boy",
 "a girl with in pink top and jeans <R> a boy with striped t-shirt and jeans",
 "a girl with in pink top and jeans <R> a boy with striped t-shirt and jeans, grey sofa in background"



"boy <R> boy",
 "a boy with shirt and trousers <R> another boy with shirt and trousers",
 "a boy with shirt and trousers <R> another boy with shirt and trousers, white background"



"bear <R> bear",
 "a bear <R> a bear in wooded area",
 "a bear <R> a bear, bush in background"



"girl <R> boy",
 "a young girl in purple dress <R> a young boy in white",
 "a young girl in purple dress <R> a young boy in white, in the field"



"girl <R> boy",
 "a teenager girl <R> a teenager boy, white background",
 "a teenager girl wearing red shirt and jeans <R> a teenager boy in blue shirt and khaki trousers, white background"



"cat <R> cat",
 "an orange cat <R> a brown and white cat",
 "an orange cat <R> a brown and white cat on a wooden bench, grasses in background"



"boy <R> boy",
 "a boy with shirt and jeans <R> a boy in shirt and jeans",
 "a boy wearing shirt and jeans <R> another boy wearing shirt and jeans, white background"

Fig. A13. **Benchmark Sample.** We present *exemplar images* and *text descriptions* that illustrate the relation where " E_A sits back to back with E_B ". The *exemplar images* feature both human figures and animals to demonstrate the invariant "back to back" relationship in various scenarios. The *text descriptions* are provided at several levels, ranging from simple class name mentions to detailed descriptions of the entities and their surroundings. During optimization, the $\langle R \rangle$ in each description will be replaced with the learnable relation prompt.

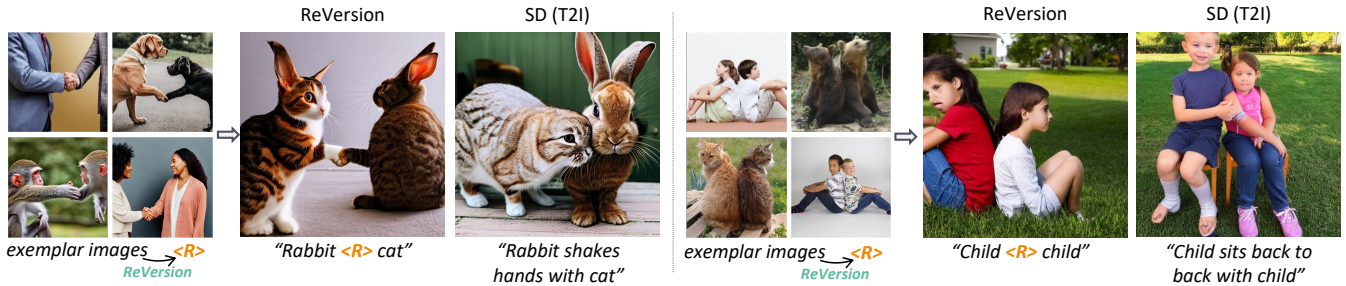


Fig. A14. **Limitations.** Although the $\langle R \rangle$ inverted by ReVersion can be applied robustly to synthesize new scenes, the image quality is limited by the generative capability of the pre-trained text-to-image model. **Left:** when tasked with depicting a “rabbit” and a “cat” together, Stable Diffusion (SD) creates entities that blend features of both - such as rabbit ears and cat-like fur color and texture. Despite ReVersion’s ability in capturing the “shake hand” relation through $\langle R \rangle$, the resulting image still has the problem of concept blending. **Right:** when SD attempts to render human faces and bodies, the outcomes are often less than ideal. Therefore, even though ReVersion effectively captures the “sitting back to back” relation, the quality of the faces and bodies of the two children remains suboptimal.

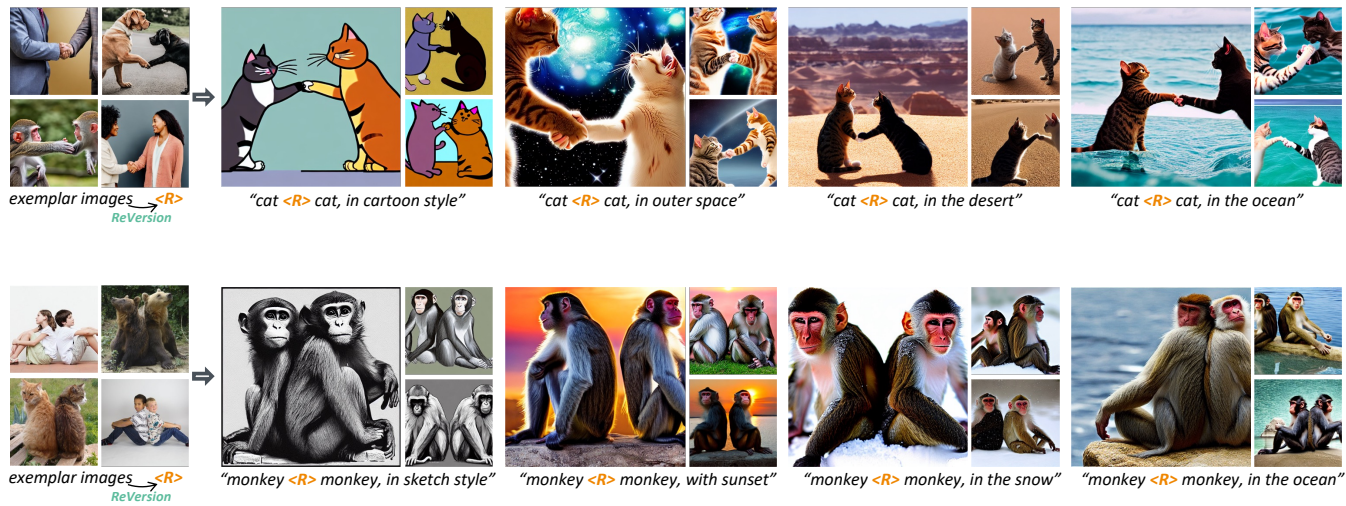


Fig. A15. **ReVersion for Diverse Styles and Backgrounds.** The $\langle R \rangle$ inverted by ReVersion can be applied robustly to relate entities in scenes with diverse backgrounds or styles.

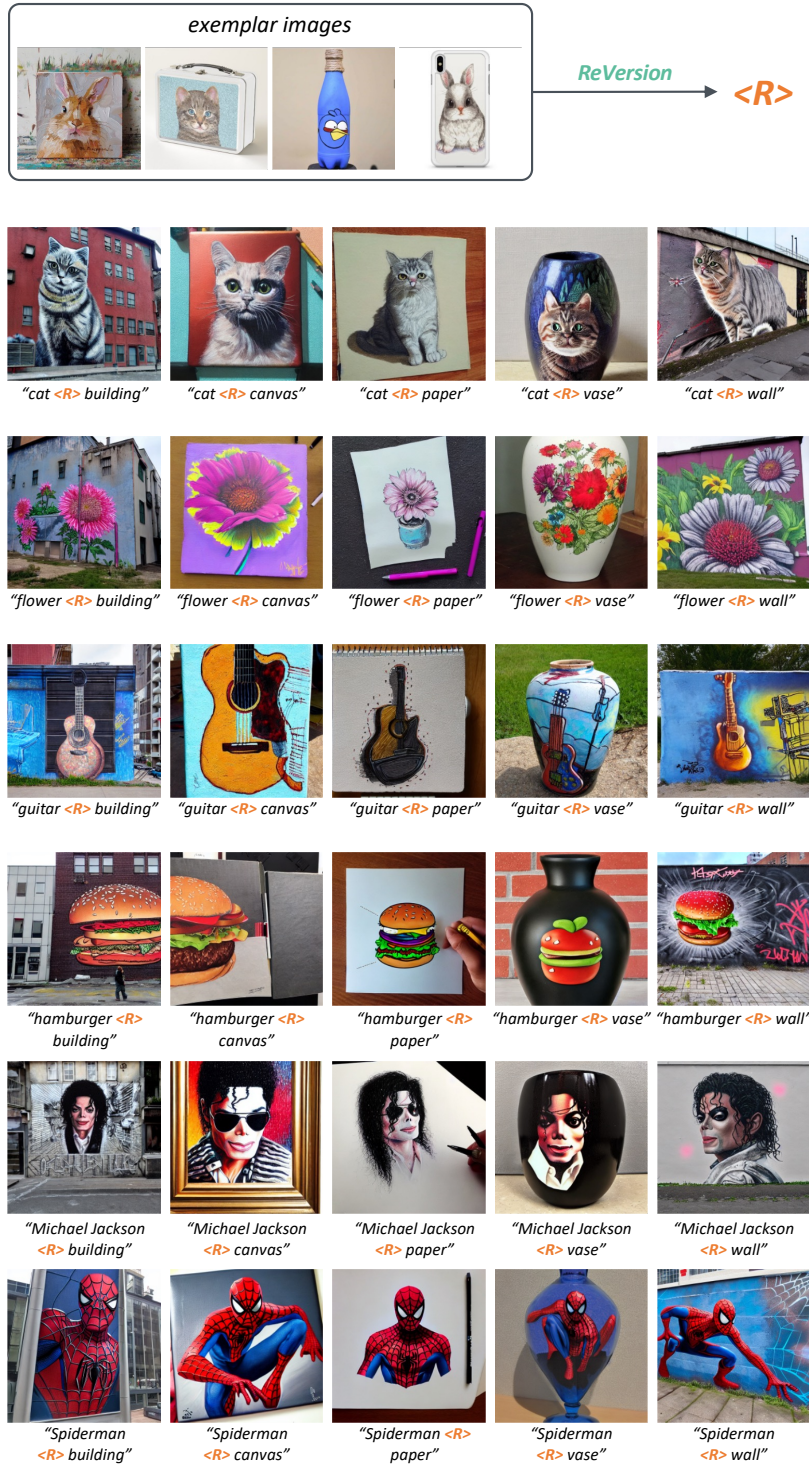


Fig. A16. **Arbitrary Entity Combinations.** The $\langle R \rangle$ inverted by ReVersion can be robustly applied to arbitrary entity combinations. For example, for the $\langle R \rangle$ extracted from the exemplar images where one entity is "painted on" the other entity, we enumerate over all combinations among $\{cat / flower / guitar / hamburger / Michael Jackson / Spiderman\} \langle R \rangle \{building / canvas / paper / vase / wall\}$, and observe that $\langle R \rangle$ successfully links these entities together via exactly the same relation in the exemplar images.

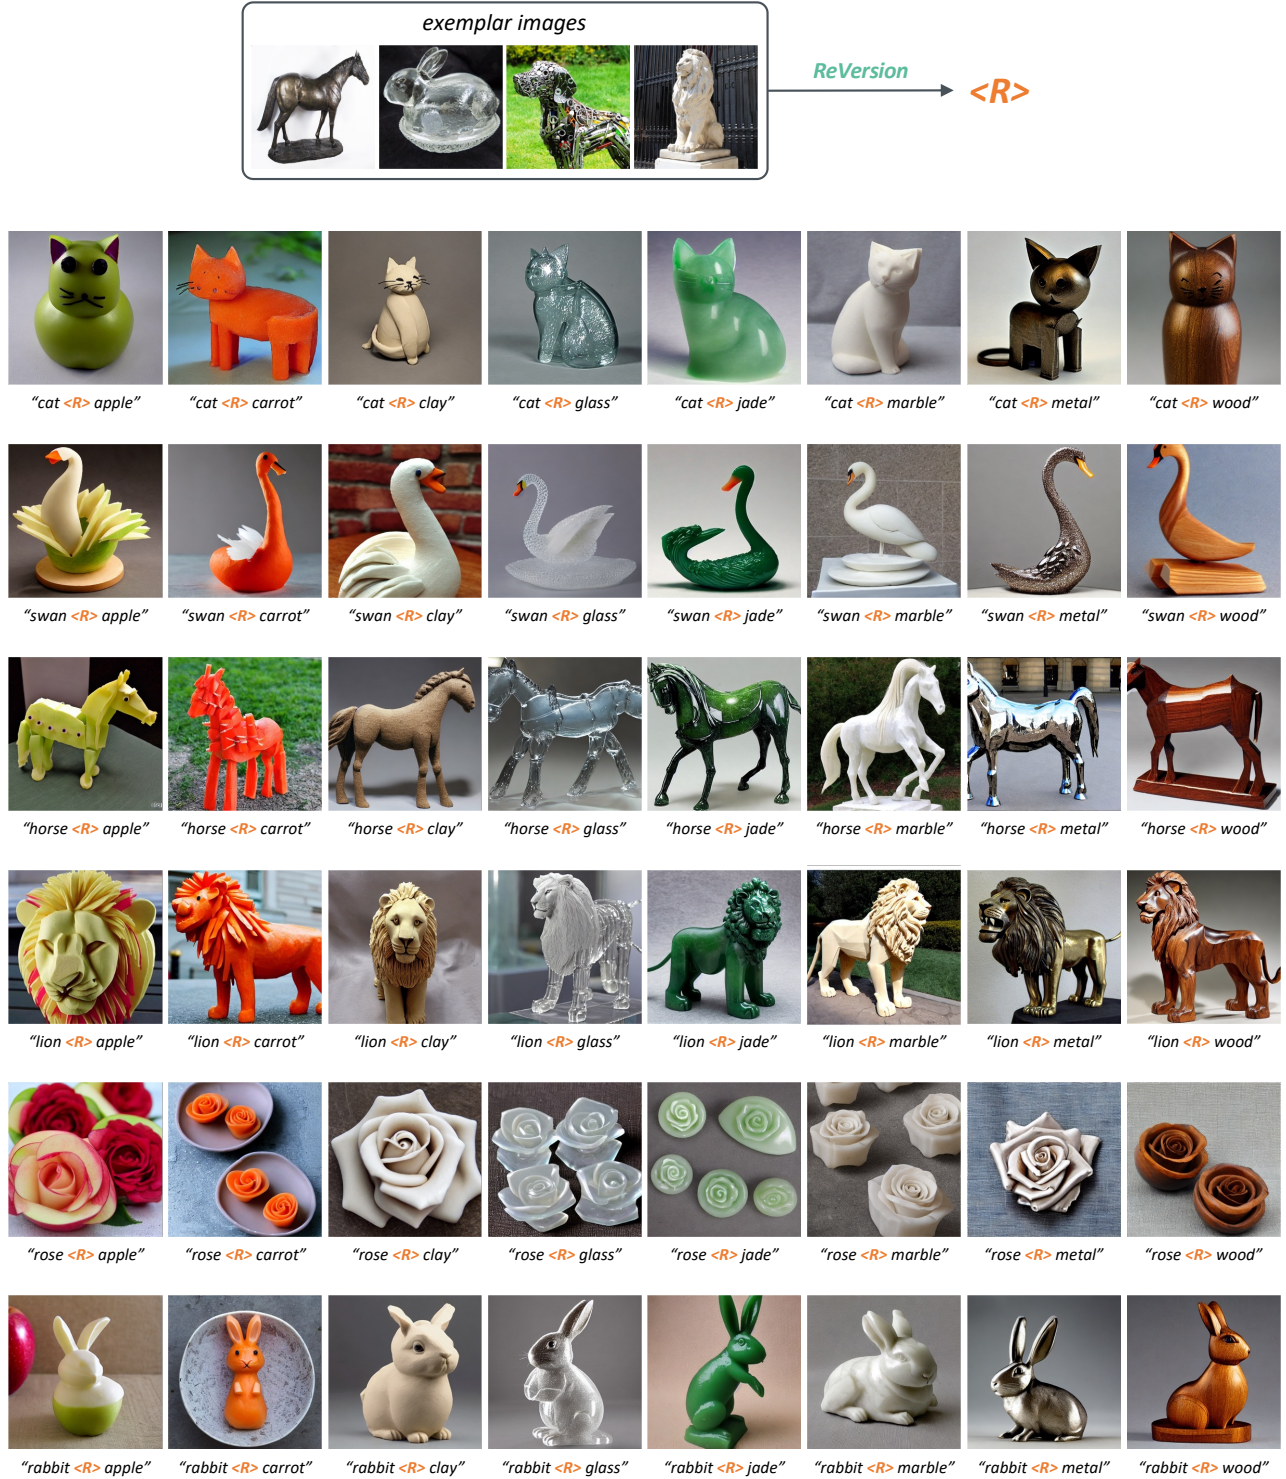


Fig. A17. **Arbitrary Entity Combinations.** The $\langle R \rangle$ inverted by ReVersion can be applied to arbitrary entity combinations. For example, for the $\langle R \rangle$ extracted from the exemplar images where one entity is “is made of the material of / is carved by” the other entity, we enumerate over all combinations among “{cat / swan / horse / lion / rose / rabbit} $\langle R \rangle$ {apple / carrot / clay / glass / jade / marble / metal / wood}”, and observe that $\langle R \rangle$ successfully links these entities together via exactly the same relation in the exemplar images.

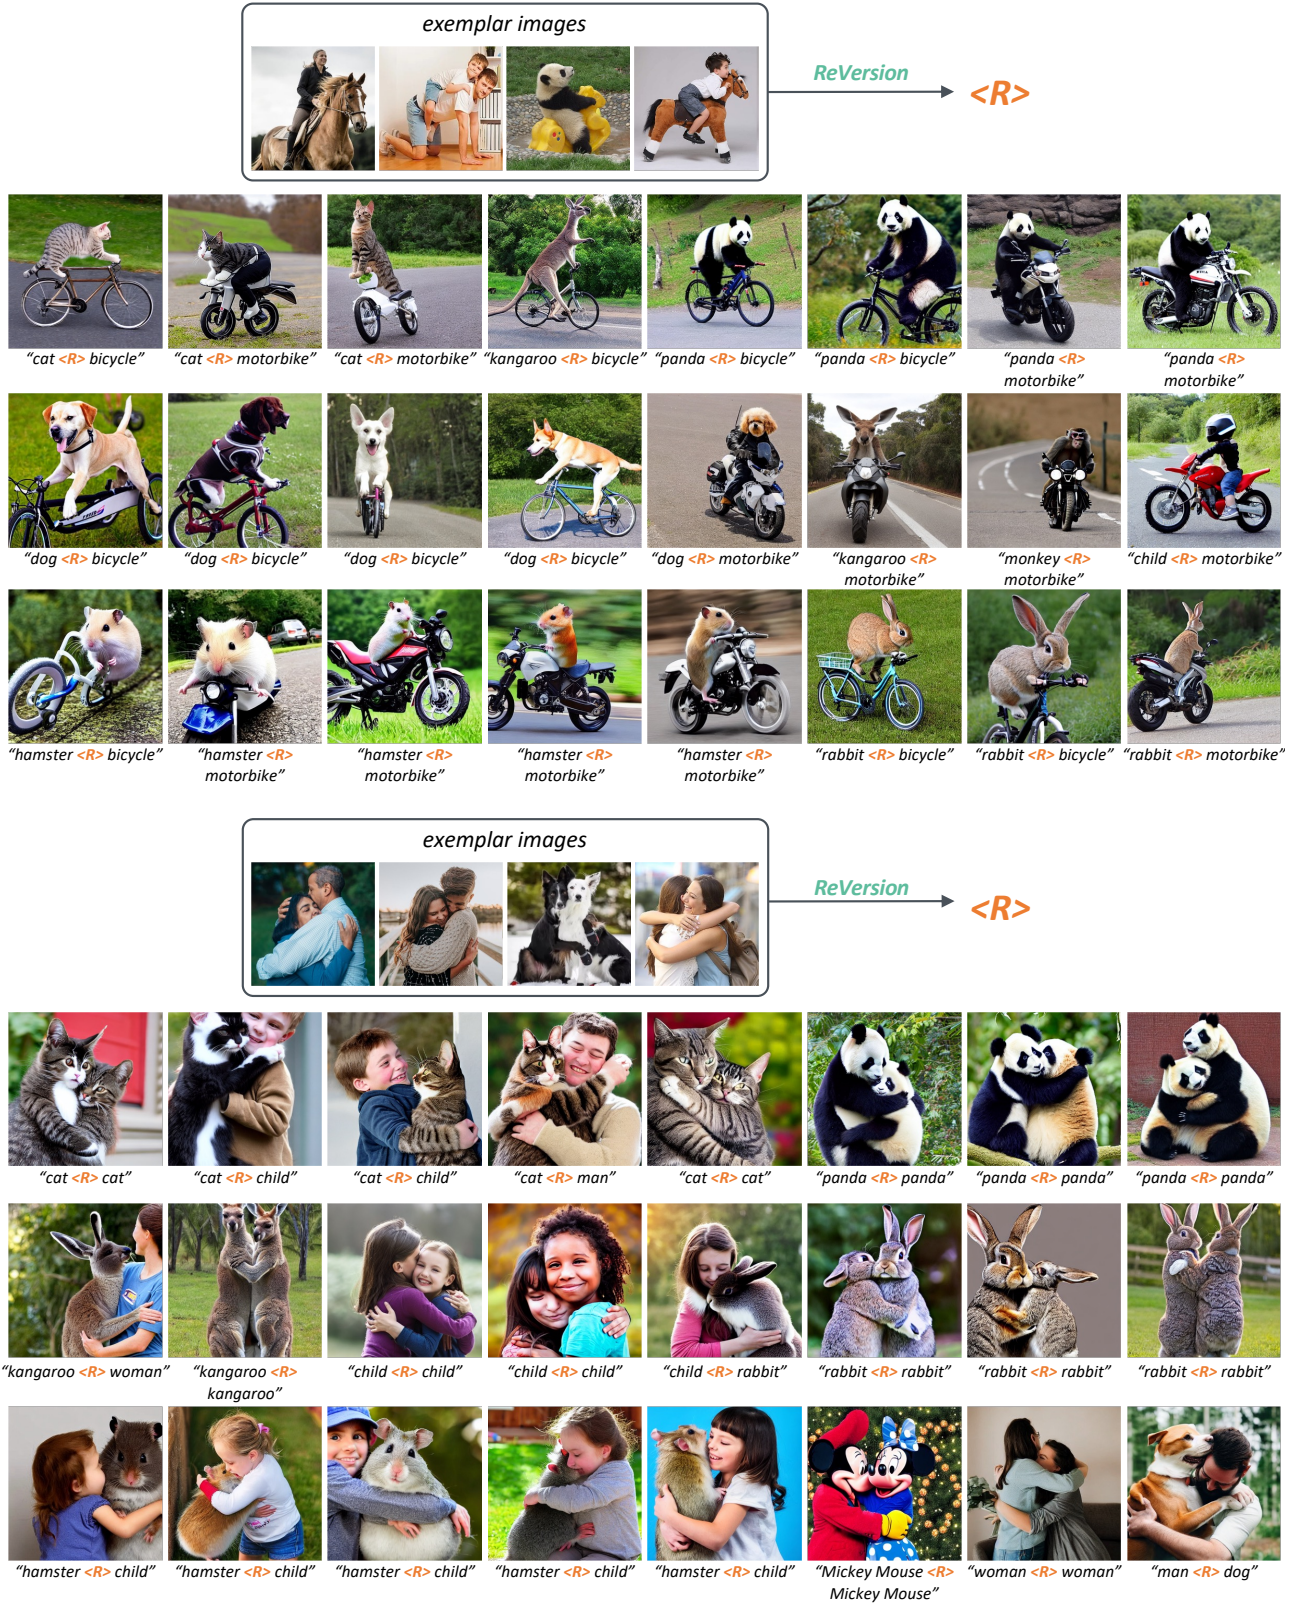


Fig. A18. More Qualitative Results.

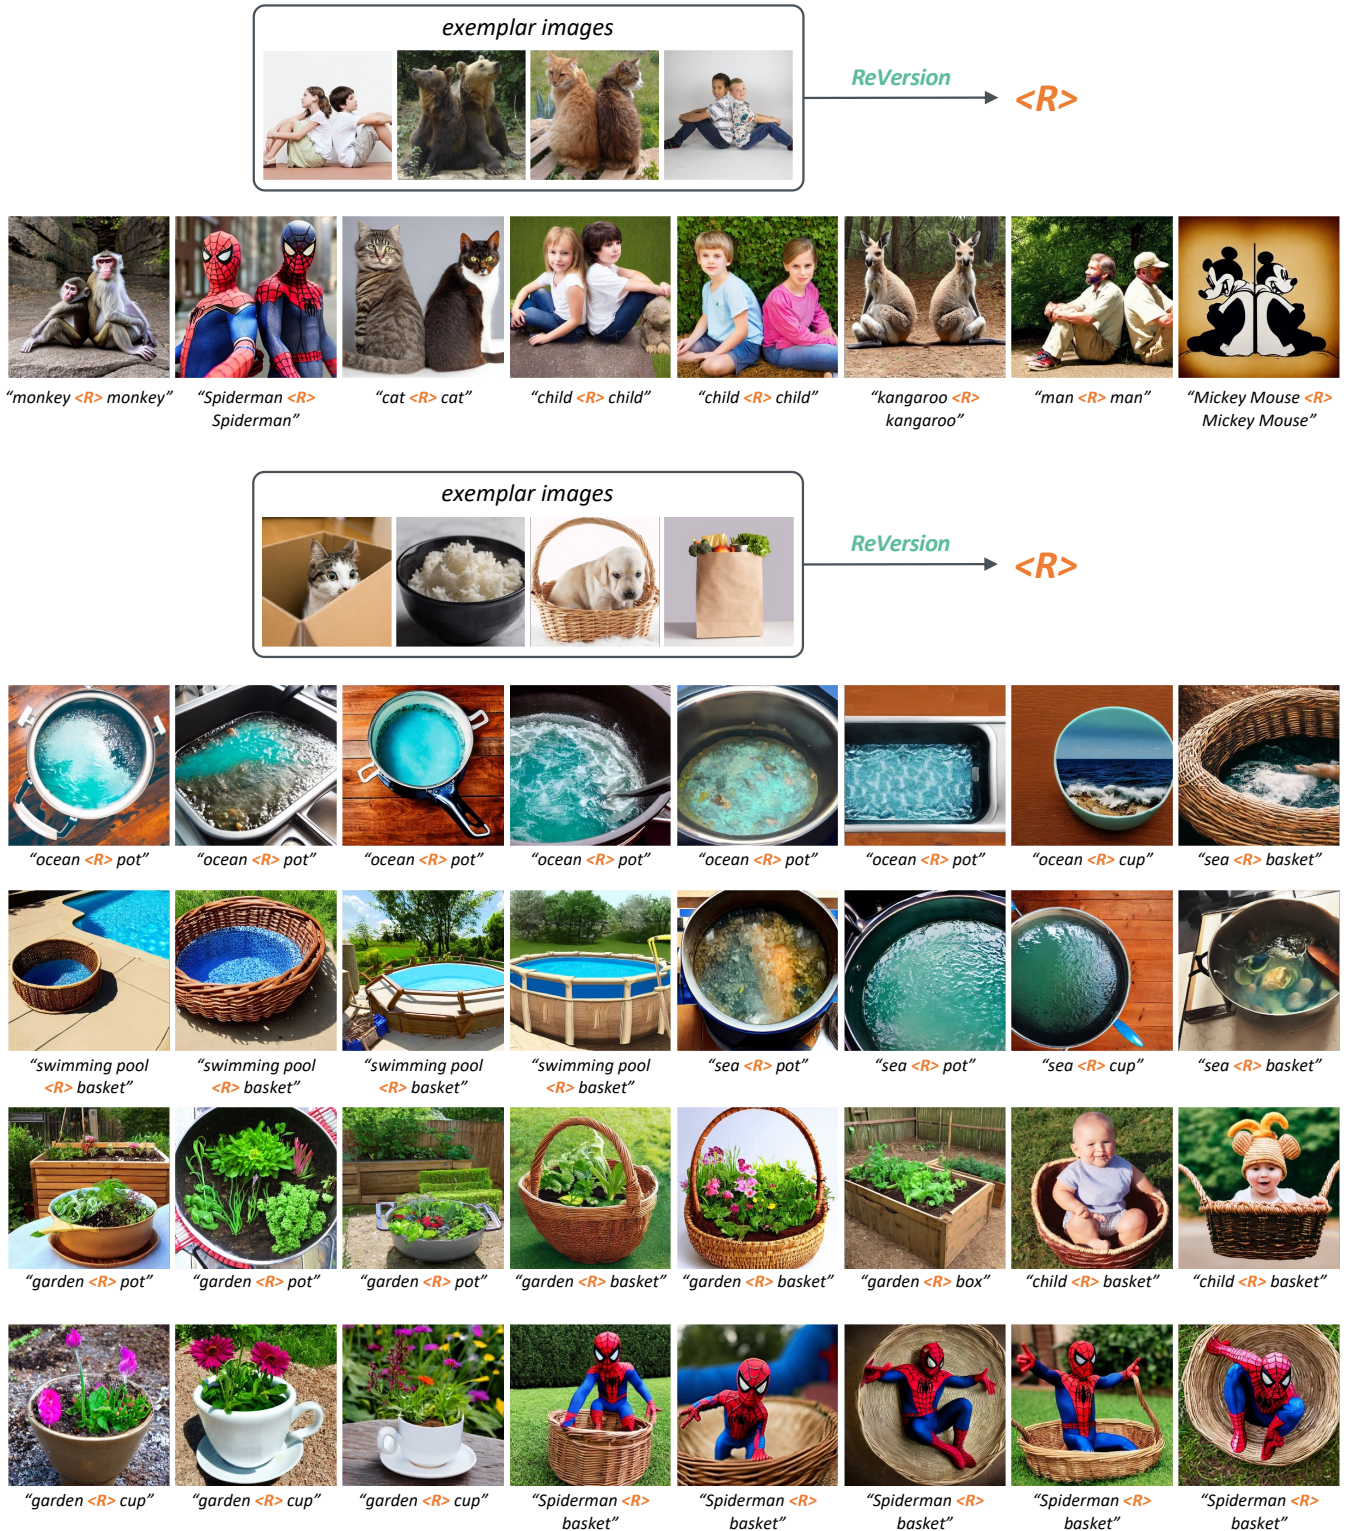


Fig. A19. More Qualitative Results.

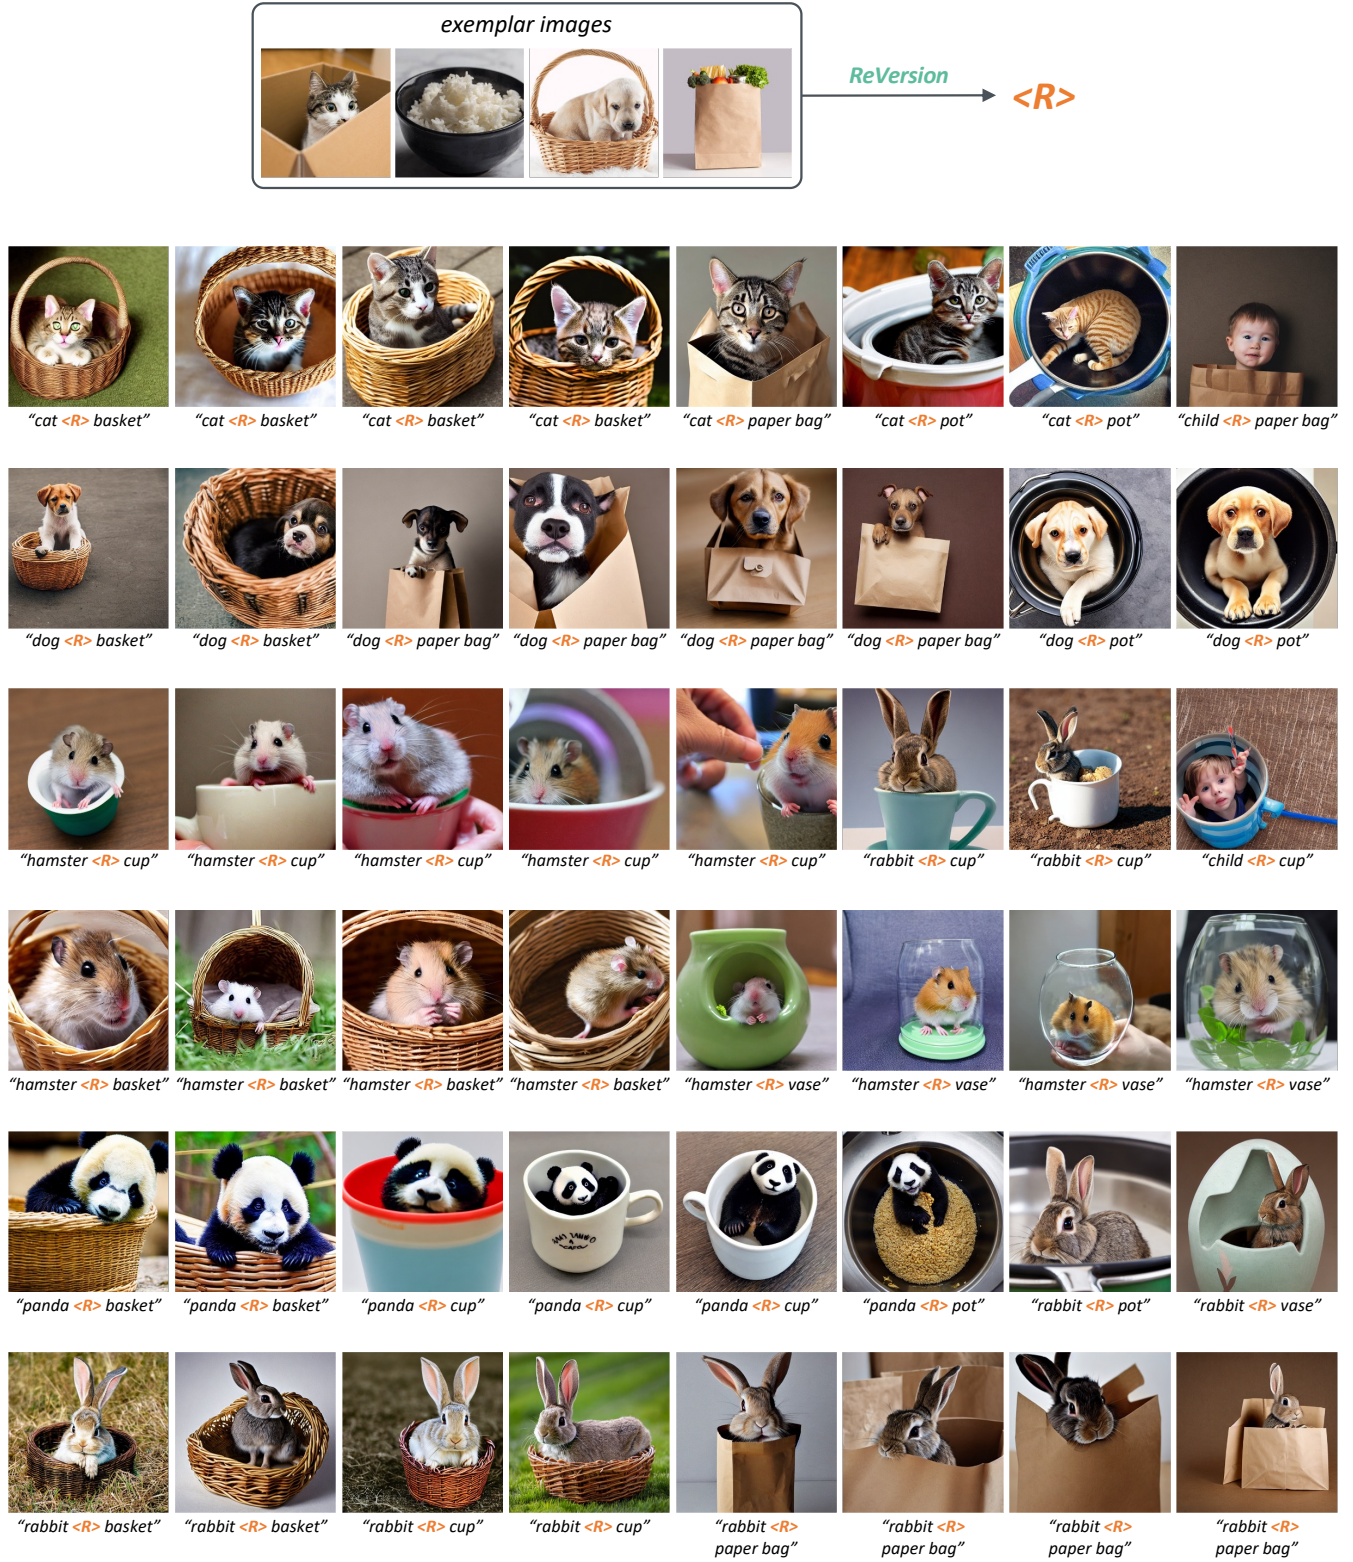


Fig. A20. More Qualitative Results.

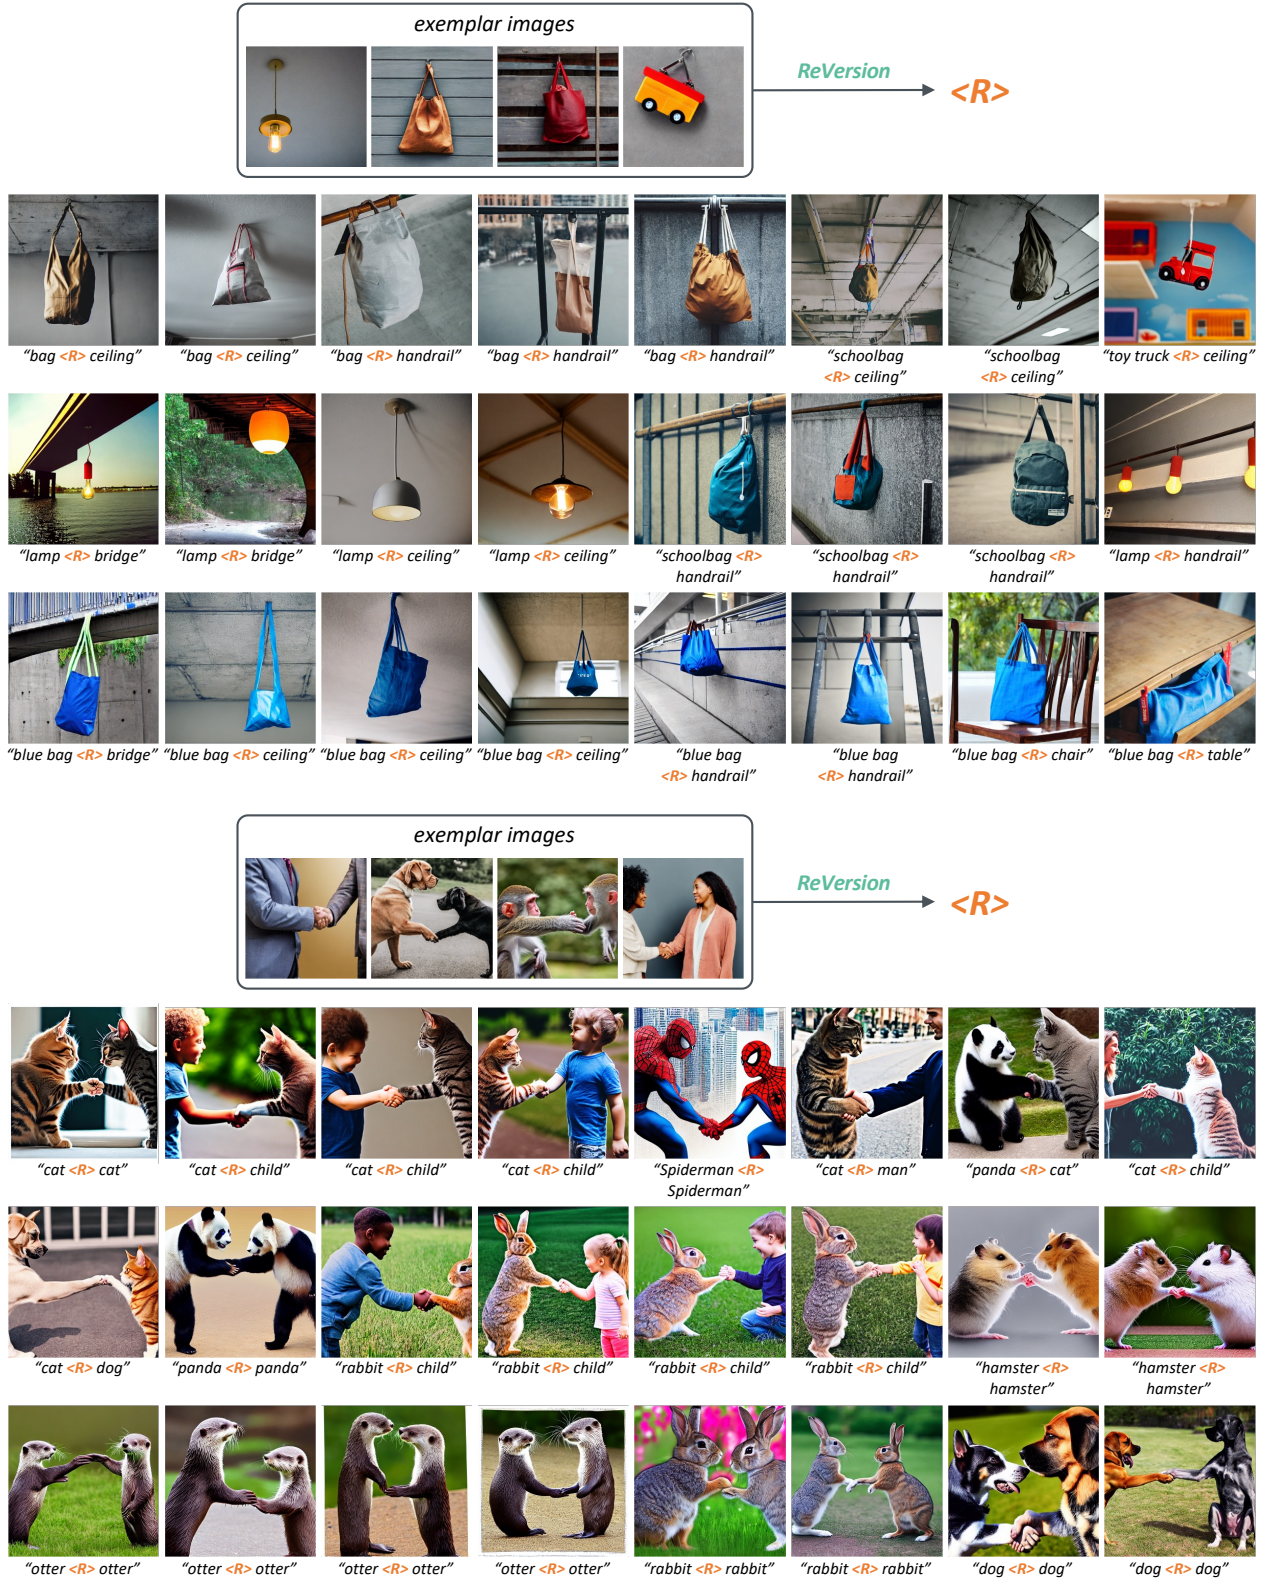


Fig. A21. More Qualitative Results.