VAR RL Done Right: Tackling Asynchronous Policy Conflicts in Visual Autoregressive Generation

Supplementary Material

6. Acknowledgments

This work is supported by the National Natural Science Foundation of China No. 62425604.

7. Details About Nextflow

7.1. Architecture

The Nextflow main Transformer is initialized from Qwen2.5-VL-7B [2] and augmented with a newly introduced visual codebook and a revised logits-prediction head.

7.2. Vision Generation

For image synthesis, the model first emits the special token ⟨boi⟩ (begin-of-image) and then switches to full attention, operating in the style of VAR [46]. Scale-specific configurations are provided in Table 8.

Table 8. Related tokens corresponding to different prefix resolutions are provided. Note that the method also supports varying aspect ratios.

Resolution	Related Schedule
64×64	(1,1) (2,2) (3,3) (4,4)
128×128	(1,1) (2,2) (3,3) (4,4) (5,5) (6,6) (7,7) (8,8)
256×256	(1,1) (2,2) (3,3) (4,4) (5,5) (6,6) (7,7) (8,8) (10,10) (12,12) (14,14) (16,16)
512×512	(1,1) (2,2) (3,3) (4,4) (5,5) (6,6) (7,7) (8,8) (10,10) (12,12) (14,14) (16,16) (20,20) (24,24) (28,28) (32,32)
1024×1024	(1,1) (2,2) (3,3) (4,4) (5,5) (6,6) (7,7) (8,8) (10,10) (12,12) (14,14) (16,16) (20,20) (24,24) (28,28) (32,32) (48,48) (64,64)

8. Captions

The detailed prompts for Fig. 5 and Fig. 6 is shown in Tab. 10 and Tab. 9.

9. Distribution Gap between the Training and Evaluation Datasets

Our in-house training corpus exhibits a pronounced imbalance in the *number of text-rendering regions* per sample: an overabundance of single-region cases and a long tail of images containing more than five regions. To mitigate this skew, we adopt a region-count–based filtering strategy for the training data. For analysis and visualization, we discretize the per-sample region count into six categories—1, 2, 3, 4, 5, and >5—using the >5 bin to summarize the long tail. The evaluation set follows a target profile restricted to bins 2–5 with probabilities 0.2, 0.3, 0.3, and 0.2, respectively. Fig. 8 compares the empirical distributions for the pre-filter training set, the post-filter training set, and the evaluation set. Importantly, this filtering does not constitute evaluation hacking or test-set leakage; rather, it calibrates the training distribution to the intended task difficulty.

10. More Visual Results

The visual samples from the HPSv3 evaluation dataset for each class are shown in Fig. 9 and Fig. 10.

11. Proofs

11.1. Proof of Theorem 1

Theorem 1 (Reverse-KL characterization of $\pi^\dagger$)**.** *At each state $\mathbf{s}_t$, the constrained optimum satisfies*

$$\pi^\dagger(\cdot \mid \mathbf{s}_t) = \arg\min_{\pi \in \mathcal{M}_\pi(\mathbf{s}_t)} \mathrm{KL}\big(\pi(\cdot \mid \mathbf{s}_t) \,\|\, \pi^*(\cdot \mid \mathbf{s}_t)\big), \tag{8}$$

where π^* is the global soft-optimal policy.

Proof. We first derive the trajectory-level variational identity, then identify the global soft-optimal policy π^* and its associated Q^*, and finally show that the constrained optimal policy $\pi^\dagger$ is a reverse-KL projection of π^* at each state.

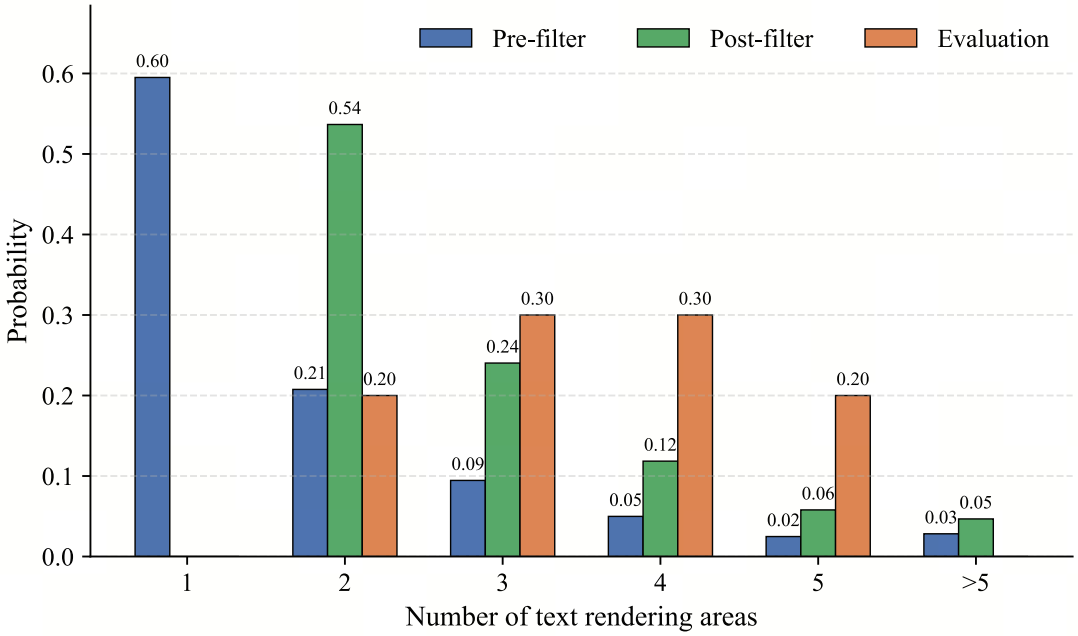

Figure 8. Distributions over the number of text-rendering regions (1, 2, 3, 4, 5, and >5) for the pre-filter training set, the post-filter training set, and the evaluation set. The evaluation distribution is constrained to bins 2–5 with probabilities 0.2, 0.3, 0.3, and 0.2.

1. Trajectory-level formulation of $J(\pi)$. Let a trajectory be

$$\tau = (\mathbf{s}_1, \mathbf{a}_1, \ldots, \mathbf{s}_T), \tag{16}$$

and denote by $p_\pi(\tau)$ the trajectory distribution induced by a policy π together with the environment dynamics. For a fixed reference policy π_{old}, let $p_{\text{old}}(\tau)$ be the corresponding trajectory distribution.

The KL-regularized objective is

$$J(\pi) = \mathbb{E}_{p_\pi(\tau)}[R(\mathbf{s}_T)] - \eta\,\mathrm{KL}\big(p_\pi(\tau)\big\|p_{\text{old}}(\tau)\big) \tag{17}$$

$$= \mathbb{E}_{p_\pi(\tau)}[R(\mathbf{s}_T)] - \eta\,\mathbb{E}_{p_\pi(\tau)}\left[\log \frac{p_\pi(\tau)}{p_{\text{old}}(\tau)}\right]. \tag{18}$$

Using the factorization

$$p_\pi(\tau) = p(\mathbf{s}_1)\prod_{t=1}^{T-1} \pi(\mathbf{a}_t \mid \mathbf{s}_t)\, p(\mathbf{s}_{t+1} \mid \mathbf{s}_t, \mathbf{a}_t), \tag{19}$$

and similarly for $p_{\text{old}}(\tau)$ with π_{old}, the transition dynamics and initial-state distribution cancel in the ratio, giving

$$\log \frac{p_\pi(\tau)}{p_{\text{old}}(\tau)} = \sum_{t=1}^{T-1} \log \frac{\pi(\mathbf{a}_t \mid \mathbf{s}_t)}{\pi_{\text{old}}(\mathbf{a}_t \mid \mathbf{s}_t)}. \tag{20}$$

Thus

$$J(\pi) = \mathbb{E}_{p_\pi(\tau)}\left[R(\mathbf{s}_T) - \eta \sum_{t=1}^{T-1} \log \frac{\pi(\mathbf{a}_t \mid \mathbf{s}_t)}{\pi_{\text{old}}(\mathbf{a}_t \mid \mathbf{s}_t)}\right]. \tag{21}$$

2. Variational identity and global soft-optimal trajectory distribution. Define an unnormalized density over trajectories

$$\tilde{p}(\tau) = p_{\text{old}}(\tau)\exp\Big(\tfrac{1}{\eta}R(\mathbf{s}_T)\Big), \tag{22}$$

with partition function

$$Z = \int \tilde{p}(\tau)\, d\tau = \mathbb{E}_{p_{\text{old}}(\tau)}\left[\exp\left(\tfrac{1}{\eta} R(\mathbf{s}_T)\right)\right]. \tag{23}$$

Define the normalized distribution

$$p^*(\tau) = \frac{1}{Z}\tilde{p}(\tau) = \frac{1}{Z} p_{\text{old}}(\tau) \exp\left(\tfrac{1}{\eta} R(\mathbf{s}_T)\right). \tag{24}$$

We now relate $J(\pi)$ to the KL divergence $\mathrm{KL}(p_\pi(\tau) \,\|\, p^*(\tau))$. Compute

$$\begin{aligned}
\mathrm{KL}\big(p_\pi(\tau) \,\big\|\, p^*(\tau)\big) &= \int p_\pi(\tau) \log \frac{p_\pi(\tau)}{p^*(\tau)}\, d\tau && (25)\\
&= \int p_\pi(\tau) \left[\log \frac{p_\pi(\tau)}{p_{\text{old}}(\tau)} - \tfrac{1}{\eta} R(\mathbf{s}_T) + \log Z\right] d\tau && (26)\\
&= \mathbb{E}_{p_\pi}\left[\log \frac{p_\pi(\tau)}{p_{\text{old}}(\tau)}\right] - \tfrac{1}{\eta}\mathbb{E}_{p_\pi}[R(\mathbf{s}_T)] + \log Z. && (27)
\end{aligned}$$

Multiplying both sides by η gives

$$\eta\, \mathrm{KL}\big(p_\pi(\tau) \,\big\|\, p^*(\tau)\big) = \eta\, \mathrm{KL}\big(p_\pi(\tau) \,\big\|\, p_{\text{old}}(\tau)\big) - \mathbb{E}_{p_\pi}[R(\mathbf{s}_T)] + \eta \log Z. \tag{28}$$

Rearranging (28) yields the trajectory-level variational identity:

$$\mathbb{E}_{p_\pi}[R(\mathbf{s}_T)] - \eta\, \mathrm{KL}\big(p_\pi(\tau) \big\| p_{\text{old}}(\tau)\big) = \eta \log Z - \eta\, \mathrm{KL}\big(p_\pi(\tau) \big\| p^*(\tau)\big). \tag{29}$$

Thus

$$J(\pi) = \eta \log Z - \eta\, \mathrm{KL}\big(p_\pi(\tau) \big\| p^*(\tau)\big), \tag{30}$$

and $J(\pi)$ is maximized exactly when $p_\pi(\tau) = p^*(\tau)$.

3. Factorization of $p^*(\tau)$ and the soft-optimal policy. Since $p_{\text{old}}(\tau)$ is induced by the Markov policy π_{old} and the environment dynamics, it factorizes as

$$p_{\text{old}}(\tau) = p(\mathbf{s}_1) \prod_{t=1}^{T-1} \pi_{\text{old}}(\mathbf{a}_t \mid \mathbf{s}_t)\, p(\mathbf{s}_{t+1} \mid \mathbf{s}_t, \mathbf{a}_t). \tag{31}$$

Hence

$$p^*(\tau) = \frac{1}{Z} p(\mathbf{s}_1) \prod_{t=1}^{T-1} \pi_{\text{old}}(\mathbf{a}_t \mid \mathbf{s}_t)\, p(\mathbf{s}_{t+1} \mid \mathbf{s}_t, \mathbf{a}_t)\, \exp\left(\tfrac{1}{\eta} R(\mathbf{s}_T)\right). \tag{32}$$

This defines a Markov trajectory distribution; thus there exists a policy π^* such that

$$p^*(\tau) = p(\mathbf{s}_1) \prod_{t=1}^{T-1} \pi^*(\mathbf{a}_t \mid \mathbf{s}_t)\, p(\mathbf{s}_{t+1} \mid \mathbf{s}_t, \mathbf{a}_t). \tag{33}$$

We now compute $\pi^*(\mathbf{a}_t \mid \mathbf{s}_t)$ explicitly. Fix a time t and a state $\mathbf{s}_t$. Using (32), the conditional distribution over $\mathbf{a}_t$ given $\mathbf{s}_t$ under p^* is proportional to the joint density over all trajectories sharing $(\mathbf{s}_t, \mathbf{a}_t)$:

$$\begin{aligned}
\pi^*(\mathbf{a}_t \mid \mathbf{s}_t) = p^*(\mathbf{a}_t \mid \mathbf{s}_t) &\propto \sum_{\text{trajectories consistent with } (\mathbf{s}_t, \mathbf{a}_t)} p^*(\tau) && (34)\\
&\propto \sum_{\text{futures}} p_{\text{old}}(\tau)\, \exp\left(\tfrac{1}{\eta} R(\mathbf{s}_T)\right), && (35)
\end{aligned}$$

where the sum is over the future part of the trajectory (from t onward), and the past part $\mathbf{s}_{1:t-1}, \mathbf{a}_{1:t-1}$ only contributes a constant factor w.r.t. $\mathbf{a}_t$. More precisely, for fixed $\mathbf{s}_t$ we have

$$\pi^*(\mathbf{a}_t \mid \mathbf{s}_t) \propto \pi_{\text{old}}(\mathbf{a}_t \mid \mathbf{s}_t) \sum_{\mathbf{s}_{t+1:T}, \mathbf{a}_{t+1:T-1}} \left[\prod_{k=t}^{T-1} p(\mathbf{s}_{k+1} \mid \mathbf{s}_k, \mathbf{a}_k) \pi_{\text{old}}(\mathbf{a}_{k+1} \mid \mathbf{s}_{k+1}) \right] \exp\Big(\tfrac{1}{\eta} R(\mathbf{s}_T)\Big) \tag{36}$$

$$= \pi_{\text{old}}(\mathbf{a}_t \mid \mathbf{s}_t)\, \mathbb{E}_{\pi_{\text{old}}} \left[\exp\Big(\tfrac{1}{\eta} R(\mathbf{s}_T)\Big) \;\Big|\; \mathbf{s}_t, \mathbf{a}_t \right]. \tag{37}$$

Thus we obtain the *general optimal solution*

$$\pi^*(\mathbf{a}_t \mid \mathbf{s}_t) \propto \pi_{\text{old}}(\mathbf{a}_t \mid \mathbf{s}_t) \exp\left(\frac{1}{\eta} Q^*(\mathbf{s}_t, \mathbf{a}_t)\right), \tag{38}$$

where we have defined the optimal soft Q-function as

$$Q^*(\mathbf{s}_t, \mathbf{a}_t) = \eta \ln \mathbb{E}_{\pi_{\text{old}}} \left[\exp\left(\frac{1}{\eta} R(\mathbf{s}_T)\right) \;\middle|\; \mathbf{s}_t, \mathbf{a}_t \right]. \tag{39}$$

Equations (1) and (3) fully characterize the global soft-optimal policy π^*.

4. Constrained optimum as trajectory-level reverse-KL projection. By the variational identity (29), for any policy π,

$$J(\pi) = \eta \log Z - \eta\, \mathrm{KL}\big(p_\pi(\tau) \big\| p^*(\tau)\big). \tag{40}$$

Therefore, when we restrict π to the constraint set $\mathcal{M}_\pi$ (Definition 3), the optimal policy

$$\pi^\dagger = \arg\max_{\pi \in \mathcal{M}_\pi} J(\pi) \tag{41}$$

is equivalently given by

$$\pi^\dagger = \arg\min_{\pi \in \mathcal{M}_\pi} \mathrm{KL}\big(p_\pi(\tau) \,\big\|\, p^*(\tau)\big). \tag{42}$$

5. Decomposition of trajectory KL into per-state KLs. Both $p_\pi(\tau)$ and $p^*(\tau)$ factorize according to the same dynamics:

$$p_\pi(\tau) = p(\mathbf{s}_1) \prod_{t=1}^{T-1} \pi(\mathbf{a}_t \mid \mathbf{s}_t)\, p(\mathbf{s}_{t+1} \mid \mathbf{s}_t, \mathbf{a}_t), \tag{43}$$

$$p^*(\tau) = p(\mathbf{s}_1) \prod_{t=1}^{T-1} \pi^*(\mathbf{a}_t \mid \mathbf{s}_t)\, p(\mathbf{s}_{t+1} \mid \mathbf{s}_t, \mathbf{a}_t). \tag{44}$$

Hence the KL divergence simplifies to

$$\mathrm{KL}\big(p_\pi(\tau) \,\big\|\, p^*(\tau)\big) = \mathbb{E}_{p_\pi(\tau)} \left[\log \frac{p_\pi(\tau)}{p^*(\tau)} \right] \tag{45}$$

$$= \mathbb{E}_{p_\pi(\tau)} \left[\sum_{t=1}^{T-1} \log \frac{\pi(\mathbf{a}_t \mid \mathbf{s}_t)}{\pi^*(\mathbf{a}_t \mid \mathbf{s}_t)} \right] \tag{46}$$

$$= \sum_{t=1}^{T-1} \mathbb{E}_{\mathbf{s}_t \sim d_\pi} \left[\mathrm{KL}\big(\pi(\cdot \mid \mathbf{s}_t) \,\big\|\, \pi^*(\cdot \mid \mathbf{s}_t)\big) \right], \tag{47}$$

where d_π denotes the (discount-free) state visitation distribution under π at time t. The key point is that the dynamics and initial-state distribution cancel and only the action distributions appear inside the KL.

6. Factorized constraint and per-state reverse-KL projection. The constraint set $\mathcal{M}_\pi$ is assumed to factorize across states:

$$\mathcal{M}_\pi = \prod_{\mathbf{s}} \mathcal{M}_\pi(\mathbf{s}), \tag{48}$$

where $\mathcal{M}_\pi(\mathbf{s})$ is the feasible set of action distributions at state $\mathbf{s}$. Because of this product structure, choosing $\pi \in \mathcal{M}_\pi$ amounts to choosing independently each $\pi(\cdot \mid \mathbf{s}) \in \mathcal{M}_\pi(\mathbf{s})$.

Since

$$\mathrm{KL}\big(p_\pi(\tau) \,\big\|\, p^*(\tau)\big) = \sum_t \mathbb{E}_{\mathbf{s}_t \sim d_\pi}\Big[\mathrm{KL}\big(\pi(\cdot \mid \mathbf{s}_t) \,\big\|\, \pi^*(\cdot \mid \mathbf{s}_t)\big)\Big], \tag{49}$$

minimizing this KL over the product set $\mathcal{M}_\pi = \prod_{\mathbf{s}} \mathcal{M}_\pi(\mathbf{s})$ decouples into independent minimizations at each state:

$$\pi^\dagger(\cdot \mid \mathbf{s}_t) = \arg\min_{\pi(\cdot|\mathbf{s}_t) \in \mathcal{M}_\pi(\mathbf{s}_t)} \mathrm{KL}\big(\pi(\cdot \mid \mathbf{s}_t) \,\big\|\, \pi^*(\cdot \mid \mathbf{s}_t)\big). \tag{50}$$

7. Conclusion. Combining the above steps, we conclude that the constrained optimal policy $\pi^\dagger$ is obtained by, at each state $\mathbf{s}_t$, projecting the global soft-optimal policy π^* onto the feasible action-distribution set $\mathcal{M}_\pi(\mathbf{s}_t)$ in the sense of reverse KL:

$$\pi^\dagger(\cdot \mid \mathbf{s}_t) = \arg\min_{\pi \in \mathcal{M}_\pi(\mathbf{s}_t)} \mathrm{KL}\big(\pi(\cdot \mid \mathbf{s}_t) \,\big\|\, \pi^*(\cdot \mid \mathbf{s}_t)\big). \tag{51}$$

This proves the theorem. □

11.2. Proof of Theorem 2

Theorem 2 (Two-stage invariance)**.** *Let $V_m^*(s_m)$ be defined as in (4). Within $\mathcal{M}_\pi$, solving the suffix problem in (5) yields $\pi^\dagger_{m:T-1}$; optimizing the prefix problem using V_m^* as its sole reward gives $\pi^\dagger_{1:m-1}$. Then, the concatenation $\pi^\dagger = \pi^\dagger_{1:m-1} \oplus \pi^\dagger_{m:T-1}$ uniquely maximizes the full-horizon objective $J(\pi)$. Within $\mathcal{M}_\pi$, replacing each subpolicy by its per-state reverse-KL projection onto the factorized family yields the constrained optimum $\pi^\dagger$.*

Proof. Write the full objective by splitting the action sequence into prefix and suffix:

$$J(\pi) = \underbrace{\mathbb{E}\big[R(s_T) \,\big|\, s_m, \pi_{m:T-1}\big] - \eta\,\mathrm{KL}\big(\pi_{m:T-1} \,\big\|\, \pi_{\mathrm{old},\, m:T-1}\big)}_{\triangleq\, J_{\mathrm{suffix}}(s_m, \pi_{m:T-1})} + \underbrace{-\eta\,\mathrm{KL}\big(\pi_{1:m-1} \,\big\|\, \pi_{\mathrm{old},\, 1:m-1}\big)}_{\text{prefix KL}}, \tag{52}$$

and take the outer expectation over s_m induced by the prefix policy $\pi_{1:m-1}$. For any fixed s_m, the inner maximization over suffix policies is exactly the standard soft-control problem whose value is $V_m^*(s_m)$. Therefore

$$\max_\pi J(\pi) = \max_{\pi_{1:m-1}} \Big\{\mathbb{E}\big[V_m^*(s_m) \,\big|\, \pi_{1:m-1}\big] - \eta\,\mathrm{KL}\big(\pi_{1:m-1} \,\big\|\, \pi_{\mathrm{old},\, 1:m-1}\big)\Big\}, \tag{53}$$

and the maximizer is obtained by (i) choosing the suffix soft-optimal $\pi^*_{m:T-1}$ for each s_m, and (ii) choosing the prefix policy $\pi^*_{1:m-1}$ that maximizes the soft value of s_m. Concatenation yields the unique full-horizon maximizer π^*. □

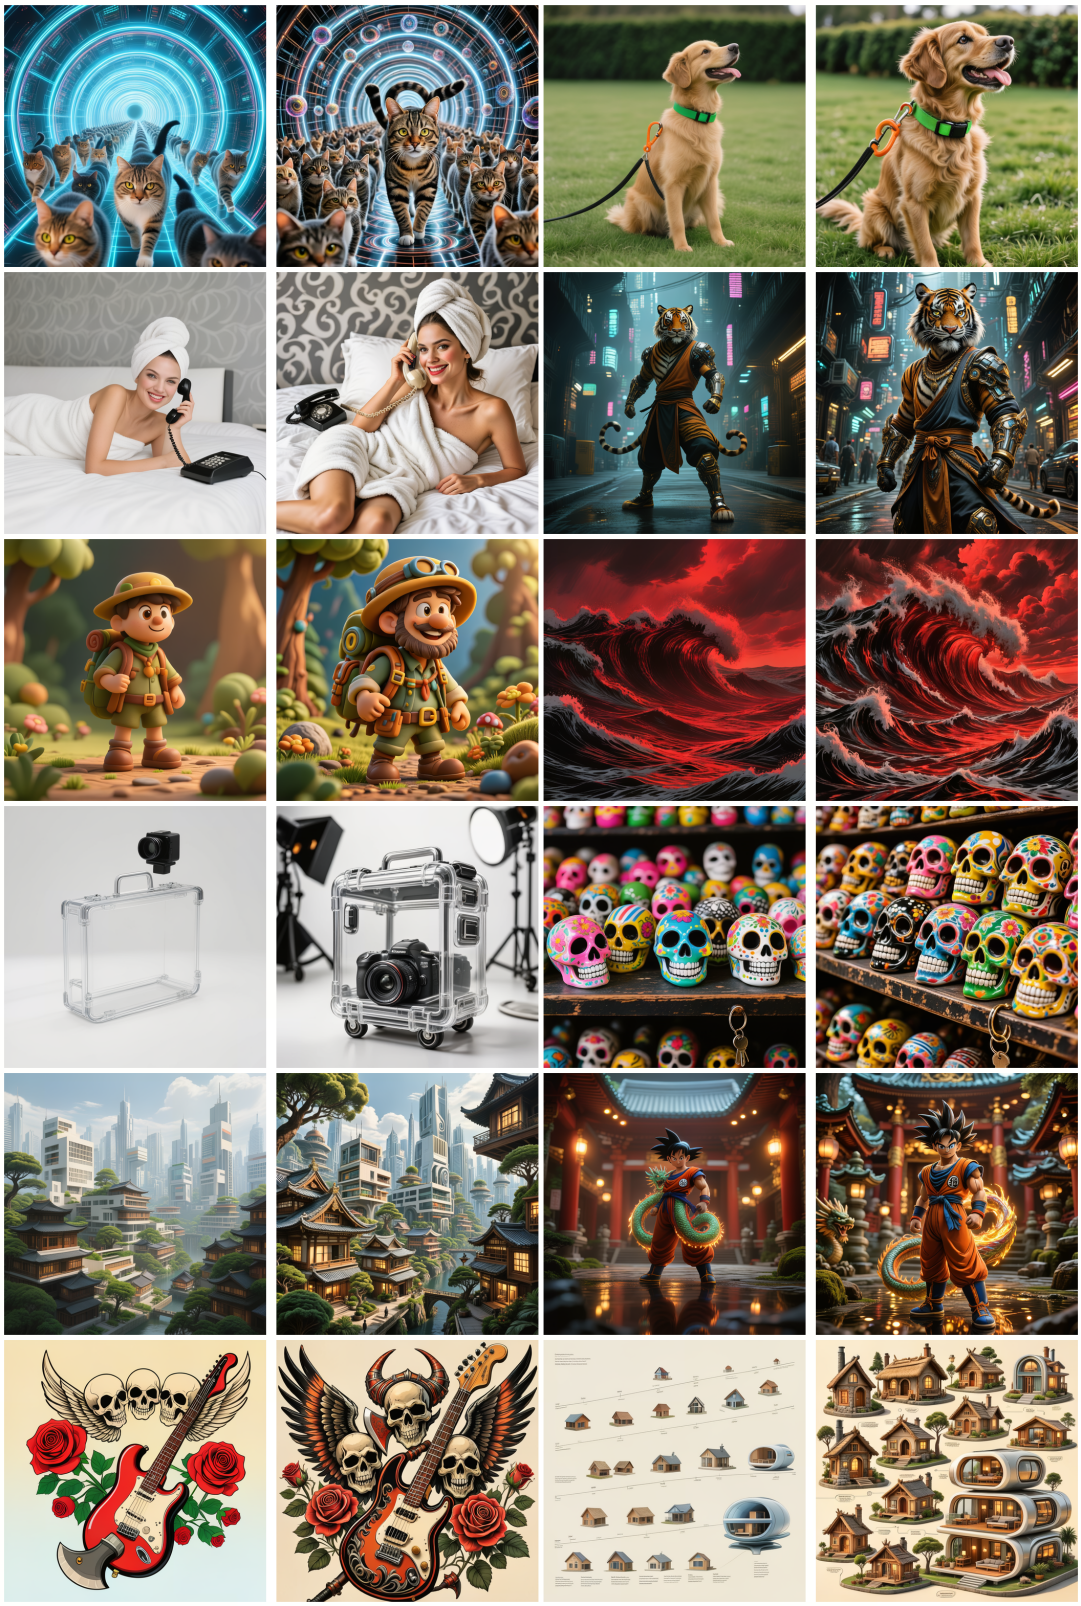

Figure 9. Random visual samples from the HPSv3 evaluation set by class. Each class includes two pairs; within each pair, the left image is before RL and the right image is after RL.

Figure 10. Random visual samples from the HPSv3 evaluation set by class. Each class includes two pairs; within each pair, the left image is before RL and the right image is after RL.

Table 9. Detailed prompts used in Fig. 5.

ID	Prompt Content
1	Six illuminated letters ('A', 'B', 'C', 'N', 'O', 'Y') in two rows on a dark blue background, outlined with white bulbs glowing blue.
2	The image is a shield-shaped graphic with a black border. Inside the shield, there are three horizontal stripes: red at the top, white in the middle, and blue at the bottom. The red stripe contains white text that reads, "Wouldn't you rather..." The white stripe in the middle has bold black text that says, "VOTE." And the blue stripe at the bottom features white text that reads, "By Mail." The overall design suggests a campaign or initiative encouraging people to vote by mail.
3	The image features two small, round, metallic tin containers placed on a dark, textured fabric background. The lids of the tins are slightly open, revealing the contents inside. The tins have a vintage or antique appearance, with one lid displaying an engraved design. The text "M-BOX 2.0" is prominently overlaid on the image in a bold, yellow font, while "by: Jimmy-Fan" is written in a smaller, yellow font below it. The overall composition suggests a promotional or product image for the tins, emphasizing their design and branding.
4	The image depicts a serene nighttime camping scene in a forest. The atmosphere is illuminated by a soft, glowing light emanating from within a tent, casting a warm and inviting glow on the surrounding area. The tents are set up on a forest floor covered with pine needles and leaves. The trees are tall and dense, creating a canopy that stretches out into the night sky. The overall mood of the image is tranquil and adventurous, capturing the essence of an outdoor camping experience. In the upper right corner, there is text in Spanish that reads "DESCUBRE LO QUE ESCONDEN NUESTRO CAMPING," and in the bottom right corner, the words "UPBAN KIDS" are displayed, with "UPBAN" in white and "KIDS" in green.
5	The image shows a cardboard sign with handwritten text in French. The text reads "L'ARANCAEY FIMMNA," with "L'ARANCAEY" written in purple and "FIMMNA" in black. The sign appears to be placed against a wall, and there are some items and containers visible in the background.
6	The image features a cartoon-style illustration of a person standing and holding a fishing rod. The person appears to be a stylized representation of a man with light skin, wearing a white shirt, a red tie, and dark pants. The fishing rod is angled upwards, and at the end of the line, there is a white fish with an upward-pointing arrow, suggesting the fish has been caught. Above the person, there is a red banner with the text "FISHING CHALLENGE" in white letters. To the right of the person, the word "GOTCHA!" is written in white text, indicating the successful catch of the fish. The background is a solid light green color.
7	The image features four packets of snacks against an orange background. Each packet has a distinct design and is labeled with different types of snacks. The first packet on the left has a black and orange swirl pattern and is labeled "Apps Way Crisps." The second packet has a yellow and black crisscross pattern and is labeled "Apps Wheat Crisps." The third packet has black diagonal stripes and is labeled "Apps Potato Crisps." The fourth packet has a black polka dot pattern and is labeled "Apps Cheese Balls." The design elements and text on each packet are consistent, suggesting they are part of a branded product line.
8	The image shows a piece of paper with handwritten notes on it, placed on a wooden surface. The text writes: 'VAR RL Done Right'. The handwriting is informal and includes some slang or abbreviations. The paper seems to be printed with a grid pattern, suggesting it might be a printed form or a piece of paper with a pre-defined layout.

Table 10. Detailed prompts used in Fig. 6.

ID	Prompt Content
1	A cheerful, anthropomorphic squirrel stands upright against a solid light green background. The squirrel has a rich brown coat with a lighter cream-colored underbelly, muzzle, and inner ears, which are pink at the tips. Its large, expressive black eyes are wide and friendly, with a small, triangular nose and a wide smile showing two prominent front teeth. The squirrel's bushy tail is curled in a spiral at the end, with a gradient from dark brown at the base to a lighter cream color at the tip. Around its neck, it wears a vibrant flower lei composed of small, colorful blossoms in red, yellow, green, purple, and blue. The squirrel holds a large, smooth, light brown egg in both front paws, which are positioned close to its chest. Its posture is upright, with its hind legs planted firmly on the ground, and it casts a soft shadow beneath it on the green background.
2	An anime-style illustration depicts a muscular, metallic tiger with sharp, angular features, standing on a rooftop. The tiger is in a dynamic pose, gripping a sleek, red electric guitar, and its mouth is open wide as if caught in the midst of a powerful roar or song. Above the tiger, a bright spotlight casts a dramatic beam of light, illuminating the scene and creating stark shadows on the surrounding rooftop features.
3	A surreal figure appears to be sculpted from intertwining tendrils of gray smoke and whirling flurries of snow, giving the impression of a man caught in a blizzard. In one hand, this ethereal being holds what looks to be a gateway to the cosmos, depicted in a photorealistic manner with vibrant nebulae and star clusters visible within its confines. The entire scene is a highly detailed octane render, showcasing sharp contrasts and the interplay of light and shadow that imbues the image with a sens.
4	An ornate royal carriage, painted in deep red with golden trim, stands prominently against a landscape blanketed in pristine snow. Behind it, the silhouettes of tall pine trees dusted with white can be discerned through the soft haze of a winter's day. In front of the carriage, the snow-covered ground glistens under the subtle light of the afternoon sun.
5	a detailed 17th-century Dutch Baroque painting depicting a chestnut horse standing amidst a vibrant field of tulips and daisies. The horse's mane and tail are elegantly captured by the artist, flowing with the gentle breeze that seems to animate the scene. In the background, a traditional windmill sits under a partly cloudy sky, completing this pastoral landscape.
6	A photograph with a standard lens style. The subject is a man standing against a background that appears to be covered in splattered paint. The man has short, dark hair and a light beard. He is wearing a long-sleeved, button-up shirt that is also splattered with paint, suggesting he might have been involved in an artistic or creative activity. The shirt is a light color, possibly beige or off-white, and has two chest pockets. The man's expression is serious and contemplative. The background is primarily white with red and black paint splatters scattered across it. The lighting in the photograph is soft and even, highlighting the details of the man's face and the texture of his shirt. The overall mood of the image is artistic and somewhat moody, with a focus on the subject's expression and the abstract splatters in the background.
7	Two young girls stand on a lush, green grassy field, both dressed in white lace sleeveless dresses adorned with a purple satin ribbon tied in a bow at the waist. The girl in the foreground has shoulder-length, wavy blonde hair and wears a floral crown featuring white and purple flowers, including roses and possibly lisianthus, with small greenery accents. She holds a small bouquet composed of white flowers, purple blooms, and clusters of dark red berries, with some green foliage. Behind her, the second girl, also with blonde hair (slightly lighter and more tousled), wears an identical floral crown and dress. She holds a bouquet of vibrant pink flowers, likely carnations, with green stems. Both girls are smiling brightly, their eyes a striking blue, and the wind gently lifts strands of their hair. The background is a soft, out-of-focus expanse of green grass, dotted with tiny white flowers, creating a serene, natural setting.
8	a breathtaking photograph capturing the vibrant hues of a sunset with streaks of pink and orange painting the sky behind the majestic Grand Canyon. The canyon's intricate rock formations are silhouetted against the illuminated backdrop, showcasing the deep crevices and towering spires. In the foreground, the Colorado River can be glimpsed winding its way through the ancient geological marvel.