

Interleaved Latent Visual Reasoning with Selective Perceptual Modeling

Shuai Dong^{1,2}, Siyuan Wang^{3*}, Xingyu Liu¹,
Chenglin Li^{2,5}, Haowen Hou^{2,6}, Zhongyu Wei^{2,4*}

¹China University of Geosciences, Wuhan ²Shanghai Innovation Institute

³University of Southern California ⁴Fudan University

⁵Zhejiang University ⁶Shanghai Jiao Tong University

{dongshuai_iu, liuxingyu}@cug.edu.cn, sw_641@usc.edu

22351307@zju.edu.cn, haowenhhou@outlook.com, zywei@fudan.edu.cn

Abstract

Interleaved reasoning paradigms enhance Multimodal Large Language Models (MLLMs) with visual feedback but are hindered by the prohibitive computational cost of re-encoding pixel-dense images. A promising alternative, latent visual reasoning, circumvents this bottleneck yet faces limitations: methods either fail to capture intermediate state evolution due to single-step, non-interleaved structures, or sacrifice precise perceptual modeling by over-compressing features. We introduce Interleaved Latent Visual Reasoning (ILVR), a framework that unifies dynamic state evolution with precise perceptual modeling. ILVR interleaves textual generation with latent visual representations that act as specific, evolving cues for subsequent reasoning. Specifically, we employ a self-supervision strategy where a momentum teacher model selectively distills relevant features from ground-truth intermediate images into sparse supervision targets. This adaptive selection mechanism guides the model to autonomously generate context-aware visual signals. Extensive experiments on multimodal reasoning benchmarks demonstrate that ILVR outperforms existing approaches, effectively bridging the gap between fine-grained perception and sequential multimodal reasoning. The code is available at <https://github.com/XD111ds/ILVR>.

1 Introduction

Multimodal Large Language Models (MLLMs) (Li et al., 2024; Bai et al., 2025a; Wang et al., 2025c; Li et al., 2025c; Wang et al., 2025a, 2026b) have demonstrated remarkable capabilities in bridging the gap between vision and language. Capitalizing on the reasoning prowess of Large Language Models (LLMs) (Chen et al., 2025b,a; Wang et al., 2026a; Zhao et al., 2025; Xu et al., 2026; Feng et al., 2025), recent works have successfully adapted

Chain-of-Thought (CoT) methodologies to the multimodal domain (Zhang et al., 2023; Bai et al., 2025b; Huang et al., 2025a; Wei et al., 2022; Zhu et al., 2025, 2026; Chen et al., 2026). This enables models to decompose complex visual tasks into sequential intermediate steps, achieving sophisticated reasoning grounded in visual content.

Recent work explores interleaved image-text reasoning by injecting intermediate visual images within textual CoTs to enhance multimodal understanding and planning (Shao et al., 2024b). These approaches generally fall into two paradigms. The first uses external tools to statically manipulate the input image, e.g., highlighting key regions (Fu et al., 2025), drawing auxiliary lines (Hu et al., 2024), or shifting image styles (Liu et al., 2025), to improve fine-grained perception. While relying on a single visual state, it cannot model evolving scenarios or simulate action outcomes crucial for sequential tasks (Li et al., 2025a). The second paradigm addresses this employing a unified model to dynamically visualizing imagined intermediate or future states (Chern et al., 2024; Deng et al., 2025). However, integrating visual generation and reasoning into a unified model often degrades reasoning performance. More critically, both paradigms incur high computational cost from iteratively re-encoding pixel-dense images, severely hindering multi-step reasoning.

Inspired by latent reasoning in LLMs (Shen et al., 2025; Hao et al., 2024), the latent visual reasoning paradigm replaces explicit images with latent representations to avoid costly pixel-level processing. However, current methods face two major limitations. First, most adopt a single-step, non-interleaved design. For instance, LVR (Li et al., 2025b) and Mirage (Yang et al., 2025) generate latent representations only once, either for a region of the static input image or the final state after all actions, and cannot model intermediate and evolving states during reasoning. In the chess puzzle in

* Corresponding authors.

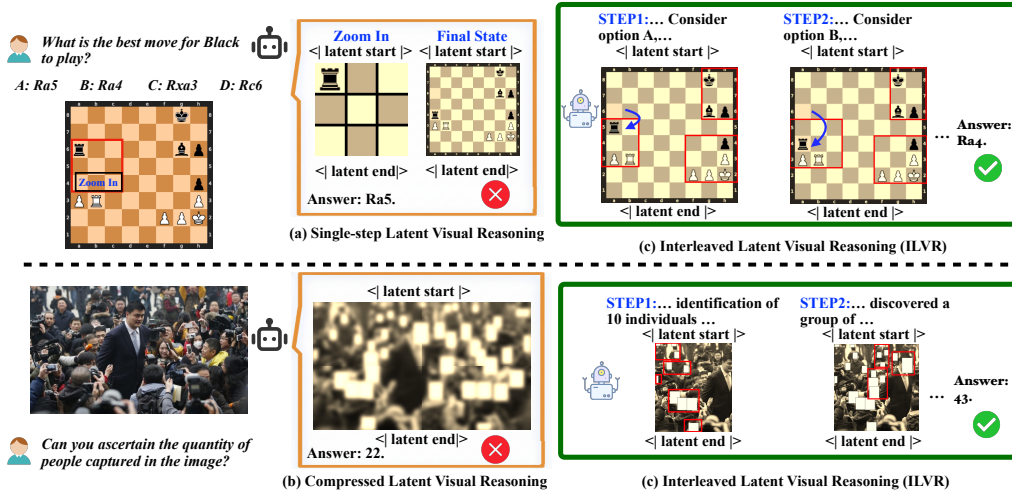


Figure 1: **Comparison of ILVR with prior latent visual reasoning methods.** In the chess puzzle (top row), single-step approaches (a) either capture static initial details (e.g., a zoomed-in rook) or jump to a predicted final state, failing to model the hypothetical states needed to evaluate move options. In the dense counting task (bottom row), methods relying on heavily compressed latent representations (b) lose fine-grained details, resulting in a hallucinated count. In contrast, our ILVR (c) succeeds by interleaving textual reasoning with dynamically updated latent states. Each latent representation provides essential visual cues for subsequent reasoning steps (highlighted in red boxes), unifying dynamic state evolution with precise perceptual modeling to reach the correct answer.

Fig. 1(a), relying on a static zoom-in or a predicted final state is insufficient, as it bypasses step-by-step verification of move legality (e.g., path obstructions), often leading to erroneous predictions. Second, methods like Mirage (Yang et al., 2025) derive latent representations by heavily compressing dense visual features from the entire image into limited latent tokens. As the counting task shown in Fig. 1(b), such over-compression discards crucial perceptual details and leads to hallucination.

To this end, we propose Interleaved Latent Visual Reasoning (ILVR) framework to integrate dynamic latent visual reasoning with selective perceptual modeling. ILVR interleaves reasoning between explicit textual generation and latent visual representations that are continuously updated to capture the most relevant visual cues at each reasoning step. We train the model to learn this interleaved paradigm by approximating ground-truth interleaved image-text trajectories, with textual outputs supervised using cross-entropy loss while latent representations are aligned with selectively extracted features from their corresponding images, which we refer to as *helper images*. Specifically, we employ a momentum teacher model (He et al., 2019), a temporally smoothed copy of the trained model, to selectively extract the most relevant features from *helper images* by aggregating highly attended patches conditioned on the ongoing rea-

soning process. By internalizing this capability, ILVR effectively unifies precise perceptual modeling with dynamic evolution of latent visual states.

In summary, our contributions are threefold:

- We propose Interleaved Latent Visual Reasoning (ILVR), a framework that interleaves explicit token generation with updated latent visual representations, enabling dynamic state evolution.
- We introduce an adaptive selection mechanism that distills the most relevant visual signals from the *helper image* into latent representations at every reasoning step, using a self-supervised strategy guided by a momentum teacher model without requiring external supervision.
- Through extensive experiments on fine-grained visual perception and sequential planning tasks, we demonstrate ILVR’s robust generalization in both in-domain and out-of-distribution (OOD) settings. By operating entirely in latent space, it achieves up to 18× inference speedup over methods requiring costly explicit image generation.

2 Related Work

2.1 Interleaved Image-Text Reasoning

Interleaved image-text reasoning refers to the capability of models to generate intermediate visual feedback (Chern et al., 2024; Li et al., 2025a; Deng et al., 2025), either directly or via external tools (Hu

et al., 2024; Shao et al., 2024b; Su et al., 2025), to enhance their reasoning abilities. Early methods used external tools for static image edits, such as cropping or OCR (Huang et al., 2025b; Zhang et al., 2025a; Wang et al., 2025b), but struggled to model evolving visual states. Recent generative approaches enable models to synthesize intermediate-state images (Chern et al., 2024; Deng et al., 2025), yet they often face a trade-off between generative fidelity and reasoning performance. Crucially, both tool-based and generative paradigms suffer from high computational overhead due to repeated pixel-level encoding of dense visual data.

2.2 Latent Reasoning

To bypass discrete token constraints, latent reasoning performs multi-step inference in continuous hidden space (Shen et al., 2025; Hao et al., 2024; Cheng and Durme, 2024). In the multimodal domain, Mirage (Yang et al., 2025) precedes textual reasoning with a latent representation formed by encoding a problem-specific helper image and aggressively pooling its patch embeddings into highly compressed vectors. LVR (Li et al., 2025b) adopts a similar strategy but isolates key visual cues within a bounding box, generating latent representations of only that targeted region. Contemporaneous with our work, Sketchpad (Zhang et al., 2025b) also explores generating visual latents to elicit reasoning. However, a fundamental limitation plagues these approaches. In their paradigm, a model generates latent representations of a helper image once, and all subsequent steps are confined to pure textual reasoning. This non-interleaved structure inherently renders the visual information static and detached from the evolving reasoning trajectory.

3 Method

In this section, we present Interleaved Latent Visual Reasoning (ILVR) framework that performs reasoning by interleaving explicit textual generation with latent visual representations, as shown in Fig. 2. We first outline the interleaved generation paradigm (Sec. 3.1). We then detail how we construct latent supervision targets by selecting key features from intermediate images (“*helper images*”) within ground-truth interleaved image-text trajectories using a momentum teacher model (Sec. 3.2). Finally, we describe the two-stage training strategy to instill this interleaved latent reasoning ability (Sec. 3.3).

3.1 Interleaved Text-Latent Paradigm

Our framework operates in an interleaved reasoning paradigm where the model autoregressively generates both text tokens and latent visual representations. The reasoning process is structured as a unified sequence \mathcal{S} that alternates between textual tokens and latent segments:

$$\mathcal{S} = [t_{1,1}, \dots, t_{1,M}, \langle |\text{latent_start}| \rangle, z_{1,1}, \dots, z_{1,K}, \langle |\text{latent_end}| \rangle, t_{2,1}, \dots, t_{2,N}, \langle |\text{latent_start}| \rangle, z_{2,1}, \dots, z_{2,K}, \langle |\text{latent_end}| \rangle, \dots] \quad (1)$$

where $t_{i,j}$ denotes discrete text tokens and $z_{i,k}$ represents continuous latent embeddings at reasoning step i . The special tokens $\langle |\text{latent_start}| \rangle$ and $\langle |\text{latent_end}| \rangle$ explicitly delimit the boundaries of latent visual reasoning phases.

During inference, the model generates text tokens as usual. When the model produces a $\langle |\text{latent_start}| \rangle$ token, it switches to a latent generation mode for a fixed length K . In this mode, instead of projecting the hidden state to the vocabulary size to sample a discrete token, the hidden state from the previous timestep \mathbf{h}_t is fed directly as the input embedding for the current timestep, effectively bypassing the discrete embedding lookup $\mathbf{e}_{t+1} = \mathbf{h}_t$. The sequence of K hidden states produced in this loop constitutes the model’s self-generated latent representations. After completing K latent generation, the model generates $\langle |\text{latent_end}| \rangle$ and resumes explicit textual reasoning, utilizing the accumulated latent information as context.

To train the model with this paradigm, we utilize pre-constructed interleaved trajectories formatted as “reasoning text→helper image→reasoning text→helper image...”. We convert each trajectory into a unified supervision sequence by replacing each *helper image* I_i at reasoning step i with a latent segment: a $\langle |\text{latent_start}| \rangle$ followed by K $\langle |\text{latent_pad}| \rangle$ tokens, and terminated by $\langle |\text{latent_end}| \rangle$. The $\langle |\text{latent_pad}| \rangle$ act as placeholders for the critical visual signals extracted from I_i . Thus, the core of our method is to select which visual features from I_i should serve as regression targets to supervise the hidden states generated at these pad positions.

3.2 Interleaved Supervision Construction

To enable the model to generate meaningful latent representations, we employ a teacher model to con-

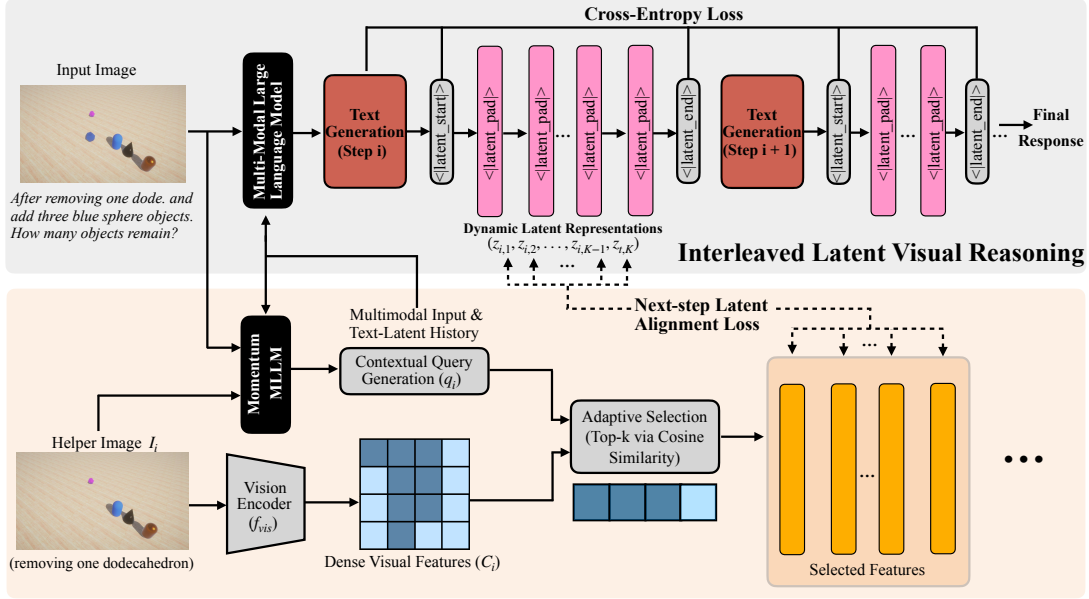


Figure 2: **The Interleaved Latent Visual Reasoning (ILVR) framework.** The model performs multi-step reasoning by interleaving textual generation with dynamically evolving latent visual representations. The momentum teacher model (bottom) utilizes the multimodal inputs and the text-latent history up to reasoning step i to form a contextual query (q_i), which selectively extracts the most relevant visual features (yellow blocks) from the *helper image*. Simultaneously, the trained model (top) generates a sequence of latent representations (pink blocks) interleaved with reasoning text. These latents are supervised via a next-step latent alignment objective that encourages them to match the teacher-selected visual features.

struct high-quality supervision targets for the latent segments. Given the same reasoning context as the model being trained, the teacher processes the *helper image* I_i at reasoning step i and extracts the most relevant visual features as ground-truth latent supervision. Meanwhile, the textual parts are trained using standard explicit text supervision.

Momentum Teacher Model We adopt a self-supervised strategy where the teacher is a momentum model, a temporally smoothed version of the model being trained (the student model). This design keeps the supervision signal stable and well-aligned with the evolving representation space of the student model. The parameters of the momentum model θ_m are updated as an Exponential Moving Average (EMA) of the student parameters θ with a decay factor τ : $\theta_m \leftarrow \tau\theta_m + (1 - \tau)\theta$.

Candidate Visual Feature Generation The goal of the momentum teacher model is to selectively distill the pixel-dense *helper image* into a sparse set of K feature vectors most relevant to the current reasoning step. The teacher first encodes a *helper image* I_i using its frozen vision encoder f_{vis} to obtain a dense pool of patch features:

$$\mathbf{C}_i = f_{\text{vis}}(I_i) = \{\mathbf{c}_{i,j} \in \mathbb{R}^H\}_{j=1}^{P_i}, \quad (2)$$

where H is the hidden dimension and P_i is the number of patches.

However, raw patch features often suffer from varying information density depending on the image resolution. In high-resolution images, individual patches may capture only local textures rather than semantic concepts. To address this, we introduce a spatial aggregation step to adapt the feature density. Specifically, we set a threshold L : if the number of raw patches $P_i \geq L$, we pool features over local spatial windows to form a refined candidate pool \mathbf{C}'_i ; otherwise, we retain the original fine-grained features. Formally,

$$\mathbf{C}'_i = \begin{cases} \text{GroupMean}(\mathbf{C}_i, L), & \text{if } P_i \geq L \\ \mathbf{C}_i, & \text{if } P_i < L \end{cases} \quad (3)$$

where GroupMean aggregates the P_i patch sequence into L semantic units, ensuring that the subsequent selection operates on robust features regardless of input resolution.

Teacher-Guided Selective Perceptual Modeling

The teacher model then identifies the most relevant candidate features from \mathbf{C}'_i as supervision. It constructs a context-aware query \mathbf{q}_i using the same context as the student model, including the multimodal inputs and the reasoning history up to step

i . By computing cosine similarity between \mathbf{q}_i and each feature in \mathbf{C}'_i , the teacher selects the top- K features to form the supervision set \mathbf{Z}_i .

To construct the query \mathbf{q}_i , we do not apply naive average pooling over the entire context that would weaken critical signals. Instead, we separately process the input text, input image, and reasoning history, with the first two forming a global intent vector and the last providing local reasoning context. For dense input text, we apply mean pooling over their final-layer hidden states to obtain \mathbf{r}_{txt} . For sparse input images, we compute text-guided attention over image to selectively emphasize informative regions, yielding \mathbf{r}_{img} . The global intent vector is obtained by averaging the representation of input text and image as $\mathbf{u} = \frac{1}{2}(\mathbf{r}_{\text{txt}} + \mathbf{r}_{\text{img}})$.

To capture evolving reasoning dynamics, we incorporate the reasoning history up to step i by averaging the final-layer hidden states of all textual rationales from step 1 to step i as $\mathbf{q}_{[1,i]}^{\text{text}}$ and all latent rationales from step 1 to step $i - 1$ as $\bar{\mathbf{z}}_{[1,i-1]}$. The final query \mathbf{q}_i is constructed by fusing the global intent, the current textual rationale, and, when available, the previous latent state as:

$$\mathbf{q}_i = \text{Average} \left(\mathbf{u}, \mathbf{q}_{[1,i]}^{\text{text}}, \mathbb{I}[i > 1] \cdot \bar{\mathbf{z}}_{[1,i-1]} \right). \quad (4)$$

Finally, the teacher computes cosine similarities between \mathbf{q}_i and each candidate feature in the refined pool \mathbf{C}'_i , and selects the top- K most relevant features to form the supervision set \mathbf{Z}_i .

3.3 Two-stage Learning

We train the model using a two-stage pipeline that progressively instills interleaved latent reasoning capabilities using constructed supervision.

Stage 1: Interleaved Text-Latent Joint Supervision In the first stage, we enforce precise perceptual modeling. The teacher-selected features \mathbf{Z}_i are used as teacher-forced inputs and supervision for the $K < |\text{latent_pad}|$ tokens at reasoning step i . The model is optimized with a joint loss: a standard cross-entropy loss \mathcal{L}_{CE} for text tokens, and a latent alignment loss that forces the student’s hidden state \mathbf{h}_{t-1} to match the teacher’s selected feature \mathbf{z}_t .

$$\mathcal{L}_{\text{S1}} = \mathcal{L}_{\text{CE}}(\mathcal{X}_{\text{text}}) + \lambda_{\text{sim}} \cdot \frac{1}{\sum_i K} \sum_i \sum_{t \in \mathcal{T}_i} \left(1 - \cos(\mathbf{h}_{t-1}, \mathbf{z}_t) \right), \quad (5)$$

where \mathcal{T}_i is the indices of the latent tokens at reasoning step i , $\mathcal{X}_{\text{text}}$ represents all textual tokens, and λ_{sim} balances the two objectives.

Stage 2: Text-Only Supervision with Latent Relaxation In the second stage, we relax the strict alignment constraint to allow the model to freely explore the latent reasoning process and use latent states as internal priors for subsequent tokens. We remove the latent alignment loss and feed self-generated hidden state as the input for the next latent position, optimizing only the textual part.

$$\mathcal{L}_{\text{S2}} = \mathcal{L}_{\text{CE}}(\mathcal{X}_{\text{text}}), \quad (6)$$

4 Experiments

4.1 Experimental Setup

Datasets We evaluate ILVR under both in-distribution (IID) and out-of-distribution (OOD) settings. IID evaluation follows the standard splits of COMT (Cheng et al., 2024) and VSP (Wu et al., 2024). For OOD evaluation, models are trained on a 10k subset of Zebra-CoT (Li et al., 2025a) spanning scientific, visual logic, and 3D reasoning tasks, then evaluated on EMMA BENCH (Hao et al., 2025), VisuLogic (Xu et al., 2025), and held-out Zebra-CoT 2D visual reasoning tasks. The OOD setting is characterized by task-type mismatch: Zebra-CoT science focuses on physics and graph problems, whereas EMMA BENCH additionally covers mathematics, chemistry, and coding; Zebra-CoT visual logic centers on maze- and game-like tasks, while VisuLogic targets positional, quantitative, and stylistic reasoning. Controlled comparisons are conducted on both Qwen2.5-VL-7B and Qwen3-VL-8B (Bai et al., 2025a) backbones to demonstrate generalization.

Baselines We compare ILVR against three categories of baselines: (1) Standard baselines, including Zero-shot, direct answer fine-tuning (Direct-FT) and CoT fine-tuning (CoT-FT). (2) Single-step latent reasoning methods, i.e., Mirage (Yang et al., 2025) and LVR (Li et al., 2025b). We report Mirage as the representative baseline in main tables, as LVR operates on pre-defined bounding boxes and models only static visual states, making it incompatible with dynamically evolving reasoning scenarios in our benchmarks.

We report additional experiments against LVR on bounding-box-annotated data in Appendix A. (3) SOTA reasoning models (OOD only) with

Methods	Paradigm	COMT				VSP	COMT				VSP		
		Creation	Deletion	Selection	Update		Avg.	Creation	Deletion	Selection		Update	Avg.
Backbones		<i>Qwen2.5-VL-7B</i>					<i>Qwen3-VL-8B</i>						
<i>Standard Baselines</i>													
Zero-shot	Direct Ans.	68.0	38.0	35.0	14.0	38.8	6.0	89.0	28.0	10.0	21.0	37.0	19.0
Direct-FT	Direct Ans.	52.0	60.0	51.0	49.0	53.0	72.0	89.0	67.0	49.0	53.0	64.5	60.8
CoT-FT	Text CoT	80.0	52.0	45.0	46.0	55.8	47.0	83.0	62.0	49.0	44.0	59.8	61.8
<i>Latent Reasoning</i>													
<i>Stage 1: Latent Alignment</i>													
Mirage	Single-step	53.0	54.0	45.0	42.0	48.5	65.8	81.0	58.0	43.0	50.0	58.0	71.3
ILVR (Ours)	Interleaved	69.0	66.0	46.0	47.0	57.0	77.3	84.0	63.0	57.0	55.0	64.8	75.0
<i>Stage 2: Latent Relaxation</i>													
Mirage	Single-step	65.0	62.0	47.0	50.0	56.0	76.0	84.0	66.0	54.0	57.0	65.3	78.3
ILVR (Ours)	Interleaved	71.0	68.0	53.0	51.0	60.8	81.5	87.0	73.0	60.0	62.0	70.5	82.8

Table 1: **IID performance comparison on COMT and VSP.** Creation, Deletion, Selection, and Update denote COMT subtasks. Backbone differences are explicitly indicated in the *Standard Baselines* header row, while latent reasoning methods are evaluated under the same column layout. “Direct Ans.” and “Text CoT” denote direct answer generation and text-only CoT, respectively. **Bold** indicates the best result. Accuracy (%) is reported.

extensive reinforcement learning (RL), including VisionR1 (Huang et al., 2025a) and PixelReasoner (Su et al., 2025). We also compare against Bagel-Zebra (Li et al., 2025a), a Bagel (Deng et al., 2025) variant fine-tuned on the complete 180k Zebra-CoT dataset to enhance reasoning capabilities. We use official checkpoints for specialized models and fine-tune all others on the same datasets as ILVR with identical implementation settings for both backbones. Notably, we omit Bagel-Zebra from Zebra-CoT OOD evaluation because it was trained on the full dataset, making the test set in-distribution and unsuitable for OOD comparison.

Implementation Details. We optimize all models using AdamW with a learning rate of $1e-5$, a cosine learning-rate scheduler, and a fixed random seed of 42. For IID tasks (COMT and VSP), training is conducted for 15 epochs. In the OOD setting, models are fine-tuned on the 10k Zebra-CoT subset for 2 epochs with a target group size $L = 784$ for adaptive feature grouping. Across all experiments, we set the latent token size to $K = 8$, the alignment weight $\lambda_{\text{sim}} = 1$, and the EMA decay to $\tau = 0.999$. Qwen2.5-VL-72B serves as the judge model for open-ended evaluations.

4.2 Main Results

Table 1 reports in-distribution results on COMT and VSP. ILVR consistently outperforms standard baselines, including Zero-shot, Direct-FT, CoT-FT, and the single-step latent method Mirage across both backbones. With the Qwen2.5-VL-7B backbone, ILVR achieves 60.8% accuracy on COMT and 81.5% on VSP, surpassing Mirage by 4.8% and

5.5%, respectively. When scaling to the stronger Qwen3-VL-8B, ILVR maintains this significant advantage, reaching 70.5% on COMT (+5.2%) and 82.8% on VSP (+4.5%). These results confirm that interleaved text-latent reasoning yields consistent and backbone-agnostic benefits, leading to stronger overall performance.

Table 2 shows that these gains transfer to OOD evaluation where ILVR consistently outperforms standard baseline and the latent method Mirage across all benchmarks, achieving an average improvement of 3.2% over Mirage. ILVR also surpasses recent state-of-the-art multimodal reasoning models VisionR1 and PixelReasoner despite their use of more stochastic reinforcement learning. In terms of average accuracy, ILVR exceeds VisionR1 by 8.0% and PixelReasoner by 5.9%. We further compare ILVR with Bagel-Zebra trained on the full Zebra-CoT dataset with 180k samples. ILVR is trained on only a 10k subset yet still outperforms Bagel-Zebra on EMMA BENCH and VisuLogic. Results on Zebra-CoT OOD are omitted for Bagel-Zebra, as its test split becomes in-distribution.

4.3 Ablation Study

We conduct ablations based on Stage 1 training and OOD benchmarks to investigate the contribution of interleaved reasoning and selective perception in our ILVR, and analyze the impact of different latent sizes K and the alignment weights λ_{sim} .

Interleaved & Selective Design. Table 3 shows that replacing mean pooling with teacher-guided selective perceptual modeling improves the overall accuracy from 31.5% to 32.4%. Adding the inter-

Model	Paradigm	EMMA BENCH					VisuLogic				Zebra-CoT (OOD)			Total
		Chem.	Code	Math	Phys.	Avg.	Pos.	Quant.	Style	Avg.	Jigsaw	Search	Avg.	
<i>SOTA Reasoning Models (Official Checkpoints - No Task-specific Fine-tuning)</i>														
VisionR1	Reasoning	15.0	30.0	32.0	20.0	24.3	18.0	13.0	14.0	15.2	25.0	65.0	45.0	29.5
PixelReasoner	Tool-use	19.0	22.0	26.0	27.0	23.5	18.0	16.0	29.0	23.4	18.0	73.0	45.5	31.6
Bagel-Zebra	Unified	23.0	28.0	29.0	32.0	28.0	28.0	39.0	21.0	28.9	-	-	-	-
<i>Standard Baselines (Fine-tuned on Zebra-CoT 10k subset)</i>														
Zero-shot	Direct Ans.	18.0	25.0	28.0	33.0	26.0	29.0	24.0	27.0	26.6	23.0	65.0	44.0	32.8
Direct-FT	Direct Ans.	16.0	27.0	28.0	32.0	25.8	25.0	23.0	23.0	23.8	17.0	73.0	45.0	32.3
CoT-FT	Text CoT	21.0	26.0	33.0	31.0	27.8	27.0	23.0	28.0	25.9	21.5	68.5	45.0	33.6
<i>Latent Reasoning (Fine-tuned on Zebra-CoT 10k subset)</i>														
<i>Stage 1: Latent Alignment</i>														
Mirage	Single-step	13.0	21.0	30.0	37.0	25.3	25.0	24.0	21.0	23.4	16.0	71.0	43.5	31.5
ILVR (Ours)	Interleaved	23.0	26.0	34.0	35.0	29.5	26.0	23.0	24.0	24.5	20.5	74.5	47.5	34.8
<i>Stage 2: Latent Relaxation</i>														
Mirage	Single-step	15.0	25.0	35.0	33.0	27.0	24.0	26.0	30.0	26.6	20.0	74.5	47.3	34.3
ILVR (Ours)	Interleaved	31.0	35.0	34.0	33.0	33.3	27.0	30.0	31.0	29.3	22.5	73.0	47.8	37.5

Table 2: **Generalization evaluation on three OOD benchmarks: EMMA BENCH, VisuLogic, and Zebra-CoT.** The table compares state-of-the-art RL-based reasoning models using official checkpoints, standard baselines fine-tuned on Zebra-CoT (10k subset), and our ILVR. **Bold** indicates the best result within each column. As Bagel-Zebra is trained on the full Zebra-CoT dataset (180k), making the Zebra-CoT test set in-distribution for this model. We thus omit its score in the OOD Zebra-CoT column to ensure a fair comparison. Accuracy (%) is reported.

Reasoning Paradigm	Perception Mechanism	Accuracy			
		VisLog	EMMA	Zebra	Total
Single-step	Mean Pooling (Mirage)	23.4	25.3	43.5	31.5
Single-step	Selective /	24.1	26.3	44.5	32.4
Interleaved	Selective (ILVR)	24.5	29.5	47.5	34.8

Table 3: **Ablation of interleaved paradigm and selection perception mechanism** against mean pooling and single-step setup (Mirage). Accuracy (%) is reported.



Figure 3: **Impact of latent size K .** Performance trends across VisuLogic, EMMA, and Zebra-CoT, as well as the overall average, as the number of latent tokens K varies. λ_{sim} is fixed at 1.0. $K = 8$ yields the most robust performance across benchmarks.


leaved reasoning paradigm yields further gains to 34.8%. These results suggest that selective perception improves the quality of latent supervision, and interleaved latent updates boost performance by explicitly modeling evolving reasoning states.

λ_{sim}	Accuracy			
	VisLog	EMMA	Zebra	Total
0.1	23.4	25.8	44.0	31.8
0.5	20.0	30.5	45.8	33.3
1 (ILVR)	24.5	29.5	47.5	34.8
2	21.7	27.8	42.5	31.6
10	21.4	27.5	48.5	33.6

Table 4: **Sensitivity to alignment loss weight λ_{sim} .** We report the average accuracy (%) across benchmarks. λ_{sim} represents the relative weight of the alignment loss relative to the text generation loss.

Latent size K . Fig. 3 reports performance under different latent sizes K , where $K = 8$ yields the best overall results. This indicates that a moderate latent budget is sufficient to capture step-specific perceptual evidence, while smaller K limits representational capacity and larger K introduces redundant latent content that weakens step-wise updates. We therefore use $K = 8$ in all experiments.


Alignment weight λ_{sim} . Table 4 reports sensitivity to λ_{sim} and shows that $\lambda_{sim} = 1$ yields the best accuracy. Smaller values weaken latent supervision and perceptual grounding, while larger values over-constrain latent representations and hinder adaptation to subsequent reasoning steps. This supports $\lambda_{sim} = 1$ as an effective trade-off between perceptual alignment and reasoning flexibility.



Navigate the character from a starting point to a goal on a grid map while avoiding obstacles like ice holes.




Please describe the sequence of robot actions from the initial state (the first image) to the goal state (the second image).



Step1: Move towards the piece of bread.

Step2: Grab the bread.



Step3: Move the bread towards pink plate.



Step4: Move away from the bread.

Figure 4: **Visualization of dynamic latent modeling.** Heatmaps depict the Gaussian-smoothed aggregation of relevant image patches for $K = 8$ generated latents. **Top** (Navigation): Latents sequentially track the character’s planned path. **Bottom** (Robotic Manipulation): Visual attention shifts from the object (bread) to the target (plate) during the task. These confirm precise alignment between generated latents and the step-wise reasoning context.

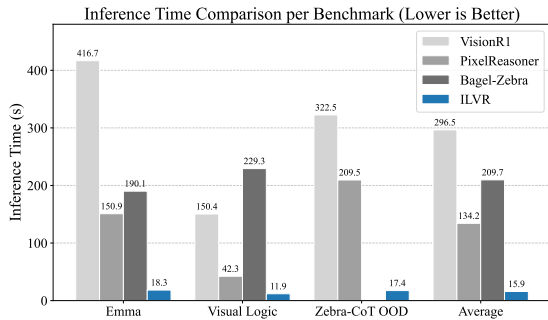


Figure 5: **Comparison of average inference time per sample.** We report the latency averaged across EMMA, VisuLogic, and Zebra-CoT OOD benchmarks.

4.4 Analysis

Efficiency. Fig. 5 reports the average inference time per sample on a same NVIDIA H200 GPU, averaged across EMMA BENCH, VisuLogic, and Zebra-CoT OOD. ILVR achieves substantially lower latency than competing methods, running orders of magnitude faster ($\times 8 \sim \times 18$ speedup) than VisionR1, PixelReasoner, and Bagel-Zebra. The key reason is that ILVR performs multi-step reasoning by updating compact latent states, which bypasses repeated pixel-level processing and intermediate image generation that dominate the runtime of these baselines. These results confirm that ILVR provides an efficient alternative to costly long-context or tool-based reasoning methods.

Qualitative visualization. Fig. 4 visualizes Gaussian-smoothed aggregation of relevant patches derived from attention weights for $K = 8$ generated latents. In the navigation example, attention evolves step by step with the planned actions. Early steps focus on both the goal and nearby ice holes to ensure safe planning, while later steps concentrate almost entirely on the goal once the path is clear. In the robotic manipulation example, attention concentrates on the bread during approaching and grasping, then shifts toward the plate during subgoal placement, and finally moves away after completion. This step-aligned evolution suggests that the latents are conditioned on the evolving reasoning context and provide subgoal-adaptive localized visual cues, which helps subsequent text generation remain grounded in the correct regions.

5 Conclusion

In this paper, we introduce Interleaved Latent Visual Reasoning (ILVR) to unify dynamic state evolution with precise perceptual modeling. Unlike single-step methods that bypass intermediate verification, ILVR interleaves textual generation with evolving latent representations to track reasoning states without costly pixel-level re-encoding. Our momentum teacher-guided selection mechanism distills step-specific visual cues, avoiding feature over-compression. Experiments confirm that ILVR

significantly outperforms single-step latent methods, validating dynamic latent reasoning as a scalable path for multimodal intelligence.

Acknowledgements

The research is supported by the AI for Science Program, Shanghai Municipal Commission of Economy and Informatization (Grant No.2025-GZL-RGZN-BTBX-02028). The project's computational resources are partially supported by CFFF platform of Fudan University.

Limitations

Despite ILVR's robust performance, three limitations remain for future work. First, while theoretically model-agnostic, our experiments currently focus on Qwen-VL backbones; validating the framework across diverse architectures and larger parameter scales is a necessary next step. Second, integrating Reinforcement Learning (RL) to directly optimize latent trajectories could further enhance multi-step planning capabilities. Finally, although attention maps provide insight, the generated latent representations are not directly human-readable; exploring decoding mechanisms to project these states back into pixel space remains an open challenge for better interpretability.

Use of AI Assistants

In adherence to the ACL Publication Ethics Policy, we did not employ AI assistants to generate the initial draft of this paper. We used AI assistants such as GPT-5.2 and Gemini3-Pro exclusively at the sentence level to enhance our writing quality and correct grammatical errors.

References

- Shuai Bai, Keqin Chen, Xuejing Liu, Jialin Wang, Wenbin Ge, Sibao Song, Kai Dang, Peng Wang, Shijie Wang, Jun Tang, Humen Zhong, Yuanzhi Zhu, Mingkun Yang, Zhaohai Li, Jianqiang Wan, Pengfei Wang, Wei Ding, Zheren Fu, Yiheng Xu, and 8 others. 2025a. [Qwen2.5-vl technical report](#). *ArXiv*, abs/2502.13923.
- Sule Bai, Mingxing Li, Yong Liu, Jing Tang, Haoji Zhang, Lei Sun, Xiangxiang Chu, and Yansong Tang. 2025b. [Univg-r1: Reasoning guided universal visual grounding with reinforcement learning](#). *ArXiv*, abs/2505.14231.
- Kedi Chen, Zhikai Lei, Xu Guo, Xuecheng Wu, Siyuan Zeng, Jia-Peng Yin, Yinqi Zhang, Qin Chen, Jie Zhou, Liang He, Qipeng Guo, Kai Chen, and Wei Zhang. 2025a. [Code-driven number sequence calculation: Enhancing the inductive reasoning abilities of large language models](#). *arXiv*, abs/2510.14620.
- Kedi Chen, Dezhao Ruan, Yuhao Dan, Yaoting Wang, Siyu Yan, Xuecheng Wu, Yinqi Zhang, Qin Chen, Jie Zhou, Liang He, Biqing Qi, Linyang Li, Qipeng Guo, Xiaomin Shi, and Wei Zhang. 2025b. [A survey of inductive reasoning for large language models](#). *arXiv*, abs/2510.10182.
- Liang Chen, Weichu Xie, Yiyan Liang, Hongfeng He, Han Zhao, Zhibo Yang, Zhiqi Huang, Haoning Wu, Haoyu Lu, Y. Charles, Yiping Bao, Yu Fan, Guopeng Li, Haiyang Shen, Xuanzhong Chen, Wen Xu, Shuzheng Si, Zefan Cai, Wenhao Chai, and 10 others. 2026. [Babyvision: Visual reasoning beyond language](#). *arXiv*, abs/2601.06521.
- Jeffrey Cheng and Benjamin Van Durme. 2024. [Compressed chain of thought: Efficient reasoning through dense representations](#). *ArXiv*, abs/2412.13171.
- Zihui Cheng, Qiguang Chen, Jin Zhang, Hao Fei, Xiaocheng Feng, Wanxiang Che, Min Li, and Libo Qin. 2024. [Comt: A novel benchmark for chain of multi-modal thought on large vision-language models](#). *ArXiv*, abs/2412.12932.
- Ethan Chern, Jiadi Su, Yan Ma, and Pengfei Liu. 2024. [Anole: An open, autoregressive, native large multimodal models for interleaved image-text generation](#). *ArXiv*, abs/2407.06135.
- Chaorui Deng, Deyao Zhu, Kunchang Li, Chenhui Gou, Feng Li, Zeyu Wang, Shu Zhong, Weihao Yu, Xiaonan Nie, Ziang Song, Shi Guang, and Haoqi Fan. 2025. [Emerging properties in unified multimodal pretraining](#). *ArXiv*, abs/2505.14683.
- Sicheng Feng, Kaiwen Tuo, Song Wang, Lingdong Kong, Jianke Zhu, and Huan Wang. 2025. [Rewardmap: Tackling sparse rewards in fine-grained visual reasoning via multi-stage reinforcement learning](#). *ArXiv*, abs/2510.02240.
- Xingyu Fu, Minqian Liu, Zhengyuan Yang, John Corring, Yijuan Lu, Jianwei Yang, Dan Roth, Dinei A. F. Florêncio, and Cha Zhang. 2025. [Refocus: Visual editing as a chain of thought for structured image understanding](#). *ArXiv*, abs/2501.05452.
- Shibo Hao, Sainbayar Sukhbaatar, DiJia Su, Xian Li, Zhiting Hu, Jason E. Weston, and Yuandong Tian. 2024. [Training large language models to reason in a continuous latent space](#). *ArXiv*, abs/2412.06769.
- Yunzhuo Hao, Jiawei Gu, Huichen Will Wang, Linjie Li, Zhengyuan Yang, Lijuan Wang, and Yu Cheng. 2025. [Can mllms reason in multimodality? emma: An enhanced multimodal reasoning benchmark](#). *ArXiv*, abs/2501.05444.

- Kaiming He, Haoqi Fan, Yuxin Wu, Saining Xie, and Ross B. Girshick. 2019. [Momentum contrast for unsupervised visual representation learning](#). *2020 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 9726–9735.
- Yushi Hu, Weijia Shi, Xingyu Fu, Dan Roth, Mari Ostendorf, Luke S. Zettlemoyer, Noah A. Smith, and Ranjay Krishna. 2024. [Visual sketchpad: Sketching as a visual chain of thought for multimodal language models](#). *ArXiv*, abs/2406.09403.
- Wenxuan Huang, Bohan Jia, Zijie Zhai, Shaoshen Cao, Zheyu Ye, Fei Zhao, Zhe Xu, Yao Hu, and Shaohui Lin. 2025a. [Vision-r1: Incentivizing reasoning capability in multimodal large language models](#). *ArXiv*, abs/2503.06749.
- Zeyi Huang, Yuyang Ji, Anirudh Sundara Rajan, Zefan Cai, Wen Xiao, Junjie Hu, and Yong Jae Lee. 2025b. [Visualtoolagent \(vista\): A reinforcement learning framework for visual tool selection](#). *ArXiv*, abs/2505.20289.
- Ang Li, Charles L. Wang, Kaiyu Yue, Zikui Cai, Ollie Liu, Deqing Fu, Peng Guo, Wang Bill Zhu, Vatsal Sharan, Robin Jia, Willie Neiswanger, Furong Huang, Tom Goldstein, and Micah Goldblum. 2025a. [Zebra-cot: A dataset for interleaved vision language reasoning](#). *ArXiv*, abs/2507.16746.
- Bangzheng Li, Ximeng Sun, Jiang Liu, Ze Wang, Jialian Wu, Xiaodong Yu, Hao Chen, Emad Barsoum, Muhao Chen, and Zicheng Liu. 2025b. [Latent visual reasoning](#). *ArXiv*, abs/2509.24251.
- Bo Li, Yuanhan Zhang, Dong Guo, Renrui Zhang, Feng Li, Hao Zhang, Kaichen Zhang, Yanwei Li, Ziwei Liu, and Chunyuan Li. 2024. [Llava-onevision: Easy visual task transfer](#). *ArXiv*, abs/2408.03326.
- Chenglin Li, Feng Han, Feng Tao, Ruilin Li, Qianglong Chen, Jingqi Tong, Yin Zhang, and Jiaqi Wang. 2025c. [Videopro: Adaptive program reasoning for long video understanding](#).
- Dairu Liu, Ziyue Wang, Minyuan Ruan, Fuwen Luo, Chi Chen, Peng Li, and Yang Liu. 2025. [Visual abstract thinking empowers multimodal reasoning](#). *ArXiv*, abs/2505.20164.
- Hao Shao, Shengju Qian, Han Xiao, Guanglu Song, Zhuofan Zong, Letian Wang, Yu Liu, and Hongsheng Li. 2024a. [Visual cot: Advancing multimodal language models with a comprehensive dataset and benchmark for chain-of-thought reasoning](#). *Advances in Neural Information Processing Systems* 37.
- Hao Shao, Shengju Qian, Han Xiao, Guanglu Song, Zhuofan Zong, Letian Wang, Yu Liu, and Hongsheng Li. 2024b. [Visual cot: Unleashing chain-of-thought reasoning in multi-modal language models](#). *ArXiv*, abs/2403.16999.
- Zhenyi Shen, Hanqi Yan, Linhai Zhang, Zhanghao Hu, Yali Du, and Yulan He. 2025. [Codi: Compressing chain-of-thought into continuous space via self-distillation](#). *ArXiv*, abs/2502.21074.
- Alex Su, Haozhe Wang, Weiming Ren, Fangzhen Lin, and Wenhui Chen. 2025. [Pixel reasoner: Incentivizing pixel-space reasoning with curiosity-driven reinforcement learning](#). *ArXiv*, abs/2505.15966.
- Dawei Wang, Chengmin Zhou, Ding Zhao, Xinyu Liu, Marci Chi Ma, Gary Ushaw, and Rich Davison. 2026a. [Towermind: A tower defence game learning environment and benchmark for llm as agents](#). In *AAAI Conference on Artificial Intelligence*.
- Diany Wang, Wei Song, Yikun Wang, Siyuan Wang, Kaicheng yu, Zhongyu Wei, and Jiaqi Wang. 2025a. [Autoregressive semantic visual reconstruction helps vlms understand better](#). *arXiv*, abs/2506.09040.
- Jiacong Wang, Zijiang Kang, Haochen Wang, Haiyong Jiang, Jiawen Li, Bohong Wu, Ya Wang, Jiao Ran, Xiao Liang, Chao Feng, and Jun Xiao. 2025b. [Vgr: Visual grounded reasoning](#). *ArXiv*, abs/2506.11991.
- Weiyun Wang, Zhangwei Gao, Lixin Gu, Hengjun Pu, Long Cui, Xingguang Wei, Zhaoyang Liu, Linglin Jing, Shenglong Ye, Jie Shao, Zhaokai Wang, Zhe Chen, Hongjie Zhang, Ganlin Yang, Haomin Wang, Qi Wei, Jinhui Yin, Wenhao Li, Erfei Cui, and 44 others. 2025c. [Internvl3.5: Advancing open-source multimodal models in versatility, reasoning, and efficiency](#). *ArXiv*, abs/2508.18265.
- Yubo Wang, Juntian Zhang, Yichen Wu, Yankai Lin, Nils Lukas, and Yuhan Liu. 2026b. [Forest before trees: Latent superposition for efficient visual reasoning](#). *arXiv*, abs/2601.06803.
- Jason Wei, Xuezhi Wang, Dale Schuurmans, Maarten Bosma, Ed H. Chi, F. Xia, Quoc Le, and Denny Zhou. 2022. [Chain of thought prompting elicits reasoning in large language models](#). *ArXiv*, abs/2201.11903.
- Qiucheng Wu, Handong Zhao, Michael Stephen Saxon, Trung M. Bui, William Yang Wang, Yang Zhang, and Shiyu Chang. 2024. [Vsp: Assessing the dual challenges of perception and reasoning in spatial planning tasks for vlms](#). *ArXiv*, abs/2407.01863.
- Weiye Xu, Jiahao Wang, Weiyun Wang, Zhe Chen, Wengang Zhou, Aijun Yang, Lewei Lu, Houqiang Li, Xiaohua Wang, Xizhou Zhu, Wenhui Wang, Jifeng Dai, and Jinguo Zhu. 2025. [Visulogic: A benchmark for evaluating visual reasoning in multi-modal large language models](#). *ArXiv*, abs/2504.15279.
- Yuqi Xu, Rizhen Hu, Zihan Liu, Mou Sun, and Kun Yuan. 2026. [Grouter: Decoupling routing from representation for accelerated moe training](#). *arXiv preprint arXiv:2603.06626*.
- Zeyuan Yang, Xueyang Yu, Delin Chen, Maohao Shen, and Chuang Gan. 2025. [Machine mental imagery: Empower multimodal reasoning with latent visual tokens](#). *ArXiv*, abs/2506.17218.

- Guanghao Zhang, Tao Zhong, Yan Xia, Zhelun Yu, Haoyuan Li, Wanggui He, Fangxun Shu, Mushui Liu, Dong She, Yi Wang, and Hao Jiang. 2025a. [Cmmcot: Enhancing complex multi-image comprehension via multi-modal chain-of-thought and memory augmentation](#). *ArXiv*, abs/2503.05255.
- Huanyu Zhang, Wenshan Wu, Chengzu Li, Ning Shang, Yan Xia, Yangyu Huang, Yifan Zhang, Li Dong, Zhang Zhang, Liang Wang, Tien-Ping Tan, and Furu Wei. 2025b. [Latent sketchpad: Sketching visual thoughts to elicit multimodal reasoning in mllms](#). *ArXiv*, abs/2510.24514.
- Zhuosheng Zhang, Aston Zhang, Mu Li, Hai Zhao, George Karypis, and Alexander J. Smola. 2023. [Multimodal chain-of-thought reasoning in language models](#). *Trans. Mach. Learn. Res.*, 2024.
- Di Zhao, Jingfeng Zhang, Hongsheng Hu, Philippe Fournier-Viger, Gill Dobbie, and Yun Sing Koh. 2025. [Balancing invariant and specific knowledge for domain generalization with online knowledge distillation](#). In *International Joint Conference on Artificial Intelligence*.
- Yakun Zhu, Yutong Huang, Shengqian Qin, Zhongzhen Huang, Shaoting Zhang, and Xiaofan Zhang. 2026. [Medmcp-calc: Benchmarking llms for realistic medical calculator scenarios via mcp integration](#). *arXiv*, abs/2601.23049.
- Yakun Zhu, Zhongzhen Huang, Qianhan Feng, Linjie Mu, Yannian Gu, Shaoting Zhang, Qi Dou, and Xiaofan Zhang. 2025. [Cp-env: Evaluating large language models on clinical pathways in a controllable hospital environment](#). *arXiv*, abs/2512.10206.

A Additional Experimental Results

In this section, we provide supplementary comparisons to further validate the effectiveness of the Interleaved Latent Visual Reasoning (ILVR) framework.

A.1 Comparison with Latent Visual Reasoning (LVR)

As discussed in the main paper, a direct comparison with the original LVR framework (Li et al., 2025b) on the Zebra-CoT (Li et al., 2025a) dataset is not feasible because LVR relies on ground-truth bounding box (BBox) annotations. To ensure a rigorous comparison, we adopted the LVR experimental protocol by training both the LVR baseline and our ILVR model on the Visual-CoT (Shao et al., 2024a) dataset (80k samples).

The results in Table 5 show that ILVR achieves superior generalization, particularly on VisuLogic (Xu et al., 2025) (+4.5% average accuracy) and Zebra-CoT (+0.3% average accuracy). We attribute this improvement to the nature of feature selection. While LVR relies on bounding box annotations to strictly localize regions deemed important by humans, such explicit supervision may not always align with the intrinsic features required by the model for reasoning. In contrast, our momentum teacher autonomously selects visual features based on the current reasoning context. This suggests that adaptively distilled features, which are optimized for the model’s own latent space, provide more effective guidance than rigid human-defined regions, thereby leading to better performance on unseen tasks.

Table 5: **Comparison with LVR fine-tuned on Visual-CoT.** Models are evaluated on OOD benchmarks.

Model	Paradigm	EMMA	VisLog	Zebra	Total
LVR	Direct	24.0%	22.1%	47.0%	31.9%
ILVR (Ours)	Interleaved	21.5%	26.6%	47.3%	32.6%

A.2 Comparison with Sketchpad

We further compare ILVR against Sketchpad (Zhang et al., 2025b). Before analyzing the results, it is important to note a disclaimer regarding the reproduction of the Sketchpad baseline. We encountered data processing discrepancies in the official repository, which prevented direct execution. We have resolved these issues to the best of

our ability to establish a functional baseline; however, these results should be considered tentative and may be updated pending future fixes to the official implementation.

Table 6 details the performance on OOD benchmarks for models fine-tuned on the Zebra-CoT (10k) subset. ILVR achieves a Total Average accuracy of 37.5%, significantly outperforming Sketchpad’s 33.0%. We attribute this performance advantage to the superior efficacy of our teacher-guided feature selection over Sketchpad’s alignment mechanism. Sketchpad operates by projecting hidden states into the vision encoder’s space (prior to LLM projection), forcing them to align with 256 visual tokens derived from a resized 448×448 helper image, and subsequently projecting them back into the LLM space to aid reasoning. This process essentially enforces a rigid alignment with the overall features of the helper image. In contrast, ILVR employs a momentum teacher to actively select features. Instead of aligning to the entire feature map, our teacher dynamically identifies and distills the specific visual cues that are most beneficial for the current reasoning context. This selective mechanism provides more precise and effective guidance than Sketchpad’s global alignment strategy.

Table 6: **Comparison with Sketchpad.** ILVR results correspond to Stage 2.

Model	Paradigm	EMMA	VisLog	Zebra	Total
Sketchpad	Direct	25.0%	25.9%	43.8%	33.0%
ILVR (Ours)	Interleaved	33.3%	29.3%	47.8%	37.5%

B Implementation Details

B.1 Training Infrastructure & Setup

All models were trained on a cluster of $8 \times$ NVIDIA H200 GPUs using DeepSpeed Zero-3 optimization with Qwen2.5-VL-7B backbone. We use the AdamW optimizer with a cosine learning rate scheduler. To prevent overfitting on the limited 10k Zebra-CoT subset, we apply a weight decay of 0.01 and a moderate warmup ratio. The specific hyperparameters for each stage are detailed in Table 7.

B.2 Data Construction Pipeline

We construct the training data to support the interleaved text-latent paradigm. Each data sample is formatted as a conversation containing a user query and a multi-step assistant response.

Table 7: Hyperparameters for ILVR Training.

Hyperparameter	Stage 1	Stage 2
Learning Rate	1×10^{-5}	1×10^{-5}
Batch Size	1	1
Gradient Accumulation	8	1
Latent Tokens (K)	8	8
Align Weight (λ_{sim})	1.0	N/A
EMA Decay (τ)	0.999	N/A
Epochs	15 / 2	15 / 1.5

Below is a simplified example of the data format using a chat template structure. We highlight the interleaved nature of the assistant’s response:

Data Sample Format (Chat Template)

```
[
  {
    "role": "user",
    "content": [
      { "type": "image", "image": "original_input.jpg" },
      { "type": "text", "text": "How many red objects are to the left...?" }
    ]
  },
  {
    "role": "assistant",
    "content": [
      { "type": "text", "text": "First, I need to locate the blue cube..." },
      { "type": "image", "image": "crop_blue_cube.jpg" }, % Becomes <latent>
      { "type": "text", "text": "Now I will scan the area to its left..." },
      { "type": "image", "image": "left_region_red_filter.jpg" }, % Becomes <latent>
      { "type": "text", "text": "I see two red spheres. The answer is 2." }
    ]
  }
]
```

C Additional Clarifications

Novelty. Existing latent visual reasoning approaches typically adopt a single-step paradigm, where a latent representation is generated once and remains static during reasoning. In contrast, ILVR introduces an interleaved latent reasoning paradigm that updates latent states at each step, enabling dynamic state evolution. In addition, ILVR employs selective perceptual modeling to distill task-relevant visual information, mitigating the over-compression issue in prior methods.

Scalability of Interleaved Training. ILVR assumes access to intermediate visual states (“helper images”), which are commonly available or can be constructed in many multimodal reasoning settings. For example, sequential or embodied tasks naturally provide intermediate observations, while in other domains, intermediate visualizations can be obtained through automated procedures such as step-wise highlighting, cropping, or model-based synthesis. Importantly, ILVR does not introduce additional supervision requirements, but instead improves the efficiency of existing interleaved reasoning pipelines by replacing pixel-level processing with latent representations.

D Detailed Dataset Composition

To evaluate the robustness and versatility of our framework, we curate a diverse suite of benchmarks encompassing both in-distribution (IID) and out-of-distribution (OOD) settings.

For in-distribution evaluation, we focus on fine-grained visual perception and sequential planning using the COMT and VSP datasets. To further assess generalizability, we construct a strictly controlled subset from the Zebra-CoT dataset as our OOD benchmark, challenging the model with multi-step reasoning across scientific and logic domains. Table 8 provides a comprehensive summary of the statistics and task definitions.

Table 8: Summary of Dataset Composition. Key characteristics for the three primary benchmarks.

Dataset	Split	Key Characteristics
COMT	3.4k / 400 (IID)	Atomic Manipulation: Fine-grained perception tasks (Creation, Deletion, Selection, Update).
VSP	1k / 400 (IID)	Sequential Planning: Tracking visual state changes over long horizons.
Zebra-CoT	10k / - (OOD)	Complex Reasoning: <ol style="list-style-type: none"> <i>Science:</i> Physics, Graphs. <i>Logic:</i> Chess, Ciphers, Maze, Tetris, RPM <i>3D:</i> Counting, Planning, Embodied.

D.1 Case Studies

To illustrate the generalization capabilities of ILVR, we present a series of inference cases using the Qwen2.5-VL-7B backbone. For visualization purposes, we represent the latent reasoning segments with <latent_start><latent_pad><latent_end>.

The tasks of the selected examples include fine-grained perception on COMT (Selection and Deletion, Fig. 6 and Fig. 7) , 2D visual search on held-out Zebra-CoT (Fig. 8 and Fig. 9) , visual logic reasoning on VisuLogic (Quantitative Reasoning, Fig. 10) , and Math reasoning on EMMA (Fig. 11). Note that ILVR effectively utilizes latent thinking to model evolving states. For example, in the Deletion and Position tasks, the generated latent representations dynamically update to reflect the removal of objects or the simulation of movement trajectories, rather than relying on static visual features.

EXAMPLE1: COMT DELETION TASK

Question: The presented image features two pictures: one aligned to the left and the other to the right, with a vertical transparent line separating them. What is the count of differences between the images? option:[6, 9, 12, 15]

Question Image



Text - Latent Interleaved Inference:

<latent_start><latent_pad><latent_end> First off, we crop the sections where the two images are not the same. <latent_start><latent_pad><latent_end> As illustrated in the above image, we have identified 4 differences. Drawing from this, we continue our exploration for the differences between the two pictures, cropping them. <latent_start><latent_pad><latent_end> According to the image above, we have discovered 4 variations. Building on this, we carry on searching for the dissimilarities between the two images and crop them. <latent_start><latent_pad><latent_end> In accordance with the above picture, we have pinpointed 4 dissimilarities. Taking this into consideration, we continue our search for the differences between the two images, cropping them. <latent_start><latent_pad><latent_end> As indicated in the image above, we have located 3 differences. In conclusion, there are 15 differences present in both images.

The final answer is: 15

Figure 6: Example of COMT Deletion Task

EXAMPLE2: COMT SELECTION TASK

Question: How many people can you identify in the image? option:[20, 30, 40, 50]

Question Image



Text - Latent Interleaved Inference:

<|latent_start|><|latent_pad|><|latent_end|> Our initial action involves removing the faces on the left side of the image. <|latent_start|><|latent_pad|><|latent_end|> The above illustration illustrates our discovery of a total of 10 people. Taking this into consideration, we extend our examination of faces to the right in the image, deleting them then. <|latent_start|><|latent_pad|><|latent_end|> In the depicted image, our findings indicate the presence of 10 individuals. Expanding on this, we proceed to the right in the image, seeking faces and removing them. <|latent_start|><|latent_pad|><|latent_end|> As shown in the above image, we have identified a total of 10 people. Based on this premise, we advance our examination of faces to the right in the image, eliminating the faces identified. <|latent_start|><|latent_pad|><|latent_end|> Based on the visual representation above, we've pinpointed a count of 10 people. Considering this development, we extend our search for faces to the right side of the image, and delete the faces found. <|latent_start|><|latent_pad|><|latent_end|> According to the illustration above, we've discovered a group of 0 individuals. In overview, the image showcases a grand total of 40 individuals.

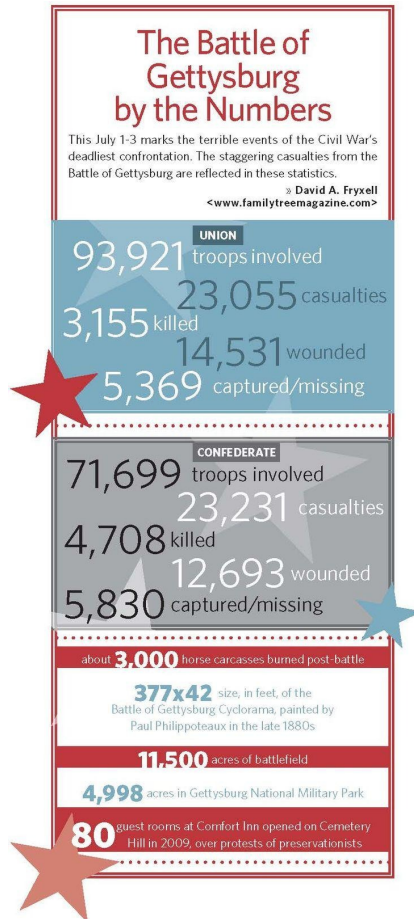
The final answer is: 40

Figure 7: Example of COMT Selection Task

EXAMPLE3: ZEBRA-COT OOD 2D VISUAL SEARCH TASK1

Question: Based on the provided infographic titled "The Battle of Gettysburg by the Numbers," how many lives were lost on the Confederate side?

Question Image



Text - Latent Interleaved Inference:

THOUGHT 0: The problem asks for the number of lives lost on the Confederate side during the Battle of Gettysburg. I need to analyze the given data to find this specific number.
THOUGHT 1: To accurately determine the number of lives lost, I will focus on the section of the infographic that provides the detailed casualty counts for both the Union and Confederate forces. This visual representation will help in clearly identifying the relevant figures.
THOUGHT 2: Upon reviewing the highlighted image, which is a close-up of the original, it's clear that the number of "killed" Confederates is explicitly stated as 4,708. This represents the total number of lives lost on the Confederate side.

Based on the reasoning, the final answer is: 4,708

Figure 8: Example of Zebra-CoT 2D Visual Search (Task 1)

EXAMPLE4: ZEBRA-COT OOD 2D VISUAL SEARCH TASK2

Question: Identify the company logo visible on the jet bridge in the provided airport image.

Question Image



Text - Latent Interleaved Inference:

THOUGHT 0: The problem asks to identify the company logo on the jet bridge shown in the image. I need to analyze the visual elements of the scene, focusing on the large white structure that appears to be a jet bridge.
THOUGHT 1: To clearly see the details and identify the logo, it's beneficial to have a closer view of the structure. I will generate an image that provides a more detailed perspective of the jet bridge.
THOUGHT 2: From this closer view, the red and blue colors of the HSBC (Hankins, Standard, Baring) logo are distinctly visible on the side of the bridge. The letters "HSBC" are prominently displayed in red, followed by the blue cross symbol, and the tagline "The world's local bank" is written below in smaller text.


Based on the reasoning, the final answer is: HSBC

Figure 9: Example of Zebra-CoT 2D Visual Search (Task 2)

EXAMPLE5: VISULOGIC QUANTITATIVE REASONING

Question: Divide the six figures below into two categories, so that each category has its own common characteristics or patterns. The correct classification is:\n\nA: ①③④, ②⑤⑥ \nB: ①③⑤, ②④⑥ \nC: ①②⑥, ③④⑤ \nD: ①④⑥, ②③⑤

Question Image



Text - Latent Interleaved Inference:

THOUGHT 0: The problem asks to identify the correct grouping of the six provided figures. I need to find which set of three figures in each option share a common characteristic or pattern.\nTHOUGHT 1: Let's first examine the shapes and lines within the figures to identify potential commonalities. I will isolate one figure to analyze its structure. <|latent_start|><|latent_pad|><|latent_end|>THOUGHT 2: The isolated figure (1) shows a single curved line (an arc) that connects to the straight lines at its two ends, forming part of the boundary. Notably, the curve does not cross the straight lines.\nTHOUGHT 3: Now, let's test option A by checking its figures (1, 3, 4). Figure 1 fits the description. I will examine figure 3 to see if it shares this "non-crossing boundary arc" pattern.<|latent_start|><|latent_pad|><|latent_end|>THOUGHT 4: As shown in the image, figure 3 also has an arc connected to the straight lines at its ends without crossing them. This matches figure 1. Now I must check the last figure in this group, figure 4.\nTHOUGHT 5: Figure 4 contains a complete circle positioned inside the straight lines. This is a closed curve, unlike the open arcs in figures 1 and 3. This does not match the identified pattern.\nTHOUGHT 6: Since figure 4 does not match, option A is incorrect. I will now evaluate option B (1, 3, 5). Figures 1 and 3 already matched. I need to check figure 5. <|latent_start|><|latent_pad|><|latent_end|>THOUGHT 7: Figure 5 contains an arc connected to the straight lines at its ends, and it is not crossed by any line. This matches the pattern found in figures 1 and 3.\nTHOUGHT 8: Since {1, 3, 5} form a consistent group, I should briefly check the other group in option B, {2, 4, 6}. In figure 2, the line crosses the curve. In figure 6, the line crosses the curve. In figure 4, the curve is closed/detached. They differ from the first group. <|latent_start|><|latent_pad|><|latent_end|> THOUGHT 9: This supports Option B. To be sure, I will check option C (1, 2, 6). Figure 1 matches. Let's examine figure 2.\nTHOUGHT 10: Figure 2 contains an oval that is crossed/intersected by a straight line. This fundamental structural difference means it does not match the "non-crossing" pattern of figure 1.\nTHOUGHT 11: Option C is incorrect because figure 2 does not match. Finally, I will check option D (1, 4, 6). Figure 1 matches. I will examine figure 4. <|latent_start|><|latent_pad|><|latent_end|> THOUGHT 12: As noted before, figure 4 contains a circle (a closed curve), which does not match the open arc pattern of figure 1.\nTHOUGHT 13: Option D is incorrect because figure 4 does not match. After evaluating all options, only option B's first group {1, 3, 5} consistently shares the specific visual characteristic.\nTHOUGHT 14: The characteristic is that the curve is an arc acting as a boundary connected at its ends, without being crossed by straight lines. The other figures {2, 4, 6} either have crossing lines or closed curves. <|latent_start|><|latent_pad|><|latent_end|> THOUGHT 15: Therefore, the correct classification groups the figures into {1, 3, 5} and {2, 4, 6}.\nTHOUGHT 16: I have verified that {1, 3, 5} are all non-crossing boundary arcs, and {2, 4, 6} are not (they are crossed or closed).\nTHOUGHT 17: This confirms that Option B is the only valid classification.\nTHOUGHT 18: No other option provides a consistent grouping based on the structural relationship between the curve and the straight lines.\nTHOUGHT 19: The logic holds up under scrutiny. Based on the reasoning, the final answer is: B


Figure 10: Example of VisuLogic Quantitative Reasoning

EXAMPLE6: EMMA BENCH MATH REASONING

Question: A road leads away from each of the six houses (see diagram). A hexagon showing the roads in the middle is however, missing. Which hexagons fit in the middle so that one can travel from A to B and to E, but not to D?

"options": ["1 and 2", "1 and 4", "1 and 5", "2 and 3", "4 and 5"]

Question Image



Text - Latent Interleaved Inference:

THOUGHT 0: The problem asks us to find which set of hexagons allows us to travel from house A to B and E, but not D. We need to analyze the given options. THOUGHT 1: Let's first consider option A, which consists of two specific hexagons. We will visualize these two hexagons together to see if they form a valid path. THOUGHT 2: The visual representation shows that hexagon 2 in option A fails to connect to house A at all. This path is invalid. THOUGHT 3: Next, let's examine option B, which also contains two hexagons. We will visualize this combination to check its connectivity. THOUGHT 4: The visualization for option B clearly shows that hexagon 4 connects A, B, and E, but it also includes D. This means we can reach D, which is not allowed by the problem statement. Therefore, option B is incorrect. THOUGHT 5: Now, let's look at option C, which has two hexagons. We will generate an image to test its validity. THOUGHT 6: The image for option C demonstrates that both hexagons 1 and 5 allow us to travel from A to B and E, but crucially, neither includes the path to D. This matches the problem's criteria. THOUGHT 7: Let's proceed to evaluate option D, which contains two hexagons. We will visualize the path formed by these two shapes. THOUGHT 8: The visual thought for option D shows that while it connects other houses, it fails to connect to house A. Thus, option D is not a valid solution. THOUGHT 9: Finally, let's analyze option E, which also consists of two hexagons. We will visualize their combined effect. THOUGHT 10: The image for option E shows that hexagon 4 connects to D, which violates the condition. Thus, option E is incorrect. THOUGHT 11: Comparing the valid options, we can see that only option C allows reaching B and E without reaching D. Options A, B, D, and E all fail due to bad connections or forbidden paths.

Based on the reasoning, the final answer is: C

Figure 11: Example of EMMA Bench Math Reasoning