

VAR-Q: Tuning-free KV Cache Quantization for Visual Autoregressive Image and Video Generation

Boxun Xu^{1,†,*} Jiaji Lu^{1,†} Yuxuan Yin¹ Zihu Wang¹
Ziyue Liu¹ Yu Wang¹ Zirui Liu² Peng Li¹

¹University of California, Santa Barbara ²University of Minnesota

[†]Equal contribution.

Abstract

Visual autoregressive generation spans two dominant paradigms: next-scale prediction for images and video, and next-frame prediction for video. Both decode block-by-block rather than token-by-token, yielding KV caches whose statistics diverge from those of language models. At scale, KV memory exceeds the weights and limits large-scale deployment caused by a large amount of visual tokens. Our analysis of these caches reveals three findings. First, a small set of *massive-value channels* carries the functional load: ablating these outlier channels collapses generation, whereas ablating a matched set of inlier channels leaves quality intact. Second, the outliers organize along two axes, feature and block, with magnitudes persistent within each cell and shifting across cells. Third, this structure arises from two independent mechanisms: projection bias and RoPE-induced scattering. Guided by these findings, we propose VAR-Q, a tuning-free KV-cache quantization framework that groups along the same two axes. The generation loop provides block boundaries for free, removing the need for outlier detection or calibration. Paired with bit-packing kernels for runtime memory reduction, VAR-Q delivers low-bit quantization across nine models spanning both paradigms, cutting KV memory by up to 78.5% with negligible quality degradation and enabling $6\times$ larger inference batches.

1 Introduction

Visual generation has consolidated around two autoregressive paradigms. *Next-scale prediction* decodes a multi-scale token map one resolution at a time, and now spans both image [4, 24] and video [12] synthesis. *Next-frame prediction* decodes a video one frame at a time, conditioned on the past frames [7, 29]. Both decode in semantically coherent *blocks*, a scale or a frame, rather than individual tokens. As representation dimensions increase and generation horizons reach hundreds of frames, the Key-Value (KV) cache accumulated across these blocks rivals or exceeds the model weights themselves, easily saturating single-GPU memory. KV-cache quantization has become the standard remedy for language models [2, 5, 9, 11, 13, 14, 21, 30, 32], where per-channel grouping for keys and per-token grouping for values now compress caches to 2–4 bits with negligible perplexity loss. None of this work targets visual AR. The cache structure of these models is not understood, and recipes calibrated to language models do not transfer: applied directly, they collapse generation.

The KV cache of a visual AR model is not a uniform memory of past tokens. Less than 1% of its outlier channels carry its functional load in a numerical manner. In LongLive [29], zeroing just the top 0.5% of channels by magnitude collapses generation entirely, while zeroing ten times as many

*Correspondence: boxunxu@ucsb.edu




Figure 1: **Generation quality under quantized KV caches.** Top: On Infinity-8B at INT2, VAR-Q reproduces the BF16 rendering of text and scene layout that FlexGen and KIVI corrupt. On Self-Forcing at INT4, VAR-Q tracks the camera trajectory properly.

inlier channels leaves generation quality largely intact. The same asymmetry holds across Infinity [4], InfinityStar [12], and SelfForcing [7], in both image and video synthesis, and at every layer we measure. Whatever the cache stores for future retrieval lives in a thin, magnitude-distinguished subspace.

This subspace is not a scattered set of outliers but a structured one, organized along two orthogonal axes. A small set of channels carries large magnitudes throughout decoding, producing persistent stripes along the channel axis. Within each block magnitudes are stationary, and across block boundaries they shift sharply, producing bands along the block axis. Both axes have architectural causes. On the channel axis, key outliers arise from a composition of three multiplicative sources: heavy-tailed $\mathbf{W}_K$ row norms seed an outlier baseline, post-projection RMSNorm gain γ_K amplifies it where present, and RoPE pair rotation [19] redistributes energy at the pair level. On the block axis, position-varying RoPE projects channel magnitudes through a sinusoidal modulation, especially at high-frequency channels in the block index, producing the cross-block shifts directly from the rotation schedule rather than from any change in input. Value caches exhibit the same two-axis structure with weaker per-axis attribution, sufficient to justify the same grouping rule without a separate mechanism.

However, existing quantization rules align with at most one axis: KIVI [14] groups keys per-channel and values per-token, leaving the block axis unprotected on both. To this end, we propose **VAR-Q**, a tuning-free KV quantization framework that groups the cache jointly along channels and blocks, with one quantization unit per channel and block pair. The block boundaries come for free from the generation loop itself, without detection, calibration, or retraining. We pair the algorithm with bit-packing kernels that turn nominal bit-width reduction into actual memory savings on current GPUs. Across nine visual AR models spanning both paradigms, VAR-Q matches the BF16 baseline at INT4 on both next-scale image and next-scale video generation, and remains usable on Infinity-8B down to INT2. On next-frame video, it sustains quality at INT4 and outperforms KV quantization work transferring from LLM domains directly such as KIVI [14] and FlexGen [18]. Complementary to concurrent cache eviction [10, 28] and attention sparsification [26] for visual generation, VAR-Q reduces KV memory by up to 78.5% relative to BF16, and inference batch size grows by up to $6\times$ on a single NVIDIA A100 80GB.

2 Background

We unify the two visual autoregressive paradigms under a common block-wise notation. *Next-scale AR* [4, 12, 24] emits one resolution of a multi-scale token map per step; *next-frame AR* [7, 29] emits one or more frames per step, and we treat each frame as a separate block. We refer to the tokens




Figure 3: Channel-wise outlier, random, and inlier KV ablation on four visual AR models. Channels are ranked by maximum absolute value over cached tokens.

emitted at step i as block i . After B steps, the key cache is

$$\mathbf{K} = \{\mathbf{K}^{(i)}\}_{i=1}^B, \quad \mathbf{K}^{(i)} \in \mathbb{R}^{n_i \times D}, \quad (1)$$

where n_i is the number of tokens in block i and D is the head-channel dimension; the value cache $\mathbf{V}$ is defined analogously. In next-scale AR, n_i grows geometrically across a fixed number of scales; in next-frame AR, n_i is constant and B grows with the video length. Both caches are produced by linear projections $\mathbf{K}^{(i)} = f_K(\mathbf{x}^{(i)} \mathbf{W}_K)$ and $\mathbf{V}^{(i)} = \mathbf{x}^{(i)} \mathbf{W}_V$, where f_K denotes any model-specific post-projection transforms applied to keys.

3 Methodology

The element-wise magnitudes of the KV cache form a two-axis grid over feature channels and decoding blocks, stable within cells and shifting sharply across block boundaries. A small subset of cells carries values orders of magnitude larger than the rest, and these dominate generation. Group-wise quantization must therefore align with both axes. Section 3.1 establishes that the massive values are critical for generation through controlled ablation; Section 3.2 characterizes the two-axis grid; Section 3.3 attributes both axes to specific architectural sources; and Section 3.4 translates the grid into VAR-Q.

3.1 Massive Values in KV Caches

Outliers in the KV cache are functional, not just numerical. Massive activations have been documented in the residual streams [22] and MLP weights [31] of language models, where direct ablations confirm them as functionally indispensable. KV cache outliers have been observed across LLMs [5, 14] and VLMs [20], motivating protect-by-isolation strategies such as per-channel grouping [14], rotation-based redistribution [1], and high-precision retention for dense outliers [5]. This protection establishes utility, not necessity: per-channel grouping and rotation succeed whether outliers carry load-bearing signal, dominate the quantization range of co-grouped channels, or both. Only direct ablation can adjudicate. Visual AR has had neither: no characterization of its KV cache outliers, no ablation of their function. Whether load-bearing outliers exist in this regime, and how they are organized across long visual sequences, remains open.

We test necessity directly through ablation. For each layer, we rank channels by their maximum magnitude across the cached sequence. The top p percent form the *outlier* set. From the remaining




Figure 2: Generation Comparison under channel-wise outlier vs. inlier KV masking.




Figure 4: Median absolute values of K (top) and V (bottom) caches across five visual AR models over the (channel, generation index) grid. Channel stripes persist along the generation axis in all models, and their magnitudes shift sharply at block boundaries. Color scales are per-panel.

channels, we draw an equal-sized random sample to form the *inlier* set. Each ablation zeros one set. Figure 2 makes the contrast visible: on Infinity-2B at $p = 5\%$ and LongLive-1.3B at $p = 0.5\%$, outlier ablation collapses generation into uniform texture or solid blocks; inlier ablation leaves outputs intact. Figure 3 extends this to four models and eight metrics. On LongLive-1.3B, masking the top 0.5% drops CLIP from 0.32 to 0.28, while ablating ten times as many inlier channels leaves CLIP at baseline (0.32). We call this small, magnitude-defined subset the *massive-value channels*.²

3.2 Massive Values are Structured Along Two Axes in KV Caches

Massive values impose a structural constraint on group-wise quantization: any group that mixes massive and regular cells must stretch its scaling factor to cover the massive cell, starving the regular cells of resolution. Resolving this requires group boundaries that follow the spatial distribution of massive values. In visual AR KV caches, that distribution is two-dimensional: massive values concentrate along a small set of channels and a small set of blocks. Both axes must be covered.

Channel stripes and block bands. The same structure appears in every model in Figure 4, organized along two orthogonal axes. A small subset of channels carries large magnitudes throughout the entire decoding trajectory, producing persistent stripes along the channel axis. Within each block, magnitudes are approximately stationary; across block boundaries, they shift sharply, producing bands along the block axis. The two axes are not redundant; they compose. Massive values live precisely at the intersections of high-magnitude channels and high-magnitude blocks. The block axis is specific to visual AR: language-AR KV caches form a homogeneous token stream with no analogous boundaries [5, 14], and prior LLM quantization methods accordingly exploit only the channel axis.

Table 1: INT4 Quantization Errors (%) on Infinity-2B, where $\epsilon_X = \|X - X'\|_F / \|X\|_F$.

Rule	Reconstruction		Attention	
	ϵ_K	ϵ_V	ϵ_A	ϵ_O
Per-tensor	24.01	43.54	29.20	32.32
Per-block	21.59	32.93	26.36	22.78
Per-channel	12.11	18.07	16.17	13.13
Per-(ch., blk.)	10.04	15.28	13.38	10.13

Aligning with both axes minimizes error. Table 1 measures how well four grouping rules contain the structure on Infinity-2B. Per-tensor collapses the entire cache into one group and incurs the largest error on every metric. Per-block aligns with the generation axis but lets each block mix high- and low-magnitude channels. Per-channel aligns with the hidden axis, the choice of KIVI [14] for LLM key caches, but forces a massive cell to share its scaling factor with regular cells of the same channel across all blocks. Only per-(channel, block) aligns with both axes, placing each massive cell in a group that contains only other massive cells from the same channel and the same block. Errors fall monotonically as more axes are covered, and the joint rule wins on every column: ϵ_O from 32.32% to 10.13%, ϵ_K from 24.01% to 10.04%.

²A per-tensor decomposition of the ablation reveals that video models are more sensitive to **K** outliers and image models to **V** outliers (Appendix A).

3.3 Why Do Massive Values Emerge in KV Caches of Visual Autoregressive Models?




Figure 5: Outlier persistence across different prompts

While Section 3.2 establishes *where* massive outliers concentrate, this section investigates *why* they emerge. We attribute these structural anomalies along the channel and block axes to specific operations within the K and V projection chains. Formally, we decompose the generation of K_t and V_t as follows:

$$K_t = \text{RoPE}_t(\text{RMSNorm}(x_t W_K^\top) \odot \gamma_K), \quad V_t = x_t W_V^\top. \quad (2)$$

K outliers are architectural and V outliers are largely input-driven. Whether an outlier is architectural or input-driven can be tested empirically through *cross-prompt persistence*: an architectural outlier should persist across prompts, while an input-driven outlier should drift. We sample N prompts across models and define outliers as the top- p channels by magnitude. **Persistence@N** is the fraction of a single prompt’s outlier channels that remain outliers under *all* N prompts (random ≈ 0). At $N=10$, $p=1\%$, K Persistence@10 reaches 0.53–0.62, a hard architectural core shared across inputs; V sits at 0.07–0.14, a $3.8 \times$ – $8.5 \times$ gap (Fig. 5). V’s near-vanishing persistence indicates input activations dominate V outlier identity.

A Compounding Architectural Effect on K-Cache Channel Outliers.

The architectural core of K identified above arises from three multiplicative sources along the projection chain (Eq. (2)): (i) heavy-tailed row norms of W_K seed an initial outlier baseline; (ii) where RMSNorm-K is present, the channel gain γ_K amplifies this baseline; and (iii) RoPE rotates each $(2i, 2i+1)$ pair, redistributing energy between the two channels while preserving the pair’s ℓ_2 norm.




Figure 6: Recall of K outliers as projection-chain components are progressively added; W_V shown as control.

We quantify this compounding by progressively adding components and measuring the recall of final K outliers (Fig. 6). At $p=1\%$, W_K alone recovers 35–42% (35–42 $\times$ random); adding γ_K lifts recall further; the full chain reaches 55–63 $\times$ random. In contrast, W_V alone yields 5–6 $\times$ random recall on every model, an order of magnitude weaker than the K chain.

Block-Axis Dynamics generated by RoPE Periodicity. A class of K block-axis outliers is generated by RoPE alone, independent of input dynamics. RoPE rotates each channel pair $(2i, 2i+1)$ at frequency ω_i by an angle $(b\omega_i)$ along the pair’s RoPE axis, where b is the block index. Within block b , every token shares the same rotation, and per-token magnitudes are

$$|K_{2i}(b, t)| = M_t |\cos(b\omega_i + \phi_t)|, \quad |K_{2i+1}(b, t)| = M_t |\sin(b\omega_i + \phi_t)|, \quad (3)$$

with M_t the pre-RoPE pair amplitude at token t and ϕ_t its initial angle. Let $c_{2i}(b)$ denote the channel’s per-token magnitudes within block b , with $\|c_{2i}(b)\|_\infty := \max_t |K_{2i}(b, t)|$. Two consequences follow and are visible in Fig. 7. (i) *Pair envelope is approximately block-invariant.* RoPE preserves the per-token pair ℓ_2 exactly, and in the massive-value regime this carries over to the pair of channel maxes (Appendix C). The two channels move in anti-phase within each block. (ii) *Fast pairs dominate the block axis.* The total angle a pair traverses across the block sequence determines its swing: slow pairs barely rotate and stay flat, fast pairs complete multiple cycles and produce block bands.




Figure 7: (a) On Self-Forcing, block-axis outliers concentrate on the lowest-index (fastest-rotating) pairs. (b) Within such a pair, $|c_0|$ and $|c_1|$ swap energy across blocks while pair ℓ_2 stays approximately flat.

3.4 VAR-Q: Algorithm and System Support

Algorithm. Section 3.2 shows that massive values in the KV cache concentrate along two axes: channels and blocks. VAR-Q aligns the quantization grouping with both. As shown in Fig. 8, on top of the per-channel grouping, VAR-Q splits each group at block boundaries. The resulting quantization unit is indexed by (i, j) , where i is the block index and j is the channel; each unit carries a single scaling factor under uniform symmetric quantization.

The rule is online without requiring auxiliary detection or calibration and retraining. It is paradigm-agnostic, applying identically whether a block is a scale or a frame. And it is unified across keys and values, since both caches share the same two-axis structure.

We further expose a single hyperparameter $r \in (0, 1]$, the sub-block ratio. Let $L_{\max}$ denote the largest block size in the sequence. Each block of size L is partitioned along the token axis into $\lceil L/(rL_{\max}) \rceil$ equal-sized sub-blocks; smaller r yields finer groups. In next-scale AR earlier scales are geometrically smaller than $L_{\max}$ and remain unsplit at typical r , so r effectively controls the granularity of the last scale; in next-frame AR every block has size $L_{\max}$, so r partitions every frame uniformly. We default to $r = 1$ throughout, recovering the per-block rule, and ablate r in Section 4.

System Support. Quantization alone does not reduce the KV-cache footprint, because current CUDA toolchains lack native sub-8-bit storage types. We pair the algorithm with a hardware-friendly Triton implementation that packs sub-byte quantized values into byte-aligned containers on write and unpacks them during attention; details in Appendix E.




Figure 8: The proposed VAR-Q pipeline.

4 Experiments

We evaluate VAR-Q on nine models across five architectures, covering both paradigms and modalities: Infinity-2B/8B [4] for next-scale image generation, InfinityStar-8B [12] at 720p and 480p for next-scale video, and Self-Forcing [7] on 5-second video clips and LongLive [29] on 30-second video clips for next-frame video generation. We additionally include VAR-d20/d24/d30 [24] on class-conditional ImageNet-1K. We compare against KIVI [14] and FlexGen [18]. KIVI uses per-channel grouping for K and per-token grouping for V; FlexGen uses 16×16 tile-wise grouping; VAR-Q uses per-channel and per-block grouping with $r = 1$. All methods share identical bit-packing kernels, and each method’s scale-factor overhead is included in the reported KV. For class-conditional image generation we report FID, IS, Precision, and Recall, together with PSNR, SSIM, and LPIPS against the BF16 output under the same seeds, isolating quantization fidelity from sampling stochasticity. For text-to-image we report GenEval [3] and DPG [6]; for text-to-video we report VBench [8] and CLIP Score on MovieGen prompts [17], and additionally report VLM-as-judge [33] scores using Claude Opus 4.7 on long-video clips. The tuning-free methods require no calibration data, and INT8 packed kernels are implemented in Triton and evaluated on NVIDIA A100 80 GB.

4.1 Next-Scale Image Generation

Table 2 reports next-scale image generation on VAR (class-conditional) and Infinity (text-to-image). The qualitative examples are shown in Appendix G.

On VAR-d20/d24/d30, VAR-Q at INT4 matches BF16 in FID across all three depths and leads KIVI and FlexGen on PSNR, SSIM, and LPIPS computed against matched-seed BF16 outputs. FID is noisy at this scale; the consistent PSNR/SSIM/LPIPS advantage shows that VAR-Q tracks the BF16 generation trajectory at the sample level rather than coincidentally landing on a similar distribution. The gap sharpens at INT3: VAR-Q stays within 1.0 FID of BF16 (3.77/3.60/3.05) while KIVI degrades by 1.2 to 2.3 FID and FlexGen collapses (FID 15.3/14.9/9.5).

On Infinity-2B and Infinity-8B, VAR-Q preserves GenEval and DPG Overall within 0.3 points of BF16 at INT4. The 8B variant additionally holds at INT3 (GenEval 83.28, $\Delta = 0.26$) while KIVI

Table 2: Next-scale image generation quality under low-precision KV-cache quantization on VAR-d20/d24/d30 and Infinity-2B/8B. Our INT4 results closely match the BF16 baseline. Bold/underline: best/second-best per precision. KV (GB): realized cache size; subscripts: savings vs. BF16.

(a) VAR-d20/d24/d30 [24] on class-conditional ImageNet 256×256.

Model	#Params	Method	Prec.	Generation Quality				Fidelity to BF16			KV (GB)↓
				FID ↓	IS ↑	Precision ↑	Recall ↑	PSNR ↑	SSIM ↑	LPIPS ↓	
VAR-d20	600M	Baseline	BF16	4.11	357.32	0.8650	0.5160	–	–	–	16.60 _{bs=128}
		FlexGen	INT4	4.03	325.39	0.8326	0.5230	<u>11.84</u>	<u>0.264</u>	0.553	4.29 _{-74.14%}
		KIVI	INT4	<u>3.55</u>	<u>341.52</u>	<u>0.8458</u>	0.5363	11.81	0.256	<u>0.552</u>	4.65 _{-71.97%}
		VAR-Q	INT4	3.44	342.77	0.8502	<u>0.5281</u>	13.22	0.308	0.462	<u>4.64</u> _{-72.06%}
		FlexGen	INT3	15.33	139.59	0.5944	0.6308	<u>10.35</u>	<u>0.212</u>	0.650	3.46 _{-79.14%}
		KIVI	INT3	<u>6.36</u>	<u>233.84</u>	<u>0.7260</u>	0.5748	9.88	0.187	0.662	3.82 _{-76.97%}
		VAR-Q	INT3	3.77	262.24	0.7706	<u>0.5767</u>	12.24	0.266	0.538	<u>3.81</u> _{-77.06%}
VAR-d24	1.0B	Baseline	BF16	2.97	357.24	0.8444	0.5605	–	–	–	23.91 _{bs=128}
		FlexGen	INT4	2.74	325.69	0.8175	0.5816	11.78	0.256	0.560	6.18 _{-74.14%}
		KIVI	INT4	<u>2.56</u>	<u>340.55</u>	<u>0.8295</u>	<u>0.5786</u>	<u>11.92</u>	<u>0.268</u>	<u>0.545</u>	6.70 _{-71.97%}
		VAR-Q	INT4	2.52	341.60	0.8321	0.5718	13.03	0.296	0.479	<u>6.68</u> _{-72.06%}
		FlexGen	INT3	14.91	144.09	0.6066	0.6424	10.63	0.201	0.643	4.99 _{-79.14%}
		KIVI	INT3	<u>4.29</u>	<u>239.40</u>	<u>0.7271</u>	<u>0.6340</u>	<u>10.88</u>	<u>0.226</u>	<u>0.631</u>	5.51 _{-76.97%}
		VAR-Q	INT3	3.60	258.33	0.7477	0.6130	11.90	0.252	0.558	<u>5.48</u> _{-77.06%}
VAR-d30	2.0B	Baseline	BF16	2.07	335.10	0.8260	0.5812	–	–	–	37.35 _{bs=128}
		FlexGen	INT4	2.00	307.92	0.8061	0.6006	12.47	0.277	0.510	9.66 _{-74.14%}
		KIVI	INT4	<u>1.99</u>	<u>321.46</u>	0.8162	<u>0.5974</u>	<u>12.86</u>	<u>0.290</u>	<u>0.484</u>	10.47 _{-71.97%}
		VAR-Q	INT4	1.97	323.67	<u>0.8156</u>	0.5961	13.54	0.312	0.437	<u>10.44</u> _{-72.06%}
		FlexGen	INT3	9.52	180.57	0.6738	0.6417	11.33	0.232	0.597	7.79 _{-79.14%}
		KIVI	INT3	<u>3.26</u>	<u>264.93</u>	<u>0.7536</u>	0.6189	<u>11.69</u>	<u>0.249</u>	<u>0.573</u>	8.60 _{-76.97%}
		VAR-Q	INT3	3.05	268.16	0.7612	<u>0.6307</u>	12.41	0.270	0.517	<u>8.57</u> _{-77.06%}

(b) Infinity-2B/8B [4] on text-to-image generation.

Model	Method	Prec.	GenEval [3] ↑				DPG [6] ↑						KV (GB)↓
			Two-Object	Position	Color	Overall	Entity	Relation	Attribute	Global	Other	Overall	
Infinity-2B	Baseline	BF16	80.61	25.51	88.03	69.61	86.55	89.51	86.62	83.24	90.14	81.97	41.10 _{bs=8}
	VAR-Q	INT8	79.34	26.26	87.23	69.06	87.69	89.39	88.70	86.40	85.85	82.21	20.65 _{-49.75%}
	FlexGen	INT4	78.57	<u>24.75</u>	84.50	<u>67.75</u>	<u>87.39</u>	91.56	85.38	82.24	<u>82.04</u>	<u>81.01</u>	<u>10.60</u> _{-74.22%}
	KIVI	INT4	80.61	24.49	87.23	67.48	88.18	<u>91.13</u>	<u>86.51</u>	<u>83.18</u>	80.81	80.04	10.78 _{-73.76%}
	VAR-Q	INT4	<u>79.34</u>	26.01	<u>86.97</u>	69.35	87.24	89.60	87.43	86.75	90.13	82.02	10.38 _{-74.75%}
	KIVI	INT3	<u>77.30</u>	26.26	<u>86.70</u>	<u>67.56</u>	<u>87.12</u>	92.39	<u>85.66</u>	<u>82.94</u>	<u>79.31</u>	<u>80.76</u>	<u>8.73</u> _{-78.76%}
	VAR-Q	INT3	78.32	<u>24.24</u>	88.03	68.04	87.74	<u>89.92</u>	87.01	85.13	87.01	81.18	8.32 _{-79.75%}
Infinity-8B	Baseline	BF16	95.41	62.12	90.43	83.54	88.27	89.30	91.57	92.83	88.61	85.72	67.43 _{bs=6}
	VAR-Q	INT8	96.43	62.88	89.89	84.00	90.57	92.57	89.23	87.12	92.68	85.97	33.88 _{-49.75%}
	FlexGen	INT4	95.41	60.89	89.02	82.45	<u>90.75</u>	<u>92.33</u>	<u>89.15</u>	85.04	<u>84.49</u>	84.58	<u>17.38</u> _{-74.22%}
	KIVI	INT4	<u>95.66</u>	65.40	<u>91.22</u>	<u>83.26</u>	91.13	94.10	88.93	<u>85.67</u>	82.04	<u>85.74</u>	17.69 _{-73.76%}
	VAR-Q	INT4	95.92	<u>63.64</u>	91.49	83.67	90.70	91.99	91.07	87.47	89.23	85.90	17.03 _{-74.75%}
	KIVI	INT3	95.66	63.38	91.49	<u>82.75</u>	<u>87.74</u>	94.33	<u>88.97</u>	<u>84.05</u>	<u>85.71</u>	<u>84.89</u>	<u>14.32</u> _{-78.76%}
	VAR-Q	INT3	<u>95.41</u>	<u>63.13</u>	<u>90.69</u>	83.28	89.65	<u>88.84</u>	90.82	89.16	91.80	85.89	13.65 _{-79.75%}
	VAR-Q	INT2	85.97	56.57	86.70	74.38	87.80	90.09	89.31	91.22	88.12	83.83	8.60 _{-87.25%}

loses 0.79 points, and remains usable at INT2 (GenEval 74.38, DPG 83.83), confirming that larger backbones tolerate more aggressive KV-cache quantization.

4.2 Next-Scale and Next-Frame Video Generation

We evaluate VAR-Q on next-scale video with InfinityStar and on next-frame video with Self-Forcing and LongLive; all three backbones use RoPE. The qualitative examples are shown in Appendix G.

InfinityStar. At INT4 on both 720p and 480p, VAR-Q matches BF16 in VBench Total Quality (81.66 and 81.32 vs. 82.33 and 81.30) and ranks first on Total at every precision and resolution. The gap widens at INT3 on the higher-resolution 720p setting: KIVI loses 6.87 points to 75.46, while VAR-Q stays within 1.74 points of BF16 at 80.59. KV memory drops by 73 to 79% across precisions.

Table 3: **Next-scale video generation on InfinityStar**. KV reports cache size (GB) with reduction vs. BF16 in subscript. **Bold** and underline indicate best and second-best, respectively.

Model	Method	Prec.	VBench [8] ↑								KV (GB)↓
			Subject Consistency	Background Consistency	Temporal Flickering	Motion Smoothness	Dynamic Degree	Aesthetic Quality	Imaging Quality	Total Quality	
InfinityStar 8B 720p	Baseline	BF16	93.69	94.62	97.71	98.86	67.50	58.49	66.79	82.33	62.72 _{bs=3}
	FlexGen	INT4	93.83	94.76	97.83	98.91	63.00	55.50	64.80	80.11	<u>16.48</u> _{-73.72%}
	KIVI	INT4	92.01	<u>94.61</u>	97.27	<u>98.70</u>	<u>68.00</u>	<u>56.12</u>	63.60	<u>81.24</u>	16.73 _{-73.33%}
	VAR-Q	INT4	<u>92.65</u>	94.09	<u>97.37</u>	98.59	74.50	56.87	<u>64.35</u>	81.66	16.00 _{-74.49%}
	FlexGen	INT3	90.12	92.68	97.84	<u>96.91</u>	62.00	56.01	60.48	79.95	<u>13.56</u> _{-78.37%}
	KIVI	INT3	<u>90.45</u>	90.42	<u>97.65</u>	96.79	74.00	50.97	61.92	75.46	13.81 _{-77.99%}
	VAR-Q	INT3	91.41	93.56	<u>97.25</u>	98.37	75.00	<u>55.55</u>	<u>61.75</u>	80.59	13.08 _{-79.15%}
InfinityStar 8B 480p	Baseline	BF16	93.01	94.53	97.25	98.68	66.00	58.13	64.01	81.30	55.87 _{bs=6}
	FlexGen	INT4	90.59	94.39	95.16	98.64	66.50	56.48	62.14	78.94	<u>15.03</u> _{-73.10%}
	KIVI	INT4	<u>92.44</u>	<u>94.42</u>	<u>96.65</u>	97.79	64.00	<u>56.97</u>	<u>62.92</u>	<u>80.46</u>	15.25 _{-72.71%}
	VAR-Q	INT4	92.75	94.51	97.21	98.67	68.50	57.84	63.62	81.32	14.60 _{-73.86%}
	FlexGen	INT3	91.89	93.91	97.24	98.62	<u>64.00</u>	<u>57.54</u>	63.69	<u>80.64</u>	<u>12.45</u> _{-77.71%}
	KIVI	INT3	93.01	94.46	97.33	98.71	64.00	57.79	63.56	78.05	12.67 _{-77.33%}
	VAR-Q	INT3	<u>92.86</u>	<u>94.30</u>	<u>97.26</u>	<u>98.64</u>	65.50	57.52	63.37	80.98	12.03 _{-78.47%}

Table 4: **Next-frame video generation on Self-Forcing (5 s) and LongLive (30 s)**. KV reports cache size (GB) with reduction vs. BF16 in subscript. On generation quality, **bold** and underline indicate best and second-best, respectively.

Model	Method	Prec.	VBench [8] ↑						Opus4.7-VL ↑			KV (GB)↓
			Subject Consistency	Background Consistency	Motion Smoothness	Dynamic Degree	Aesthetic Quality	Total Quality	Visual	Dynamic	Text	
Self-Forcing 1.3B (5-second)	Baseline	BF16	93.57	94.07	98.43	75.00	60.07	82.94	3.81	3.10	3.62	67.49 _{bs=12}
	FlexGen	INT6	93.52	94.11	98.44	<u>76.00</u>	60.02	<u>82.98</u>	3.79	<u>3.15</u>	3.73	<u>27.95</u> _{-58.59%}
	KIVI	INT6	93.58	94.12	98.49	73.00	60.07	82.85	3.82	3.14	<u>3.63</u>	28.22 _{-58.18%}
	VAR-Q	INT6	93.96	94.76	98.49	76.50	61.25	83.34	3.89	3.17	3.62	27.50 _{-59.25%}
	FlexGen	INT4	86.39	85.84	95.40	73.00	57.39	75.91	3.60	3.40	<u>3.75</u>	<u>17.40</u> _{-74.22%}
	KIVI	INT4	<u>90.53</u>	<u>91.86</u>	<u>98.29</u>	<u>74.50</u>	57.64	<u>80.33</u>	3.54	<u>3.44</u>	3.76	17.68 _{-73.81%}
	VAR-Q	INT4	91.28	92.55	98.35	75.50	<u>57.40</u>	80.99	<u>3.55</u>	3.50	3.73	16.95 _{-74.88%}
LongLive 1.3B (30-second)	Baseline	BF16	91.44	93.83	98.94	44.50	60.74	80.92	3.75	3.25	3.69	64.28 _{bs=20}
	FlexGen	INT6	91.33	93.49	98.95	<u>46.50</u>	60.29	80.60	3.78	3.33	3.78	26.62 _{-58.59%}
	KIVI	INT6	91.14	93.45	<u>98.94</u>	43.50	60.36	<u>80.65</u>	3.73	3.27	3.70	26.88 _{-58.18%}
	VAR-Q	INT6	91.26	93.57	98.93	47.50	60.17	80.94	3.84	<u>3.31</u>	<u>3.71</u>	26.19 _{-59.25%}
	FlexGen	INT4	<u>86.72</u>	90.19	95.83	45.00	55.16	76.59	<u>3.48</u>	3.48	3.77	<u>16.57</u> _{-74.22%}
	KIVI	INT4	85.83	92.09	<u>97.80</u>	49.50	<u>56.59</u>	<u>78.95</u>	<u>3.48</u>	<u>3.55</u>	<u>3.79</u>	16.84 _{-73.81%}
	VAR-Q	INT4	88.39	<u>91.97</u>	98.80	<u>47.50</u>	57.11	79.60	3.54	3.57	3.80	16.15 _{-74.88%}

Self-Forcing and LongLive. At INT6, VAR-Q is indistinguishable from BF16 on VBench Total (83.34 and 80.94 vs. 82.94 and 80.92). Opus4.7-VL ratings on LongLive corroborate this: 3.84 / 3.31 / 3.71 against BF16’s 3.75 / 3.25 / 3.69. At INT4, VAR-Q’s gap to BF16 stays within 2 VBench points (1.95 and 1.32), against 2.61 and 1.97 for KIVI and 7.03 and 4.33 for FlexGen; cache savings reach 75%. Together with the VAR result in Section 4.1, these results establish that VAR-Q’s two-axis grouping transfers across both visual AR paradigms.

4.3 Ablation Studies

Grouping rule. Table 5 ablates the two axes of VAR-Q at 4-bit on VAR-d30 (next-scale) and Self-Forcing INT4 (next-frame). On VAR-d30, removing the channel axis collapses generation (FID 12.07, IS 159.52); removing the block axis raises FID by 0.86 (1.97 to 2.83). The channel axis is the dominant factor, the block axis secondary. The trend persists on Self-Forcing: removing the channel axis drops Subject Consistency from 91.22 to 79.17 and Dynamic Degree from 74.00 to 27.00. VAR-Q covers both axes by construction and is therefore the safe paradigm-agnostic default. Among LLM-style alternatives, per-head per-token trails by 0.24 FID on VAR-d30 and 1.04 points on Subject; per-token alone collapses on both backbones with FID of 7.13 and Subject of 78.55.

Table 5: Grouping ablation at INT4 on next-scale VAR-d30 and next-frame Self-Forcing. VAR-Q (full) is the proposed method. Shaded rows ablate one axis; the bottom rows compare to LLM-style alternatives. Bold marks the best in each column.

Grouping	VAR-d30 (next-scale)						Self-Forcing (next-frame)				
	FID↓	IS↑	PSNR↑	SSIM↑	LPIPS↓	KV (GB)↓	Subject↑	Dynamic↑	Aesthetic↑	Total↑	KV (GB)↓
VAR-Q (full)	1.97	323.67	13.54	0.312	0.437	10.44 _{-72.06%}	91.28	75.50	57.40	80.99	16.95
w/o block-axis	2.83	292.13	12.07	0.257	0.538	9.45 _{-74.71%}	88.49	69.00	54.51	79.38	16.90
w/o channel-axis	12.07	159.52	10.42	0.198	0.645	9.34 _{-75.00%}	79.17	27.00	33.86	67.22	16.87
<i>Additional baselines (LLM-style groupings):</i>											
Per-token & per-head	2.21	292.12	11.97	0.261	0.545	10.51 _{-71.87%}	90.18	67.00	54.96	79.46	17.93
Per-token	7.13	203.50	10.51	0.208	0.635	9.38 _{-74.90%}	78.55	25.50	39.51	67.72	16.96

Effect of sub-block ratio r . Beyond the default per-block grouping ($r = 1$), finer sub-blocks improve quality on every backbone we test (Figure 9). On Self-Forcing INT4, VBench Quality rises from 80.99 at $r = 1$ to 81.34 at $r = 1/4$ and saturates at $r = 1/8$. At $r = 1/8$, gains range from +0.34 for Infinity-2B INT3 to +1.07 for InfinityStar INT3, with Infinity-8B INT2 at +0.88. Crucially, r partitions within existing block boundaries, so this refinement adds no detection or calibration cost.

Effect of Boundary Alignment. We define two schemes with exactly matched groupings. *Aligned* places group boundaries on block boundaries; *misaligned* shifts every boundary by half a block. This minimal perturbation reduces the generation quality of Self-Forcing by 0.45 in Table 6 and uniformly degrades all dimensions.




Figure 9: The effect of sub-block ratio r .

5 Related Work

Visual autoregression has consolidated around two block-wise paradigms: *next-scale* prediction [4, 12, 24], which decodes images coarse-to-fine across resolutions, and *next-frame* prediction [7, 29], which decodes video frame-by-frame. Both emit a block of tokens per autoregressive step, so the KV cache scales with horizon and dominates memory in long-form synthesis. KV-cache compression for language models is mature, spanning low-precision quantization [2, 5, 13, 14, 21, 30], attention-guided eviction [11, 32], and hybrid quantization-sparsification [9]. These methods inherit token-centric assumptions of uniform importance, linear cache growth, and stationary statistics across steps, all of which fail under the block-wise dynamics of visual AR, where magnitudes shift sharply at block boundaries and concentrate on a small set of channels. Extensions to diffusion transformers [15] and multimodal LLMs [25] target adjacent regimes but inherit the same statistics. The closest concurrent work addresses visual AR directly. KV eviction [10, 28] and sparse-attention accelerators [26] reduce *which* entries are retained rather than their precision, and thus compose orthogonally with VAR-Q. Concurrent quantization efforts [16, 27] target next-scale image and next-frame video respectively, but neither attributes the cache structure to its architectural sources nor exploits the joint channel-block grouping; VAR-Q reuses block boundaries from the generation loop and applies one calibration-free rule across both paradigms.

6 Conclusion

VAR-Q groups KV activations along channel and block axes, reusing boundaries supplied by the generation loop and requiring no calibration. Across nine visual AR models spanning both next-scale and next-frame paradigms, it cuts KV memory by up to 78.5% with negligible quality degradation, enables $6\times$ larger inference batches, and makes large visual autoregressive models more practical to deploy.

Table 6: Boundary alignment ablation.

Grouping	Subj ↑	Dyn ↑	Aes ↑	Quality ↑
Aligned	91.23	74.50	57.41	80.91
Misaligned	90.80	73.50	56.75	80.46

References

- [1] Saleh Ashkboos, Amirkeivan Mohtashami, Maximilian L Croci, Bo Li, Pashmina Cameron, Martin Jaggi, Dan Alistarh, Torsten Hoefer, and James Hensman. Quarot: Outlier-free 4-bit inference in rotated llms. *Advances in Neural Information Processing Systems*, 37:100213–100240, 2024.
- [2] Daehyeon Baek, Jieun Choi, Jimyoung Son, Kyungmin Bin, Seungbeom Choi, Kihyo Moon, Minsung Jang, and Hyojung Lee. Fireq: Fast int4-fp8 kernel and rope-aware quantization for llm inference acceleration. *arXiv preprint arXiv:2505.20839*, 2025.
- [3] Dhruva Ghosh, Hannaneh Hajishirzi, and Ludwig Schmidt. Geneval: An object-focused framework for evaluating text-to-image alignment. *Advances in Neural Information Processing Systems*, 36:52132–52152, 2023.
- [4] Jian Han, Jinlai Liu, Yi Jiang, Bin Yan, Yuqi Zhang, Zehuan Yuan, Bingyue Peng, and Xiaobing Liu. Infinity: Scaling bitwise autoregressive modeling for high-resolution image synthesis. In *Proceedings of the Computer Vision and Pattern Recognition Conference*, pages 15733–15744, 2025.
- [5] Coleman Hooper, Sehoon Kim, Hiva Mohammadzadeh, Michael W Mahoney, Yakun S Shao, Kurt Keutzer, and Amir Gholami. Kvquant: Towards 10 million context length llm inference with kv cache quantization. *Advances in Neural Information Processing Systems*, 37:1270–1303, 2024.
- [6] Xiwei Hu, Rui Wang, Yixiao Fang, Bin Fu, Pei Cheng, and Gang Yu. Ella: Equip diffusion models with llm for enhanced semantic alignment. *arXiv preprint arXiv:2403.05135*, 2024.
- [7] Xun Huang, Zhengqi Li, Guande He, Mingyuan Zhou, and Eli Shechtman. Self forcing: Bridging the train-test gap in autoregressive video diffusion. *arXiv preprint arXiv:2506.08009*, 2025.
- [8] Ziqi Huang, Yinan He, Jiashuo Yu, Fan Zhang, Chenyang Si, Yuming Jiang, Yuanhan Zhang, Tianxing Wu, Qingyang Jin, Nattapol Chanpaisit, et al. Vbench: Comprehensive benchmark suite for video generative models. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 21807–21818, 2024.
- [9] Hao Kang, Qingru Zhang, Souvik Kundu, Geonhwa Jeong, Zaoxing Liu, Tushar Krishna, and Tuo Zhao. GEAR: An efficient error reduction framework for KV cache compression in LLM inference. In Mehdi Rezagholizadeh, Peyman Passban, Soheila Samiee, Vahid Partovi Nia, Yu Cheng, Yue Deng, Qun Liu, and Boxing Chen, editors, *Proceedings of The 4th NeurIPS Efficient Natural Language and Speech Processing Workshop*, volume 262 of *Proceedings of Machine Learning Research*, pages 305–321. PMLR, 14 Dec 2024. URL <https://proceedings.mlr.press/v262/kang24a.html>.
- [10] Kunjun Li, Zigeng Chen, Cheng-Yen Yang, and Jenq-Neng Hwang. Memory-efficient visual autoregressive modeling with scale-aware KV cache compression. In *Advances in Neural Information Processing Systems (NeurIPS)*, 2025.
- [11] Yuhong Li, Yingbing Huang, Bowen Yang, Bharat Venkitesh, Acyr Locatelli, Hanchen Ye, Tianle Cai, Patrick Lewis, and Deming Chen. Snapkv: LLM knows what you are looking for before generation. *Advances in Neural Information Processing Systems*, 37:22947–22970, 2024.
- [12] Jinlai Liu, Jian Han, Bin Yan, Hui Wu, Fengda Zhu, Xing Wang, Yi Jiang, Bingyue Peng, and Zehuan Yuan. Infinitystar: Unified spacetime autoregressive modeling for visual generation. In *The Thirty-ninth Annual Conference on Neural Information Processing Systems*, 2025.
- [13] Zechun Liu, Changsheng Zhao, Igor Fedorov, Bilge Soran, Dhruv Choudhary, Raghuraman Krishnamoorthi, Vikas Chandra, Yuandong Tian, and Tijmen Blankevoort. Spinquant: Llm quantization with learned rotations. *arXiv preprint arXiv:2405.16406*, 2024.

- [14] Zirui Liu, Jiayi Yuan, Hongye Jin, Shaochen Zhong, Zhaozhuo Xu, Vladimir Braverman, Beidi Chen, and Xia Hu. Kivi: A tuning-free asymmetric 2bit quantization for KV cache. *arXiv preprint arXiv:2402.02750*, 2024.
- [15] Xinyin Ma, Gongfan Fang, and Xinchao Wang. Deepcache: Accelerating diffusion models for free. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 15762–15772, 2024.
- [16] Jaehyeon Moon and Bumsu Ham. Shift-and-sum quantization for visual autoregressive models. In *The Fourteenth International Conference on Learning Representations*.
- [17] Adam Polyak, Amit Zohar, Andrew Brown, Andros Tjandra, Animesh Sinha, Ann Lee, Apoorv Vyas, Bowen Shi, Chih-Yao Ma, Ching-Yao Chuang, et al. Movie gen: A cast of media foundation models. *arXiv preprint arXiv:2410.13720*, 2024.
- [18] Ying Sheng, Lianmin Zheng, Binhang Yuan, Zhuohan Li, Max Ryabinin, Beidi Chen, Percy Liang, Christopher Ré, Ion Stoica, and Ce Zhang. Flexgen: High-throughput generative inference of large language models with a single gpu. In *International Conference on Machine Learning*, pages 31094–31116. PMLR, 2023.
- [19] Jianlin Su, Murtadha Ahmed, Yu Lu, Shengfeng Pan, Wen Bo, and Yunfeng Liu. Roformer: Enhanced transformer with rotary position embedding. *Neurocomputing*, 568:127063, 2024.
- [20] Zunhai Su, Wang Shen, Linge Li, Zhe Chen, Hanyu Wei, Huangqi Yu, and Kehong Yuan. Akvq-vl: Attention-aware kv cache adaptive 2-bit quantization for vision-language models. In *2025 IEEE International Conference on Multimedia and Expo (ICME)*, pages 1–6. IEEE, 2025.
- [21] Zunhai Su, Hanyu Wei, Zhe Chen, Wang Shen, Linge Li, Huangqi Yu, and Kehong Yuan. Rotatekv: accurate and robust 2-bit kv cache quantization for llms via outlier-aware adaptive rotations. In *Proceedings of the Thirty-Fourth International Joint Conference on Artificial Intelligence, IJCAI '25*, 2025. ISBN 978-1-956792-06-5. doi: 10.24963/ijcai.2025/690. URL <https://doi.org/10.24963/ijcai.2025/690>.
- [22] Mingjie Sun, Xinlei Chen, J Zico Kolter, and Zhuang Liu. Massive activations in large language models. In *First Conference on Language Modeling (COLM)*, 2024.
- [23] Qian Tao, Wenyuan Yu, and Jingren Zhou. Asymkv: Enabling 1-bit quantization of kv cache with layer-wise asymmetric quantization configurations. In *Proceedings of the 31st International Conference on Computational Linguistics*, pages 2316–2328, 2025.
- [24] Keyu Tian, Yi Jiang, Zehuan Yuan, Bingyue Peng, and Liwei Wang. Visual autoregressive modeling: Scalable image generation via next-scale prediction. *Advances in neural information processing systems*, 37:84839–84865, 2024.
- [25] Zhongwei Wan, Ziang Wu, Che Liu, Jinfa Huang, Zhihong Zhu, Peng Jin, Longyue Wang, and Li Yuan. Look-m: Look-once optimization in kv cache for efficient multimodal long-context inference. In *Findings of the Association for Computational Linguistics: EMNLP 2024*, pages 4065–4078, 2024.
- [26] Haocheng Xi, Shuo Yang, Yilong Zhao, Chenfeng Xu, Muyang Li, Xiuyu Li, Yujun Lin, Han Cai, Jintao Zhang, Dacheng Li, Jianfei Chen, Ion Stoica, Kurt Keutzer, and Song Han. Sparse video-gen: Accelerating video diffusion transformers with spatial-temporal sparsity. In Aarti Singh, Maryam Fazel, Daniel Hsu, Simon Lacoste-Julien, Felix Berkenkamp, Tegan Maharaj, Kiri Wagstaff, and Jerry Zhu, editors, *Proceedings of the 42nd International Conference on Machine Learning*, volume 267 of *Proceedings of Machine Learning Research*, pages 68208–68224. PMLR, 13–19 Jul 2025. URL <https://proceedings.mlr.press/v267/xi25c.html>.
- [27] Haocheng Xi, Shuo Yang, Yilong Zhao, Muyang Li, Han Cai, Xingyang Li, Yujun Lin, Zhuoyang Zhang, Jintao Zhang, Xiuyu Li, et al. Quant videogen: Auto-regressive long video generation via 2-bit kv-cache quantization. *arXiv preprint arXiv:2602.02958*, 2026.

- [28] Boxun Xu, Yu Wang, Zihu Wang, and Peng Li. Ams-kv: Adaptive kv caching in multi-scale visual autoregressive transformers. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 40, pages 27206–27214, 2026.
- [29] Shuai Yang, Wei Huang, Ruihang Chu, Yicheng Xiao, Yuyang Zhao, Xianbang Wang, Muyang Li, Enze Xie, Yingcong Chen, Yao Lu, et al. Longlive: Real-time interactive long video generation. *arXiv preprint arXiv:2509.22622*, 2025.
- [30] Dingyu Yao et al. VecInfer: Efficient LLM inference with low-bit KV cache via outlier-suppressed vector quantization. *arXiv preprint arXiv:2510.06175*, 2025.
- [31] Mengxia Yu, De Wang, Qi Shan, Colorado J Reed, and Alvin Wan. The super weight in large language models. *arXiv preprint arXiv:2411.07191*, 2024.
- [32] Zhenyu Zhang, Ying Sheng, Tianyi Zhou, Tianlong Chen, Lianmin Zheng, Ruisi Cai, Zhao Song, Yuandong Tian, Christopher Ré, Clark Barrett, et al. H2o: Heavy-hitter oracle for efficient generative inference of large language models. *Advances in Neural Information Processing Systems*, 36:34661–34710, 2023.
- [33] Lianmin Zheng, Wei-Lin Chiang, Ying Sheng, Siyuan Zhuang, Zhanghao Wu, Yonghao Zhuang, Zi Lin, Zhuohan Li, Dacheng Li, Eric Xing, et al. Judging llm-as-a-judge with mt-bench and chatbot arena. *Advances in neural information processing systems*, 36:46595–46623, 2023.

A Asymmetric Roles of K and V Across Modalities

A.1 K/V asymmetry between image and video AR

The joint result above masks a per-tensor split: **video models are more sensitive to K outliers, while image models are more sensitive to V outliers**. Mechanistically, **K** determines which past tokens are attended to and **V** determines what content is retrieved. The split therefore suggests that video generation places a heavier demand on precise temporal addressing into the cache, whereas image generation tolerates coarser addressing but depends more on the fidelity of the content retrieved at each cached position. The same per-tensor ablation on LLMs reproduces an LLM-side **K**-dominance consistent with prior reports [14, 23], placing video AR closer to the LLM regime and image AR farther from it.

This appendix provides per-tensor evidence for the **K/V** asymmetry discussed in the main paper, separating **K**-only and **V**-only ablations across image and video AR, and extending the comparison to LLaMA-3-8B-Instruct. We follow the channel-wise ablation protocol of the main paper but mask **K** and **V** independently. Channels are ranked by per-token absolute maximum aggregated over a calibration set; the top, bottom, or random $p\%$ are zeroed at inference, with $p \in \{0.1, 0.5, 3, 5\}$. Four visual AR backbones (Infinity-2B, Infinity-8B, LongLive-30s, InfinityStar-720p) are evaluated with the perceptual and VBench metrics used in the main paper, and LLaMA-3-8B-Instruct is evaluated on WikiText-2 perplexity (4096-token chunks), MMLU (200 sampled questions), and GSM8K (50 sampled questions, chat-template greedy generation). Random and inlier ablation arms remain within metric-level noise of the unmasked baseline across all backbones and ratios, confirming that the observed degradation is selective to high-magnitude channels rather than to channel count. One measurement caveat applies to video: the VBench dimensions *dynamic_degree* and *temporal_flickering* register increases under **K**-outlier ablation that do not reflect genuine motion, since the per-frame chroma drift induced by **K**-side collapse is misclassified as motion by the underlying classifiers; we therefore rely on *imaging_quality* (MUSIQ) and *overall_consistency* (ViCLIP) as the most reliable VBench dimensions for this perturbation regime.

A.2 Extension to LLaMA-3-8B

The top two rows of Figure 10 show that image and video AR sit on opposite sides of the **K/V** axis. Image AR is **V**-dominant: on Infinity-2B at $p = 5\%$, **V**-only ablation drives ΔLPIPS to 0.20 against 0.08 for **K**-only, and the LAION aesthetic score drops by 1.3 under **V**-only against near-zero under **K**-only. Video AR reverses the direction: on LongLive-30s at $p = 5\%$, **K**-only ablation drives $\Delta\text{LPIPS}_{\text{mid}}$ to 0.40 against 0.14 for **V**-only, and the same direction holds on InfinityStar-720p. The image-side **V**-dominance is consistent across both Infinity backbones and the video-side **K**-dominance across both video backbones, indicating that the asymmetry tracks the modality rather than per-model idiosyncrasy. Figure 11 renders the same contrast at the sample level: on Infinity-2B, **V**-only masking eliminates recognizable scene content while **K**-only masking is nearly invisible; on LongLive-1.3B, **K**-only masking collapses frames into uniform color while **V**-only masking preserves scene layout with degraded texture. The joint **K+V** mask in both cases tracks the more damaging single-tensor mask, consistent with the dominant tensor carrying the load-bearing signal.

The bottom row of Figure 10 extends the same protocol to LLaMA-3-8B-Instruct. The model behaves like video AR in direction but more sharply in magnitude: at $p = 0.5\%$, **K**-only ablation increases WikiText-2 perplexity from 8.4 to 82.4 ($\sim 10\times$), while **V**-only ablation leaves it at 11.5 ($1.4\times$); MMLU and GSM8K accuracy degrade in the same direction. Figure 12 summarizes the damage ratio $\Delta_{\mathbf{K}}/\Delta_{\mathbf{V}}$ at $p = 0.5\%$: image AR sits at or below $1\times$, video AR ranges from $2.5\times$ to $61\times$, and LLaMA-3-8B reaches $320\times$ on perplexity. The LLM observation aligns with prior reports of **K**-side outlier dominance in LLM KV caches [14, 23], placing video AR closer to the LLM regime while leaving image AR as the singular setting in which **V** outliers carry the bulk of the load.

B Persistency Across Different Prompts

As shown in Figure 13, the shapes of the **K** and **V** curves differ qualitatively, not just in magnitude. **V** (dashed) collapses on the first added prompt: even a single second prompt removes about two-thirds of the outlier channels ($\text{Persistence}@2 = 0.31$ to 0.35 across the three backbones), and the curve continues to decay toward 0.07 to 0.14. **K** (solid) instead bends gently, retaining 0.77 to 0.82 at $N = 2$

Cross-paradigm K vs V outlier-channel prune: damage as Δ vs inlier-prune baseline




Figure 10: **K**-only (red) versus **V**-only (blue) outlier ablation. **Top row**: image AR backbones on perceptual and aesthetic metrics. **Middle row**: video AR backbones on perceptual and VBench-MUSIQ metrics. **Bottom row**: LLaMA-3-8B-Instruct on language benchmarks (PPL on log scale). All metrics are oriented so that higher values indicate worse generation. The dominant color flips between image and video AR.

and asymptoting at 0.53 to 0.62. The K-V gap is already $2.3\times$ at $N=2$ and widens monotonically with N , so the $N=10$ cross-section reported as Persistence@10 is the rightmost point of two curves whose decay rates differ throughout. The plateau shape on K is the signature of an architectural core: most outlier channels survive every additional prompt, consistent with their being selected by the static W_K , γ_K , and RoPE-pair chain (Section 3.3). The crash on V is the signature of input-driven outliers: each new prompt is a fresh sample of token amplitudes, and the V channel identity tracks those samples.

C Approximate ℓ_2 Invariance of Channel Max under RoPE

Within one block of T tokens, let pair (c_t, d_t) at token t have pre-RoPE amplitude $M_t = \sqrt{c_t^2 + d_t^2}$, with order statistics $M_{(1)} \geq M_{(2)} \geq \dots \geq M_{(T)}$. After RoPE rotation by a shared angle θ , write $a_t = M_t \cos(\theta + \phi_t)$ and $b_t = M_t \sin(\theta + \phi_t)$, where ϕ_t is the per-token pre-RoPE angle. The pair of channel maxes satisfies

$$M_{(1)}^2 \leq \|a\|_\infty^2 + \|b\|_\infty^2 \leq M_{(1)}^2 + M_{(2)}^2. \quad (4)$$

Lower bound. Let $t^* = \arg \max_t M_t$. Per-token RoPE invariance gives $a_{t^*}^2 + b_{t^*}^2 = M_{(1)}^2$. Since $\|a\|_\infty \geq |a_{t^*}|$ and $\|b\|_\infty \geq |b_{t^*}|$, the sum is at least $M_{(1)}^2$.




Figure 11: Generation under **K**-only, **V**-only, and joint **K+V** outlier masking. **(a)** Infinity-2B at $p = 5\%$: **V**-only masking collapses outputs into texture or symmetric patterns, while **K**-only masking leaves outputs visually close to the baseline. **(b)** LongLive-1.3B at $p = 0.5\%$: **K**-only masking collapses frames into solid color and chroma stripes, while **V**-only masking preserves scene layout with degraded fine detail. Inlier masking matches the baseline in both cases.




Figure 12: Δ_K / Δ_V damage ratio at $p = 0.5\%$ (log scale). Bars below the dashed **K=V** line are **V**-dominant; bars above are **K**-dominant. Image AR sits at or below 1x, video AR ranges from 2.5x to 61x, and LLaMA-3-8B reaches 320x on perplexity.

Upper bound. Let $t_a = \arg \max_t |a_t|$ and $t_b = \arg \max_t |b_t|$. Per-token, $|a_t|, |b_t| \leq M_t$, so $\|a\|_\infty^2 \leq M_{t_a}^2$ and $\|b\|_\infty^2 \leq M_{t_b}^2$. If $t_a = t_b$, the sum equals $M_{t_a}^2 \leq M_{(1)}^2$. If $t_a \neq t_b$, at least one of them is not t^* , so $\min(M_{t_a}, M_{t_b}) \leq M_{(2)}$ and the sum is at most $M_{(1)}^2 + M_{(2)}^2$.




Figure 13: Outlier persistence across different amounts of prompts

Massive-value regime. When pre-RoPE amplitudes are heavy-tailed within the block, $M_{(2)}^2 \ll M_{(1)}^2$ and the bounds collapse, leaving channel-max pair ℓ_2 approximately invariant under RoPE rotation. This is the mathematical basis for the envelope-flat behavior in Fig. 7(b).

D Prompt Template for VLM-based Video Evaluation

Evaluation Prompt. For each video, we partition the timeline into 12 equal segments and sample one frame at the midpoint of each segment, i.e., at relative timestamps $(i + 0.5)/12$ for $i = 0, \dots, 11$. This produces samples spanning 4.17% to 95.83% of the video while excluding boundary regions, where transition artifacts are more common. Frames are extracted with `ffmpeg`, resized to 360 pixels in height with aspect ratio preserved, and stored as JPEG images using `-q:v 3`. To ensure a fair comparison, we keep the 12 sampling locations fixed across all quantization settings for the same prompt, so the VLM judge is exposed to identical temporal evidence. This corresponds to an average spacing of approximately 0.42 seconds for 5-second Self-Forcing videos and 2.5 seconds for 30-second LongLive videos. We provide the exact prompt template used to query the VLM judge for rating text alignment, dynamics, and visual quality.

You are a video quality evaluation expert. Below are 12 frames uniformly sampled from a generated video, plus the text prompt that was used to produce it. Rate the video on three dimensions (1-5 integer each):

1. Text Alignment - consistency between video content and the text description.
 - 1: Completely Irrelevant. 2: Mostly Mismatched. 3: Partially Matched.
 - 4: Largely Consistent. 5: Perfectly Aligned.
2. Dynamics - dynamism/fluidity of the scene (camera, objects, transitions).
 - 1: Static/Disjointed. 2: Mostly Static. 3: Moderately Dynamic.
 - 4: Largely Dynamic. 5: Highly Dynamic.
3. Visual Quality - technical execution (clarity, color, composition, artifacts).
 - 1: Very Poor. 2: Poor. 3: Fair. 4: Good. 5: Excellent.

Output STRICTLY in this format (nothing else, no markdown fences):

Text Alignment: <1-5>

Dynamics: <1-5>

Visual Quality: <1-5>

Text Prompt: ""{TEXT_PROMPT}""

Frames (in temporal order): @f00.jpg @f01.jpg ... @f11.jpg

E VAR-Q System Supports

Algorithm 1: VAR-Q Quantization

Input: KV tensor X , bit-width b , quantization method M , qkv layout F , packing flag pack

Output: quantized tensor $\hat{Q}$, scale S , quantization metadata m_q , packing metadata m_p

```

1  $q_{\min} \leftarrow -2^{b-1}, \quad q_{\max} \leftarrow 2^{b-1} - 1$ 
2  $D \leftarrow \text{REDUCEDIMS}(M, F)$ 
3  $S \leftarrow \max(|X|, \dim = D, \text{keepdim} = \text{True}) / q_{\max}$ 
4  $S \leftarrow \max(S, \epsilon)$ 
5  $Q \leftarrow \text{round}(X/S)$ 
6  $Q \leftarrow \text{clamp}(Q, q_{\min}, q_{\max})$ 
7  $Q \leftarrow \text{int8}(Q)$ 
8 if pack then
9    $(\hat{Q}, m_p) \leftarrow \text{PACKLASTDIM}(Q, b)$ 
10 else
11    $\hat{Q} \leftarrow Q, \quad m_p \leftarrow \emptyset$ 
12  $m_q \leftarrow (M, F, \text{shape}(X))$ 
13 return  $\hat{Q}, S, m_q, m_p$ 

```

Algorithm 2: VAR-Q Dequantization

Input: quantized tensor $\hat{Q}$, scale S , quantization metadata m_q , packing metadata m_p , output dtype τ
Output: dequantized tensor $\tilde{X}$

```
1 if  $m_p \neq \emptyset$  then  
2    $Q \leftarrow \text{UNPACKLASTDIM}(\hat{Q}, m_p)$   
3 else  
4    $Q \leftarrow \hat{Q}$   
5  $S \leftarrow \text{BROADCASTSCALE}(S, m_q, \text{shape}(Q))$   
6  $\tilde{X} \leftarrow \text{Cast}_{\tau}(\text{float32}(Q) \odot S)$   
7 return  $\tilde{X}$ 
```

The full quantization stack consists of four operators (Algorithms 1–4) that together turn the per-(channel, block) grouping rule of Section 3.4 into actual KV-memory savings on GPU. Algorithm 1 performs uniform symmetric quantization on each newly emitted block: the reduction axes D are derived from the quantization method M and qkv layout F so that one scaling factor is computed per quantization unit, an ϵ floor guards against numerically zero groups, and the rounded integers are clamped to the signed range $[-2^{b-1}, 2^{b-1} - 1]$ and held in `int8`. Algorithm 2 is the symmetric inverse: it broadcasts S back to the cache shape recorded in m_q , multiplies in `float32`, and casts to the requested attention dtype τ . Both operators are stateless with respect to other blocks, which matches the online property of the algorithm: a block is quantized once at write time and never revisited.

Quantization alone does not shrink the cache, because current CUDA toolchains expose no native sub-8-bit storage type. Algorithms 3 and 4 close this gap by packing $w = \lfloor 32/b \rfloor$ signed b -bit values into a single `uint32` word along the last (channel) dimension. We support packed storage for INT8, INT6, INT4, INT3, and INT2: INT8, INT4, and INT2 exactly fill a 32-bit word, whereas INT6 and INT3 leave two unused bits per word, yielding effective storage ratios of 40% and 20% relative to BF16, respectively. The packer pads the channel axis up to a multiple of w , masks each lane to b bits, and ORs the lanes into their assigned slots via the shifts kb for $k = 0, \dots, w - 1$; the original channel count C_{orig} , pad amount p , lane count w , and bit-width b are recorded in m_p so that the unpacker can shift, mask, sign-extend, and trim back to the exact original shape. This converts the nominal bit-width reduction into a real cache-byte reduction, and the lane layout is byte-aligned so loads remain coalesced.

The four operators compose into the runtime pipeline expected by visual AR inference. At write time, a freshly produced block is passed through Algorithm 1 with `pack = True`, which internally invokes Algorithm 3; the resulting `uint32` tensor, scale, and metadata (m_q, m_p) are appended to the cache. At attention time, Algorithm 2 is invoked on the cached blocks, calling Algorithm 4 only when $m_p \neq \emptyset$ and otherwise falling back to a plain `int8` dequant path. The whole stack is implemented in Triton and runs without calibration data, retraining, or model-specific kernel variants, so it drops into Infinity, InfinityStar, Self-Forcing, LongLive, and VAR with no code changes outside the cache module.

Algorithm 3: Pack Last Dimension into 32-bit Words

Input: quantized tensor Q stored in int8, bit-width b

Output: packed tensor P stored in uint32, packing metadata m_p

```
1  $C_{\text{orig}} \leftarrow \text{LastDim}(Q)$ 
2  $w \leftarrow 32/b$  // number of  $b$ -bit values per 32-bit word
3  $p \leftarrow (-C_{\text{orig}}) \bmod w$ 
4  $C_{\text{pad}} \leftarrow C_{\text{orig}} + p$ 
5  $C_w \leftarrow C_{\text{pad}}/w$ 
6  $Q' \leftarrow \text{PADLASTDIM}(Q, p)$ 
7  $Q' \leftarrow \text{RESHAPELASTDIM}(Q', [C_w, w])$ 
8  $P \leftarrow 0 \in \text{uint32}^{\text{shape}(Q)[:-1] \times C_w}$ 
9 for  $k \leftarrow 0$  to  $w - 1$  do
10    $U_k \leftarrow Q'[\dots, k] \& ((1 \ll b) - 1)$ 
11    $P \leftarrow P \mid (\text{uint32}(U_k) \ll (kb))$ 
12  $m_p \leftarrow (C_{\text{orig}}, p, w, b)$ 
13 return  $P, m_p$ 
```

Algorithm 4: Unpack 32-bit Words along the Last Dimension

Input: packed tensor P stored in uint32, packing metadata $m_p = (C_{\text{orig}}, p, w, b)$

Output: recovered quantized tensor Q stored in int8

```
1  $mask \leftarrow (1 \ll b) - 1$ 
2  $U \leftarrow 0 \in \text{uint32}^{\text{shape}(P) \times w}$ 
3 for  $k \leftarrow 0$  to  $w - 1$  do
4    $U[\dots, k] \leftarrow (P \gg (kb)) \& mask$ 
5  $Q' \leftarrow \text{SIGNEXTEND}(U, b)$ 
6  $Q' \leftarrow \text{FLATTENLASTTWO DIMS}(Q')$ 
7  $Q \leftarrow Q'[\dots, 0 : C_{\text{orig}}]$ 
8 return  $Q$ 
```




Figure 14: Pre-RoPE vs. post-RoPE quantization on Self-Forcing at INT6. Since RoPE preserves the pair-level envelope but redistributes magnitude across paired channels, quantizing after RoPE better matches the final attention-ready representation. Post-RoPE quantization better preserves dynamic behavior and overall visual quality in the shown examples.

F Pre-RoPE vs. post-RoPE quantization.

We ablate the quantization point of VAR-Q on Self-Forcing at INT6. As shown in Table 7, post-RoPE quantization consistently outperforms pre-RoPE on the metrics most related to temporal behavior and overall generation quality: Subject Consistency, Background Consistency, Dynamic Degree, Aesthetic Quality, and Total Quality. The gain in Dynamic Degree (76.50 vs. 74.50) indicates that post-RoPE quantization better preserves motion cues instead of over-smoothing the generated sequence.

Table 7: Pre-RoPE vs. post-RoPE quantization at INT6.

Variant	Prec.	VBench [8] ↑					KV (GB) ↓
		Subj.	Bg.	Dyn.	Aes.	Quality	
post-RoPE INT6	93.96	94.76	76.50	61.25	83.34	27.50 _{-59.25%}	
pre-RoPE INT6	93.63	93.97	74.50	60.00	82.93	27.50 _{-59.25%}	

Meanwhile, the improvements in Subject and Background Consistency suggest that quantizing the attention-ready representation also stabilizes the appearance of both foreground objects and scene layout across frames. Post-RoPE further improves Aesthetic Quality and Total Quality (83.34 vs. 82.93), showing that the benefit is not limited to motion but also carries over to perceptual quality.

This supports quantizing the final post-RoPE cache, which is directly consumed by attention, rather than the pre-RoPE cache. We therefore use post-RoPE quantization as the default.

G Qualitative Examples

This appendix provides additional qualitative comparisons across next-scale image generation, next-scale video generation, and next-frame video generation. For VAR, each figure arranges different quantization configurations by row and ImageNet classes by column, with the class ID shown above each column; INT8 examples include only BF16 and VAR-Q. For Infinity, each row corresponds to one text prompt and each column to one configuration, so images in the same row are generated from the same prompt. For video models, each page shows one prompt, with configurations arranged by row and temporally sampled frames arranged by column.

Qualitative observations. Across these examples, the visual trends are consistent with the quantitative results in the main paper. At moderate precision, VAR-Q largely preserves the appearance, layout, and temporal evolution of the BF16 outputs, while competing KV-cache quantization baselines more often introduce visible softening, local texture loss, or later-frame drift. We therefore keep individual figure captions concise and summarize the recurring visual patterns below, rather than describing every sample independently.

Next-scale image generation. For VAR and Infinity, the main visible difference appears in local fidelity rather than coarse semantics. Most methods preserve the high-level object category or prompt content, but FlexGen and KIVI more frequently lose fine textures, weaken small structures, or alter local details under low-bit KV caches. VAR-Q remains visually closer to BF16 across the shown samples, especially at INT4 and INT3, where it better preserves object boundaries, repeated patterns, and prompt-specific visual attributes.



Figure 15: Qualitative comparison on VAR d20 with INT8 KV cache.




Figure 16: Qualitative comparison on VAR d20 with INT4 KV cache.




Figure 17: Qualitative comparison on VAR d20 with INT3 KV cache.




Figure 18: Qualitative comparison on VAR d24 with INT8 KV cache.




Figure 19: Qualitative comparison on VAR d24 with INT4 KV cache.




Figure 20: Qualitative comparison on VAR d24 with INT3 KV cache.

Infinity2B Q8 qualitative comparison

BF16

VAR-Q



a futuristic spaceship with a design reminiscent of the iconic Sydney Opera House, featuring multiple white, shell-like structures that form its hull. The vessel hovers above the ground, with a slight indescendent sheen on its surface that reflects the light of a distant sun. It is surrounded by a barren landscape, which contrasts sharply with the spaceship's smooth, curved architecture.



A detailed vector illustration showcasing four separate square compositions, each highlighting different aspects of waste management and recycling processes. The vibrant images feature a variety of flat design icons and pictograms that represent sorting bins, recycling symbols, and cleaning equipment. Each square has a distinct color palette to categorize the type of waste, ranging from organic to plastic, metal, and paper, emphasizing the importance of recycling and proper disposal.



A figure depicted with a dynamic pose: its left arm extends forward while the right arm is elongated back, parallel with its spine. Both legs are positioned closely together in a near vertical line, with the right leg gracefully stacked atop the left. The silhouette reflects a ballet dancer's poise during a challenging balancing act.

Figure 21: Qualitative comparison on Infinity 2B with INT8 KV cache.




Figure 22: Qualitative comparison on Infinity 2B with INT4 KV cache.



Figure 23: Qualitative comparison on Infinity 2B with INT3 KV cache.

Infinity8B Q8 qualitative comparison

BF16

VAR-Q



A baffling scene where smoke is inexplicably wafting from the filter end of a cigarette between a person's fingers, rather than the lit end. The cigarette is resting in an ashtray that's placed on a round, glass-topped table. Stray ashes can be seen scattered around the ashtray, highlighting the peculiarity of the situation.



A man with bright red boxing gloves awkwardly attempts to play a glossy black grand piano in the center of a room. Around him, the floor is a polished hardwood, reflecting the soft overhead lighting. Despite the gloves, he appears concentrated, engaged in his unique challenge, with sheet music propped up on the piano's stand.



In the center of a bustling intersection, a large tree with a thick trunk and sprawling branches stands out amidst the concrete. Its green leaves contrast sharply with the grey asphalt roads that converge around it. Traffic lights and street signs are positioned awkwardly around the tree's base, creating an unusual juxtaposition of nature and urban infrastructure.

Figure 24: Qualitative comparison on Infinity 8B with INT8 KV cache.

Infinity8B Q3 qualitative comparison

BF16

KIVI

VAR-Q



A baffling scene where smoke is inexplicably wafting from the filter end of a cigarette between a person's fingers, rather than the lit end. The cigarette is resting in an ashtray that's placed on a round, glass-topped table. Stray ashes can be seen scattered around the ashtray, highlighting the peculiarity of the situation.



A man with bright red boxing gloves awkwardly attempts to play a glossy black grand piano in the center of a room. Around him, the floor is a polished hardwood, reflecting the soft overhead lighting. Despite the gloves, he appears concentrated, engaged in his unique challenge, with sheet music propped up on the piano's stand.



In the center of a bustling intersection, a large tree with a thick trunk and sprawling branches stands out amidst the concrete. Its green leaves contrast sharply with the grey asphalt roads that converge around it. Traffic lights and street signs are positioned awkwardly around the tree's base, creating an unusual juxtaposition of nature and urban infrastructure.

Figure 25: Qualitative comparison on Infinity 8B with INT3 KV cache.

Infinity8B Q2 qualitative comparison

BF16

VAR-Q



A baffling scene where smoke is inexplicably wafting from the filter end of a cigarette between a person's fingers, rather than the lit end. The cigarette is resting in an ashtray that's placed on a round, glass-topped table. Stray ashes can be seen scattered around the ashtray, highlighting the peculiarity of the situation.



A man with bright red boxing gloves awkwardly attempts to play a glossy black grand piano in the center of a room. Around him, the floor is a polished hardwood, reflecting the soft overhead lighting. Despite the gloves, he appears concentrated, engaged in his unique challenge, with sheet music propped up on the piano's stand.



In the center of a bustling intersection, a large tree with a thick trunk and sprawling branches stands out amidst the concrete. Its green leaves contrast sharply with the grey asphalt roads that converge around it. Traffic lights and street signs are positioned awkwardly around the tree's base, creating an unusual juxtaposition of nature and urban infrastructure.

Figure 26: Qualitative comparison on Infinity 8B with INT2 KV cache.




Figure 27: Qualitative comparison on InfinityStar 480p with INT8 KV cache.



Figure 28: Qualitative comparison on InfinityStar 480p with INT4 KV cache.

Next-scale video generation. For InfinityStar, the sampled frames show that quantization errors are more apparent in temporal consistency and imaging quality than in single-frame semantics. FlexGen and KIVI often preserve the rough scene content, but their later frames more frequently exhibit softened details or less stable background structure. VAR-Q better maintains the BF16 trajectory across frames, with smaller visible drift and more consistent subject/background appearance.




Figure 29: Qualitative comparison on InfinityStar 480p with INT3 KV cache.




Figure 30: Qualitative comparison on InfinityStar 720p with INT8 KV cache.



Figure 31: Qualitative comparison on InfinityStar 720p with INT4 KV cache.



Figure 32: Qualitative comparison on InfinityStar 720p with INT3 KV cache.



Figure 33: Qualitative comparison on Self-Forcing with INT6 KV cache.



Figure 34: Qualitative comparison on Self-Forcing with INT6 KV cache.

Next-frame video generation. For Self-Forcing and LongLive, the qualitative comparison emphasizes long-horizon stability. In the 5-second Self-Forcing examples, VAR-Q keeps motion and scene structure close to BF16 under INT6 and INT4, while the baselines show stronger degradation in texture sharpness and frame-to-frame consistency. In the 30-second LongLive examples, differences become more visible in later sampled frames: VAR-Q better preserves subject appearance and background configuration, whereas the baselines more often accumulate drift or smoothing over time.




Figure 35: Qualitative comparison on Self-Forcing with INT6 KV cache.



Figure 36: Qualitative comparison on Self-Forcing with INT6 KV cache.



Figure 37: Qualitative comparison on Self-Forcing with INT4 KV cache.




Figure 38: Qualitative comparison on Self-Forcing with INT4 KV cache.




Figure 39: Qualitative comparison on Self-Forcing with INT4 KV cache.




Figure 40: Qualitative comparison on LongLive with INT6 KV cache.




Figure 41: Qualitative comparison on LongLive with INT6 KV cache.

Prompt: In an ancient library, books float and glow as they drift through the air, occasionally landing softly on the tables, where curious individuals reach out to read their contents.




Figure 42: Qualitative comparison on LongLive with INT6 KV cache.

Prompt: A goat operating a food truck, serving gourmet grilled cheese sandwiches to a line of animals.



Figure 43: Qualitative comparison on LongLive with INT6 KV cache.

Prompt: A bird landing on water and turning into a fish.



Figure 44: Qualitative comparison on LongLive with INT4 KV cache.



Figure 45: Qualitative comparison on LongLive with INT4 KV cache.

H Evaluation Prompts Adapted from MovieGen

For the MovieGen-based evaluation, we use a fixed subset of 200 prompts randomly sampled from the original **MovieGenVideoBench** prompt set. We adopt this reduced prompt set purely for efficiency, so that all methods can be evaluated under the same protocol with substantially lower inference cost and turnaround time. The sampled 200 prompts are kept identical across all experiments, models, and quantization settings to ensure fair comparison. No manual prompt selection or prompt-specific tuning is applied after sampling.

Table 8: Complete list of the 200 prompts used for MovieGen-based evaluation. The original line number is 1-based in the full MovieGenVideoBench prompt file.

ID	Original Line	Prompt
0	621	A person rowing a boat through a river of liquid gold, with shimmering banks.
1	117	Over the shoulder shot of a woman running and watching a rocket in the distance.
2	705	A trampoline surface bending under the weight of a person jumping on it, then springing back up as they jump off.
3	338	A woman's face twisted in disgust when she tasted the spoiled food.
4	50	A little man with blocks visiting an art gallery
5	331	A man's eyes filled with tears of frustration after failing the exam.
6	394	a man and a woman in their 20s are dining in a futuristic restaurant materialized out of nanotech and ferrofluids
7	513	A deep underground cave filled with glowing crystals and hidden treasures.
8	172	A spaceship being pulled into a blackhole.
9	98	A middle-aged sad bald man becomes happy as a wig of curly hair and sunglasses fall suddenly on his head.
10	850	A tracking shot following a skateboarder performing tricks down a city street, keeping pace with their fluid movements.
11	906	a toy robot wearing purple overalls and cowboy boots taking a pleasant stroll in Mumbai India during a colorful festival
12	923	a woman wearing blue jeans and a white t shirt taking a pleasant stroll in Mumbai India during a beautiful sunset
13	378	Gold coins are falling out when elevator door opens
14	756	A low-angle view of a majestic lion standing on a rocky outcrop, looking regal and powerful against the horizon.
15	984	an old man wearing blue jeans and a white t shirt taking a pleasant stroll in Antarctica during a colorful festival
16	305	A gardener planting seeds in a garden bed, their hands gently pressing the soil over the seeds.
17	440	A princess blowing out birthday candles on a cake.
18	920	a toy robot wearing a green dress and a sun hat taking a pleasant stroll in Antarctica during a beautiful sunset
19	834	A push-in through an open window, entering a cozy room lit by the warm glow of a fireplace.
20	901	a toy robot wearing blue jeans and a white t shirt taking a pleasant stroll in Johannesburg South Africa during a winter storm
21	875	Macro shot. Man in an antique scuba helmet with dark glass walking out of a flower
22	991	an old man wearing purple overalls and cowboy boots taking a pleasant stroll in Johannesburg South Africa during a winter storm
23	738	A single ice cube placed in a warm drink, slowly melting and sending gentle ripples through the liquid as it transforms.
24	407	A surreal collage of a whirlwind of colorful fabrics and clothing items, fluttering and swirling in mid-air. The scene is dynamic and fashionable, with vibrant textile patterns. A sense of motion and style create a visually striking and complex scene. Pitch black background.
25	883	A mother dog gently picks up a piece of meat and carefully places it in her puppy's bowl, her eyes filled with warmth and care as she watches her little one eat.
26	14	A beautiful homemade video showing the people of Lagos, Nigeria in the year 2056. Shot with a mobile phone camera.
27	785	A pan left across a serene beach at sunrise, moving from the darkened shore to the brightening horizon.
28	231	The person's nose wrinkled in distaste as he smelled something unpleasant.
29	496	A quaint village nestled in a valley is surrounded by blooming cherry blossoms, with petals drifting through the air as villagers go about their daily activities, adding life to the scene.
30	424	Close-up of a man eating an apple.
31	701	A rubber band being stretched to its maximum length and then released, snapping back to its original shape.
32	403	Cinematic trailer for a group of adventurous puppies exploring ruins in the sky
33	908	a toy robot wearing purple overalls and cowboy boots taking a pleasant stroll in Johannesburg South Africa during a beautiful sunset
34	828	A zoom-in on a drop of morning dew on a leaf, showing the reflection of the surrounding world within it.
35	927	a woman wearing blue jeans and a white t shirt taking a pleasant stroll in Johannesburg South Africa during a colorful festival
36	933	a woman wearing purple overalls and cowboy boots taking a pleasant stroll in Mumbai India during a colorful festival
37	177	A softball player sliding safely into second base.

Continued on next page

Table 8 – continued from previous page

ID	Original Line	Prompt
38	35	Five gray wolf pups frolicking and chasing each other around a remote gravel road, surrounded by grass. The pups run and leap, chasing each other, and nipping at each other, playing.
39	301	A person is rapidly typing on a keyboard
40	313	A woman is search her bag trying to find something.
41	398	a man BASE jumping over tropical hawaii waters. His pet macaw flies alongside him
42	728	A crystal ball sitting on a table, with sunlight streaming through it and casting a circle of rainbow colors on the floor.
43	892	A girl is peeling an orange
44	173	a real girl frantically running a dense forest with bushes, trees, in rainy day, the animals are running after her and she is screaming and shouting
45	843	A pull-out from a child's hands holding a small seashell, moving back to show the beach and the waves around them.
46	590	A panda in a scientist's lab coat, conducting experiments with beakers and test tubes.
47	831	A push-in through a garden archway, revealing a secret, tranquil garden filled with blooming flowers.
48	631	A person holding an umbrella that turns rain into colorful confetti.
49	622	A person playing a harp strung with rainbows, creating music that colors the air.
50	260	In a vibrant theater, a magician in dazzling attire stands center stage, pulling a comically oversized rubber chicken from an ornate, old-fashioned box. His costume shimmers under the stage lights, adding to the spectacle. The crowd erupts in laughter and applause, their faces filled with joy and amazement. The magician's expression hints at mischievous delight as he holds up the rubber chicken, his performance bringing cheer to the audience.
51	296	A man is putting a ring on a woman's finger
52	783	An aerial shot of a field of blooming wildflowers, creating a patchwork of colors in the landscape.
53	750	An arc shot around a couple standing under a cherry blossom tree, petals falling around them as they embrace.
54	673	A cyclist accelerating out of the saddle during a steep climb.
55	968	an adorable kangaroo wearing a green dress and a sun hat taking a pleasant stroll in Mumbai India during a beautiful sunset
56	460	A video of a diver creating bubbles underwater, with bubbles rising and interacting with each other.
57	151	A chimp on the sidewalk doing a backflip on a skateboard.
58	752	An arc shot around a lone tree in a vast, foggy field at dawn, revealing the changing light and shadows.
59	279	A bunny puts the bright moon on its back and flies into the distance.
60	409	A dynamic motion shot of a broom morphing surreal and magically into a peacock. The handle shortens and curves into a slender neck, the bristles fanning out into a magnificent tail. The camera moves around as vibrant colors and eye-shaped patterns emerge on the expanding feathers. A small head forms at the top, complete with a delicate crest. The transformation completes as the peacock proudly displays its newly formed plumage.
61	134	a cactus with googly eyes dancing in the breeze
62	306	A pair of hands skillfully knitting a colorful scarf, the yarn winding through their fingers with each stitch.
63	283	A car crashes into a barrier at high speed.
64	323	A man is sipping a hot cup of coffee, steam rising from the mug.
65	91	A cinematic wide portrait of a man with his face lit by the glow of a TV.
66	777	An over-the-shoulder shot of a photographer adjusting their camera, framing a beautiful sunset scene.
67	374	A man is running in the forest and transforms into a wolf.
68	652	A person on a hoverboard colliding with a wall, the board stopping abruptly.
69	861	Camera tracking shot. An early 19th century scuba diver with a huge iron helmet and an iron body suit lounges on an antique lawn chair. The light is diffused and gray, casting soft shadows along the scene. The diver brings a martini glass to his helmet and puts it back down. The year is 1912. The diver is in a grassy tree-filled park. People in period-accurate dress mill around, wearing long dresses and suits, holding parasols. The diver's suit is burnished and old, held together with rusted bolts. The diver tips the martini toward his helmet and clinks the glass against the glass. The scene should feel serene and beautiful, evoking the feeling of an impressionist painting.
70	114	Aerial shot of the ocean. a maelstrom forms in the water swirling around until it reveals the fiery depths below.
71	604	A person playing chess with a robot on a floating platform above the ocean.
72	910	a toy robot wearing purple overalls and cowboy boots taking a pleasant stroll in Johannesburg South Africa during a winter storm
73	790	A truck left through a bustling city street, following the flow of traffic and pedestrians during rush hour.
74	965	an adorable kangaroo wearing purple overalls and cowboy boots taking a pleasant stroll in Antarctica during a beautiful sunset
75	42	This close-up shot of a chameleon showcases its striking color changing capabilities. The background is blurred, drawing attention to the animal's striking appearance.
76	299	A woman is ascending to the sky from the ground
77	369	A paraglider descending to a landing zone.
78	137	A police helicopter hovers above a high-speed chase, guiding officers on the ground to apprehend a suspect.
79	385	a giant duck walks through the streets in Boston
80	832	A push-in towards a lone figure standing at the edge of a cliff, overlooking a vast, fog-covered valley.
81	146	A cook flipping pancakes on a griddle.
82	139	a woman jumps over a bear
83	784	An aerial view of a snow-covered mountain range, with the peaks and valleys forming intricate patterns in the snow.
84	978	an old man wearing blue jeans and a white t shirt taking a pleasant stroll in Mumbai India during a colorful festival

Continued on next page

Table 8 – continued from previous page

ID	Original Line	Prompt
85	297	A man is playing the drums under the water
86	308	A person using a screwdriver to assemble a piece of furniture, carefully tightening each screw.
87	220	The person's lips curled up in a sly grin as he shared a secret joke.
88	510	A bustling ancient marketplace with merchants selling spices and fabrics.
89	408	A dynamic motion shot of a lamp transforming into a flamingo. The curved neck of the lamp elongates, its shade flattening into a delicate head. The camera circles as the base splits into two spindly legs, the bulb socket becoming a beak. Pink hues wash over the metal surface, transforming into soft feathers. The power cord coils and disappears as the transformation completes, revealing a graceful flamingo balancing on one leg.
90	207	A man is eating spaghetti
91	821	A pedestal down from the top of a waterfall, descending to show the pool of water and mist at its base.
92	104	A slow cinematic push in on an ostrich standing in a 1980s kitchen.
93	77	A FPV shot zooming through a tunnel into a vibrant underwater space.
94	566	A fish driving a tiny submarine, exploring an underwater city.
95	443	A woman drinking water from a glass, her lips touching the rim.
96	107	A cyclone of broken glass in an urban alleyway. dynamic movement.
97	325	A singer is belting out a high note on stage.
98	81	Zooming in hyper-fast to a dandelion to reveal macro dream-like abstract world.
99	737	The slow crystallization of a water droplet turning into ice on a frosty morning, with delicate patterns forming across its surface.
100	155	A chef drizzling sauce onto a plate with precision.
101	856	A little boy is riding a dragon in the sky to a castle
102	400	a dark neon rainforest aglow with fantastical fauna and animals
103	452	A young boy is playing a harmonica at sunset, with his dog sitting quietly beside him, listening.
104	286	The tall skyscrapers in Hong Kong suddenly transform into a moving Gundam robot, cinematic CGI
105	969	an adorable kangaroo wearing a green dress and a sun hat taking a pleasant stroll in Mumbai India during a colorful festival
106	333	A man sighed in relief as the doctor delivered the good news.
107	390,695	pov footage of an ant navigating the inside of an ant nest
108	836	A zoom-out from a detailed shot of an intricate snowflake, pulling back to show a snowy landscape.
109	228	The person's mouth twisted in disgust as he tasted something bitter.
110	232	The person's forehead furrowed in concern as he listened to a friend's problems.
111	803	A truck right through a tranquil garden, moving past blooming flowers, trees, and a small fountain.
112	420	A marble goes through a glass cup, breaking it into pieces.
113	937	a woman wearing purple overalls and cowboy boots taking a pleasant stroll in Johannesburg South Africa during a winter storm
114	51	A white cat driving in a car through a busy downtown street with tall buildings and pedestrians in the background
115	335	A man looked away in shame when confronted with his wrongdoing.
116	499	Robots move efficiently through a futuristic laboratory, adjusting holographic displays and conducting experiments, while scientists observe and interact with the high-tech equipment.
117	357	A high-speed video of a soap bubble popping, with the soapy liquid dispersing in all directions.
118	918	a toy robot wearing a green dress and a sun hat taking a pleasant stroll in Johannesburg South Africa during a colorful festival
119	699	A magician takes off his performing mask.
120	271	Penguins roller skating
121	879	An astronaut fighting a monster
122	493	In an ancient library, books float and glow as they drift through the air, occasionally landing softly on the tables, where curious individuals reach out to read their contents.
123	164	A chef skillfully tossing a salad in a bowl.
124	891	A mother is peeling an apple for her daughter
125	106	Zooming through a dark forest with neon light flora lighting up.
126	565	A horse wearing roller skates, gracefully gliding through a city park.
127	215	The person walked into the room, his face lighting up with a warm smile.
128	193	A pair of lovebirds preening each other's feathers.
129	634	A person skydiving from a hot air balloon into a sea of clouds.
130	46	Tiltshift of a construction site filled with workers, equipment, and heavy machinery.
131	444	A woman singing softly to a baby, her lips forming gentle words.
132	156	A frog that gets kissed and turns into a chocolate milkshake.
133	581	A kangaroo wearing boxing gloves, sparring with a punching bag in a gym.
134	131	Hyperrealistic monster that closes its mouth
135	458	A time-lapse video of a river flowing through a forest, with changing water levels and currents.
136	797	A pan right across a field of tall grass swaying gently in the wind, with a setting sun in the background.
137	725	Sunlight piercing through a canopy of bamboo, casting long, linear shadows and patches of light on the forest floor.
138	574	A peacock wearing a crown, sitting on a throne and holding court with other animals.
139	90	A FPV drone shot through a castle on a cliff.

Continued on next page

Table 8 – continued from previous page

ID	Original Line	Prompt
140	38	A grandmother with neatly combed grey hair stands behind a colorful birthday cake with numerous candles at a wood dining room table, expression is one of pure joy and happiness, with a happy glow in her eye. She leans forward and blows out the candles with a gentle puff, the cake has pink frosting and sprinkles and the candles cease to flicker, the grandmother wears a light blue blouse adorned with floral patterns, several happy friends and family sitting at the table can be seen celebrating, out of focus. The scene is beautifully captured, cinematic, showing a 3/4 view of the grandmother and the dining room. Warm color tones and soft lighting enhance the mood.
141	774	A close-up of a key turning in a lock, showing the subtle movements of the key and the intricate details of the mechanism as it turns into place.
142	23	A cat waking up its sleeping owner demanding breakfast. The owner tries to ignore the cat, but the cat tries new tactics and finally the owner pulls out a secret stash of treats from under the pillow to hold the cat off a little longer.
143	782	An aerial perspective of a group of dolphins swimming near the surface of a crystal-clear ocean, their movements synchronized.
144	596	A person walking up a staircase made of clouds leading to a floating castle.
145	322	A little child let out a big yawn
146	863	Over the shoulder camera shot. A huge lizard creature sits in a midcentury orange swivel chair. The light is dim and volumetric, casting an eerie glow across the scene. The creature uses its arms to maniacally push buttons on a gigantic control panel. Above the control panel is a panoramic window looking out and down on 1940s new york city. The room should invoke midcentury science fiction aesthetics, like rusty orange colors, bright flashing control buttons, and space-age flair. As the creature continues to quickly push buttons, the New York City scene out of the window moves closer, as though the creature is in a gigantic robot stomping through the city. The scene should give the feeling of frantic action, highlighting the intensity of piloting a giant robot. The scene should take inspiration from midcentury japanese monster films.
147	281	Whilst a man and woman are walking through a city street in a dream, the man shows the woman how to fold the entire street upwards at a 90-degree angle and connect it with the sky. This creates a visually stunning effect, with the buildings and road bending and defying gravity. The scene highlights the limitless possibilities and creativity within the dream world.
148	764	A first-person view of an artist sketching in a notebook, the pencil moving swiftly across the page as the drawing takes shape.
149	393	an alien blending in naturally with new york city, paranoia thriller style, 35mm film
150	687	A woman is playing a soft melody on his lute while sitting by the fountain in the castle courtyard.
151	830	A push-in through a dense crowd at a festival, moving towards a performer on stage who is captivating the audience.
152	82	Internal window of a train moving at hyper-speed in an old European city.
153	32	New York City submerged like Atlantis. Fish, whales, sea turtles and sharks swim through the streets of New York.
154	199	A martial artist breaking a board with a powerful punch.
155	153	A fish walking into a coffee shop and asking for a cup of coffee.
156	872	A computer mouse with legs running on a treadmill
157	206	a woman eating ice scream
158	63	Carefully pouring the milk into the cup, the milk flowed smoothly and the cup was gradually filled with a milky white color
159	184	A soccer goalie making a diving save with outstretched arms
160	171	A penguin flies into the mouth of a blue whale breaking the surface of the water.
161	471	A close-up of a drink being poured over ice, capturing the detailed flow and interaction with the ice cubes.
162	128	A curious cat peering out from a cozy hiding spot.
163	975	an adorable kangaroo wearing a green dress and a sun hat taking a pleasant stroll in Antarctica during a colorful festival
164	257	A bird made of fresh oranges rushes out of the orange
165	570	A bear wearing a superhero cape, flying through the sky over a bustling city.
166	718	Light filtering through a canopy of autumn leaves, casting warm, dappled patterns of yellow, orange, and red onto the ground.
167	654	A person on a scooter colliding with a park bench, the scooter tipping over.
168	677	A bobsled team accelerating down an icy track.
169	816	A tilt-down from the towering peak of a mountain to the winding path leading up to it.
170	255	Camera tracking shot. A man walking down a city street, holding a coffee cup in his hand. He is wearing a dark suit and red tie.
171	379	A rose is growing out of a stone
172	886	A mother bird feeding her chicks in the nest, delicately placing food into their wide-open beaks as they chirp eagerly.
173	217	The person's eyebrows furrowed in concentration as he worked on a puzzle.
174	218	The person's mouth dropped open in surprise as he watched a magic trick.
175	929	a woman wearing blue jeans and a white t shirt taking a pleasant stroll in Antarctica during a beautiful sunset
176	578	A fox wearing a pirate hat and eyepatch, steering a ship through a stormy sea.
177	823	A pedestal down through a field of sunflowers, showing their tall stalks and bright petals against the sky.
178	337	A man grinned with satisfaction after completing the challenging task.
179	747	A thin sheet of ice on a lake cracking and breaking as the sun warms it, creating a mosaic of shifting patterns.
180	122	A bird landing on water and turning into a fish.
181	384	a photorealistic video of a butterfly that can swim navigating underwater through a beautiful coral reef

Continued on next page

Table 8 – continued from previous page

ID	Original Line	Prompt
182	266	A red panda taking a bite of a pizza
183	202	A child’s face lighting up with joy as they blow out the candles on their birthday cake.
184	191	A person opens a book and turns it upside-down and characters from the book begin to fall out of it.
185	244	The person jumps up and down excitedly, expressing happiness through dance moves.
186	208	A person takes a big bite of a juicy burger, the meat and cheese filling his mouth.
187	935	a woman wearing purple overalls and cowboy boots taking a pleasant stroll in Johannesburg South Africa during a beautiful sunset
188	431	A man pushes away a huge stone with superhuman strength.
189	109	Highly detailed close up of a bacteria.
190	573	A goat operating a food truck, serving gourmet grilled cheese sandwiches to a line of animals.
191	472	A mesmerizing video of a whirlpool forming in a sink as water drains.
192	97	An astronaut walking between stone buildings.
193	610	A person doing yoga on top of a giant lily pad in the middle of a serene pond.
194	636	A person riding a giant tortoise through a desert of glass sand.
195	614	A person playing hopscotch on the rings of Saturn.
196	938	a woman wearing purple overalls and cowboy boots taking a pleasant stroll in Antarctica during a beautiful sunset
197	844	A pull-out from a dancer’s feet moving gracefully, expanding to show the entire stage and audience.
198	463	A close-up of paint being mixed, showing the detailed interaction of colors and textures.
199	442	A person chewing on a pencil while deep in thought.

I Broader Societal Impact

The proposed VAR-Q framework significantly reduces the memory footprint and computational overhead of deploying large-scale visual autoregressive models. By compressing the KV cache by up to 78.5% with negligible quality degradation, this work democratizes access to state-of-the-art image and video generation. Positive societal impacts include empowering researchers, independent creators, and under-resourced institutions to utilize high-fidelity visual AI on commodity hardware. This can foster innovation in creative industries, education, and accessibility tools.

However, by drastically lowering the computational barrier to generating highly realistic synthetic media, this technology simultaneously exacerbates the risk of malicious applications. Bad actors can leverage these efficiency gains to mass-produce Deepfakes and realistic synthetic videos for disinformation campaigns, fraud, or non-consensual explicit content. Furthermore, because VAR-Q makes these models cheaper and faster to run, any inherent biases within the foundational models’ training distributions (e.g., demographic misrepresentations) could be propagated more widely and rapidly, raising fairness concerns.

To mitigate these negative impacts, we advocate that the deployment of highly efficient visual generators be coupled with robust provenance techniques, such as invisible watermarking and cryptographic media authentication, to ensure synthetic content remains identifiable. Additionally, developers of foundational visual AR models should consider gated access, strict licensing, or API-monitored deployments to prevent the automated, large-scale generation of harmful content, ensuring the broader community can safely benefit from the efficiency gains of methods like VAR-Q.