

VITAL: Visual-Semantic Dual Supervision for Enhanced and Interpretable Latent Reasoning in Medical MLLMs

Qiaoru Li^{1,3*}, Shaotian Liang¹, Jintao Chen¹, Haoran Sun², Yuxiang Cai^{1,4,5†},
Jianwei Yin¹, Yankai Jiang^{1,2†}

¹Zhejiang University, ²Shanghai Artificial Intelligence Laboratory, ³Tencent,

⁴Ningbo Global Innovation Center, Zhejiang University,

⁵Zhejiang Key Laboratory of Digital-Intelligence Service Technology

Correspondence: caiyuxiang@zju.edu.cn, jiangyankai@pjlab.org.cn

Abstract

Latent reasoning enables reasoning over continuous hidden states rather than explicit tokens, avoiding the language bottleneck and inference overhead of chain-of-thought for medical VQA. However, existing methods suffer from modality collapse, insufficient visual supervision, and train–inference mismatch. Moreover, their opaque latent states offer no interpretability, which is critical in clinical applications. We propose VITAL, a latent-space reasoning framework for medical MLLMs with visual-semantic dual supervision: an auxiliary text decoder reconstructs reasoning chains from latent states, while a visual projector regresses ROI features from a frozen, independent medical vision encoder. Both modules are discarded at inference with zero overhead, yet can be re-attached post-hoc for dual interpretability, providing textual and visual explanations of the reasoning process without sacrificing efficiency. We construct a 61K dataset spanning 9 imaging modalities, exceeding prior medical visual latent reasoning datasets by an order of magnitude. Experiments on 7 benchmarks show that VITAL consistently and substantially outperforms the backbone, all latent reasoning baselines, and medical MLLMs trained on far larger data, achieving state-of-the-art results competitive with trillion-parameter proprietary models. Code is available at <https://github.com/EmberQR/vital2026>.

1 Introduction

Medical multimodal large language models (MLLMs) have rapidly evolved from image-level question answering toward fine-grained visual understanding (Li et al., 2023; Chen et al., 2024; Xu et al., 2025; Shu et al., 2025), yet inadequate visual grounding remains a primary bottleneck (Liu et al., 2026). This has motivated Chain-of-Thought (CoT)

reasoning for medical VQA, decomposing diagnosis into structured multi-step processes (Qiao et al., 2025; Wang et al., 2025b; Fan et al., 2026) with expert-annotated visual evidence at each step (LeDuc et al., 2025). However, explicit CoT in the multimodal setting suffers from two inherent limitations: (i) a language bottleneck: fine-grained visual evidence (lesion texture, boundary morphology, subtle intensity gradients) must be compressed into discrete text tokens, inevitably losing perceptual details critical for clinical judgment; (ii) inference overhead: generating verbose reasoning chains substantially increases latency and cost in time-sensitive clinical environments.

These limitations motivate latent reasoning, where the model reasons over continuous hidden states rather than explicit text tokens (Figure 1). Coconut (Hao et al., 2024) first demonstrated multi-step latent reasoning by iteratively feeding hidden states back as inputs; subsequent work refined this via self-distillation (Shen et al., 2025b), compact hidden thinking states (Shen et al., 2025a), and step-level supervision to prevent state collapse (Wei et al., 2025). The paradigm has since been extended to vision-language models (Li et al., 2025; Pham and Ngo, 2025). In the medical domain, MedLVR (Xi et al., 2026) applies latent visual reasoning to medical VQA by aligning latent states with ROI patch tokens from the VLM’s own visual encoder and applying RL-based policy optimization. However, MedLVR exhibits three key limitations: (i) its visual supervision is self-referential, aligning to the model’s own patch embeddings rather than an independent expert medical vision encoder; (ii) its latent states remain black box, offering no post-hoc interpretability; and (iii) its two-stage pipeline introduces different computation graphs at training and inference, creating train–test mismatch.

We propose **VITAL** (Visual-Implicit Thinking with Aligned Latents), a latent-space reasoning framework for medical MLLMs that addresses all

*Work done during an internship.

†Corresponding authors.




Figure 1: Comparison of reasoning paradigms.

three limitations. VITAL features a recurrent latent reasoning loop with strict train–inference parity, and introduces visual-semantic dual supervision: an auxiliary text decoder reconstructs the reasoning chain from each latent state, while a visual projector regresses ROI features from a frozen, independent medical vision encoder. Both auxiliary modules serve as training-time scaffolding, constraining latent states to encode both textual logic and visual evidence. They are **entirely discarded at inference** with **zero additional overhead**. Crucially, these modules can be re-attached post-hoc, yielding dual interpretability: a textual window (what the model reasons) and a visual window (where the model looks), which is a remarkable level of transparency for latent reasoning.

Our contributions are summarized as follows:

- We propose **VITAL**, a novel latent reasoning framework with visual-semantic dual supervision for medical MLLMs. VITAL maintains strict train–inference parity through a single recurrent loop, prevents modality collapse via orthogonal dual supervision anchored to an independent medical vision encoder, and offers dual interpretability, all with zero inference overhead via its plug-and-play scaffolding design.
- We construct a large-scale dataset of **~61K** five-tuple training samples (image, question,

answer, K -step reasoning chain, ROI visual feature) from two public medical segmentation datasets, spanning **9 imaging modalities**. To our knowledge, this is the largest dataset specifically designed for training latent reasoning in the medical vision domain, exceeding prior efforts by $6\times$ – $57\times$ in scale.

- Extensive experiments on **7 benchmarks** (2 in-domain testsets, 3 medical VQA benchmarks, 1 visual grounding benchmark, and 1 held-out in-house testset) demonstrate that VITAL consistently improves over the backbone on every evaluation, surpasses all latent reasoning baselines and medical MLLMs trained on far larger data, and achieves performance competitive with trillion-parameter proprietary models while exceeding them on visual grounding by a large margin.

2 Related Work

2.1 Multimodal Latent-Space Reasoning

Coconut (Hao et al., 2024) pioneered multi-step latent reasoning in LLMs; Heima (Shen et al., 2025a) compressed reasoning into compact hidden states with reconstructable traces; and SIM-CoT (Wei et al., 2025) showed that step-level supervision is essential to prevent latent state collapse. In the multimodal setting, Latent Visual Reasoning (Li et al., 2025) interleaves reasoning in visual embed-

ding space with text generation, and MedLVR (Xi et al., 2026) first applies this paradigm to medical VQA. However, MedLVR relies on self-referential visual supervision from the model’s own encoder, offers no post-hoc interpretability for latent states, and introduces train–inference mismatch through its two-stage pipeline.

2.2 Medical Visual Grounding and Multimodal Alignment

Medical MLLMs have evolved from image-level QA (Li et al., 2023; Chen et al., 2024) toward fine-grained multimodal reasoning with broader domain coverage (Xu et al., 2025; Deria et al., 2026). Meanwhile, grounding capabilities have gained prominence: LISA (Lai et al., 2024) and GLaMM (Rasheed et al., 2024) introduced reasoning segmentation and pixel-level grounding in the general domain, with medical extensions (MedPLIB (Huang et al., 2025), UniBiomed (Wu et al., 2025), MMedAgent (Li et al., 2024), Citrus-V (Wang et al., 2025a), and IBISAgent (Jiang et al., 2026b)) demonstrating the value of coupling MLLMs with specialized tools and pixel prompts. However, these methods rely on explicit mask annotations or external tool invocation; direct reasoning in latent visual space for medical understanding remains underexplored, which is the gap our work targets.

3 Method

3.1 Overview of VITAL

VITAL is a latent-space reasoning framework built on a frozen MLLM backbone. As shown in Figure 2, the forward pass has three stages: (1) the multimodal backbone encodes the image and question into a prefix KV-cache; (2) a recurrent latent loop produces K continuous states $\{z_1, \dots, z_K\}$ (§3.2); (3) the answer is auto-regressively decoded from the enriched cache.

Without explicit supervision, latent states may collapse into a text-dominant subspace, discarding fine-grained visual evidence. VITAL addresses this with visual-semantic dual supervision (§3.3): an auxiliary text decoder and a visual projector constrain latent states from orthogonal directions, yet are **entirely discarded at inference** (§3.5).

3.2 Recurrent Latent Reasoning

Prefix encoding. Given a medical image $\mathbf{I}$ and a question $\mathbf{q}$, the frozen vision encoder extracts visual tokens which are fused with text tokens

through the multimodal backbone. A full forward pass over the concatenated prefix yields the hidden states and a prefix KV-cache $\mathcal{C}_{\text{prefix}}$.

Recurrent latent loop. We adopt a recurrent mechanism to produce K latent reasoning states. Let $z_0 = h_{\text{prefix}}^{\text{last}} \in \mathbb{R}^d$ denote the last hidden state of the prefix. At each step $k = 1, \dots, K$, the language model performs a single-token forward pass:

$$z_k = f_{\theta}(z_{k-1}; \mathcal{C}_{\text{prefix} \cup \{z_{<k}\}}), \quad (1)$$

where f_{θ} denotes one transformer forward step and the KV-cache is updated in place. Each step’s input is the output of the previous step, creating an auto-regressive causal chain over continuous vectors rather than discrete tokens.

Train–inference parity. A critical design is that both training and inference traverse the identical recurrent path: there is no teacher forcing, scheduled sampling, or parallel computation of latent states. This eliminates train–test mismatch and distinguishes VITAL from prior latent reasoning methods that use different computation graphs at training and inference time.

Answer decoding. After the K -step loop, the answer token embeddings are fed into the same KV-cache and decoded by the language model head. Let $\mathbf{a} = (a_1, \dots, a_T)$ denote the answer token sequence of length T . Notably, the first answer token is predicted directly from the final latent state z_K , ensuring a smooth transition from continuous reasoning to discrete generation:

$$p(\mathbf{a}) = \prod_{t=1}^T p(a_t \mid z_{1:K}, a_{<t}). \quad (2)$$

3.3 Visual-Semantic Dual Supervision

Without explicit constraints, latent states in multimodal models tend to collapse into a text-dominant subspace, discarding subtle visual cues such as lesion texture and boundary morphology. VITAL prevents this modality collapse by supervising every latent state z_k from two directions: semantic and visual (see Figure 2).

Semantic Supervision. To ensure that each latent state encodes interpretable reasoning logic, we attach an auxiliary text decoder, an independent, smaller causal language model with its own embedding layer and prediction head (no weight sharing




Figure 2: **Overview of VITAL.** The multimodal backbone encodes the input into a prefix KV-cache. A recurrent latent loop iterates K steps ($z_k = f_\theta(z_{k-1}; \mathcal{C})$) with identical paths at training and inference. Latent states are supervised by $\mathcal{L}_{\text{task}}$ (answer), $\mathcal{L}_{\text{text}}$ (auxiliary text decoder), and $\mathcal{L}_{\text{visual}}$ (visual projector regressing ROI features). Both auxiliary modules are discarded at inference with zero overhead.

with the backbone). Let e_k denote the ground-truth reasoning text for step k in the K -step reasoning chain. For each latent step k , the hidden state is first projected by a learnable linear layer $\text{pj_in} : \mathbb{R}^d \rightarrow \mathbb{R}^{d_{\text{dec}}}$ to match the decoder’s hidden dimension, then concatenated with the decoder’s own embeddings of e_k :

$$\mathbf{h}_k^{\text{dec}} = [\text{pj_in}(z_k); \text{Embed}_{\text{dec}}(e_k)]. \quad (3)$$

The decoder auto-regressively predicts e_k , and we compute a cross-entropy loss averaged over all K valid steps:

$$\mathcal{L}_{\text{text}} = \frac{1}{K} \sum_{k=1}^K \text{CE}(\text{Decoder}(\mathbf{h}_k^{\text{dec}}), e_k). \quad (4)$$

Visual Supervision. To anchor each latent state to fine-grained visual evidence, we project it into the feature space of a frozen medical vision encoder and regress a pre-extracted ROI feature. A naïve linear projection, however, is prone to training instability due to the severe norm imbalance between the LLM’s hidden states and the vision encoder’s feature space: [Pham and Ngo \(2025\)](#) report that the last hidden state norm can be $4\times$ larger than the vision feature norm, making direct mapping numerically fragile.

We address this with a visual projector (VP) based on a norm-stabilized residual MLP. Let d_v denote the feature dimension of the frozen medical vision encoder. The projector is defined as:

$$\text{VP}(z) = \text{LN}(W_2(\text{GELU}(W_1 \text{LN}(z)) + z)), \quad (5)$$

where $W_1 \in \mathbb{R}^{d \times d}$ and $W_2 \in \mathbb{R}^{d_v \times d}$. The design is motivated by three considerations: (i) the **input LayerNorm** normalizes the LLM hidden state before any linear transformation, eliminating the norm imbalance at source; (ii) the **residual connection** in d -dimensional space (around the expansion layer $W_1 + \text{GELU}$) preserves the representational richness of z_k while allowing the non-linear pathway to learn a modality-bridging correction; (iii) the **output LayerNorm** re-scales the projected d_v -dimensional output to match the ℓ_2 -normalized scale of the vision encoder features, ensuring stable L1 regression.

The supervision target $\mathbf{f}_{\text{ROI}} \in \mathbb{R}^{d_v}$ is a pre-extracted, ℓ_2 -normalized ROI feature for the target anatomical region (§3.4). All K latent steps share the same target, and an L1 loss is applied:

$$\mathcal{L}_{\text{visual}} = \frac{1}{K} \sum_{k=1}^K \|\text{VP}(z_k) - \mathbf{f}_{\text{ROI}}\|_1. \quad (6)$$

We use a shared target for all steps: $\mathcal{L}_{\text{text}}$ already differentiates steps, so the shared visual anchor

provides a stable spatial signal without imposing an artificial curriculum (ablation in §C.6).

Dual Interpretability. As an important by-product of dual supervision, both auxiliary modules can be re-attached during post-hoc analysis. They are discarded at inference time. The text decoder can decode any latent state into a human-readable reasoning chain, providing a textual window into the model’s implicit thought process. Simultaneously, the visual projector output can be used for nearest-neighbor retrieval in the medical vision encoder’s feature space, recovering the visual content that the latent state encodes, which serves as a visual window into what the model “sees in its mind.” To the best of our knowledge, this work is among the first to achieve both textual and visual interpretability for latent reasoning. We provide a qualitative example in Figure 3.

3.4 Training Data Construction

VITAL requires each training sample to be a five-tuple: (image, question, answer, K -step reasoning chain, ROI visual feature). We build an automated, reproducible pipeline from two public medical segmentation datasets: MSD (Antonelli et al., 2022) and BiomedParse (Zhao et al., 2025), producing 61K samples in total.

Teacher distillation. Textual reasoning chains and answers are generated by a state-of-the-art proprietary MLLM serving as the teacher. The teacher receives a teacher-only overlay image in which the segmentation mask is rendered as a semi-transparent color layer on top of the original image, together with the target identity as hidden metadata. Crucially, the teacher is instructed to produce outputs as if it can only see the raw, unannotated image, leveraging the overlay solely for internal correctness while never leaking annotation cues. Each question is classified into one of five types (yes/no, identification, location, description, reasoning), and the reasoning depth $K \in \{1, \dots, 4\}$ is set accordingly. A multi-round quality gate verifies every generated sample against a forbidden-word list to catch annotation leakage, validates reasoning step counts, and normalizes answer format. A subsequent 5% human audit confirms a leakage rate below 0.3%. Full pipeline details, prompt templates, and quality statistics are provided in §A.

Adaptive ROI feature extraction. For each sample, we extract a d_v -dimensional ROI feature from

a frozen medical vision encoder as the visual supervision target. To handle the wide range of target sizes (from large organs spanning the full image to small nodules occupying only a few patches), we employ an adaptive extraction strategy: when the mask covers $\geq 20\%$ of the image area, we extract patch features from the full image and pool over mask-aligned patches; when the mask is smaller, we tightly crop the ROI region (with a 5% margin), pad it to a square, and pool over valid patches. All features are ℓ_2 -normalized to the unit hypersphere. See §A.5 for algorithmic details and threshold ablations.

3.5 Training Objective and Plug-and-Play Design

Joint objective. The total training loss combines the three supervision signals:

$$\mathcal{L} = \mathcal{L}_{\text{task}} + \lambda_1 \mathcal{L}_{\text{text}} + \lambda_2 \mathcal{L}_{\text{visual}}, \quad (7)$$

where $\mathcal{L}_{\text{task}}$ is the standard next-token cross-entropy loss computed only over answer tokens, and λ_1, λ_2 are scalar weights that balance the auxiliary losses.

Curriculum learning. Since longer latent chains are harder to optimize from scratch, we adopt a three-phase curriculum: **Phase 1** ($K=1$) teaches single-step latent reasoning; **Phase 2** ($K \leq 2$) introduces two-step chains; **Phase 3** ($K \leq 4$) trains on all reasoning depths. Each phase warm-starts from the previous one.

Plug-and-play inference. The auxiliary text decoder, visual projector, and `pj_in` serve exclusively as training-time scaffolding. At inference, all these scaffoldings are discarded: the deployed model consists solely of the frozen backbone augmented with LoRA, the latent loop, and the language model head. The resulting inference cost is identical to that of a standard LoRA-tuned model with no auxiliary modules: **zero extra latency, zero extra memory, zero extra parameters.**

4 Experiment

4.1 Experimental Setup

Implementation. We build VITAL on top of Qwen3-VL-8B (Bai et al., 2025). The vision encoder and LLM backbone are both frozen; only LoRA adapters ($r=64, \alpha=128$) are inserted into all attention and feed-forward projections of the

36 transformer layers. The training-time scaffolding consists of an auxiliary text decoder (Qwen3-1.7B, fully fine-tuned), a visual projector, and a projection layer π_{in} . In total, 1.93B out of 10.07B parameters are trainable (19.1%); at inference all scaffolding is discarded, leaving only the LoRA weights (174.6M). We train on our 61K five-tuple dataset (§3.4) with a three-phase curriculum: Phase 1 ($K=1$, 5 epochs), Phase 2 ($K \leq 2$, 5 epochs), and Phase 3 (all K , 10 epochs), using AdamW with learning rate 2×10^{-5} , 5% warmup, and bf16 precision. All experiments are conducted on 6×NVIDIA RTX PRO 6000 (Blackwell). Complete hyperparameters and per-module parameter breakdowns are provided in §B.

Benchmarks. We evaluate on **7 benchmarks** spanning four categories: (i) In-domain testsets: held-out splits of MSD (Antonelli et al., 2022) and BiomedParse (Zhao et al., 2025); (ii) Medical VQA: VQA-RAD (Lau et al., 2018), PathVQA (He et al., 2020), and PMC-VQA (Zhang et al., 2024); (iii) Visual grounding: GEMeX-RMCoT (Liu et al., 2025); and (iv) Held-out in-house testset: a private evaluation set of 1,000 brain neuroimaging VQA pairs (MRI and PET), curated to avoid training-set contamination from public benchmarks. For closed-ended questions (yes/no and choice) we report accuracy; for open-ended questions we report token-level F1 (see §B.4).

Baselines. We compare against four groups of methods (Table 1). **(1) Text latent reasoning:** Coconut (Hao et al., 2024), CODI (Shen et al., 2025b), and SIM-CoT (Wei et al., 2025). These methods operate on text-only LLMs and thus have no original multimodal weights; we **re-implement** them on our backbone (Qwen3-VL-8B + LoRA) with identical training data, differing only in the reasoning paradigm. **(2) Visual latent reasoning:** MCOU (Pham and Ngo, 2025) and LVR (Li et al., 2025). We report both their original weights and re-implementations on our backbone. **(3) Medical MLLMs:** LLaVA-Med (Li et al., 2023), HuatuoGPT-Vision (Chen et al., 2024), Fleming-VL (Shu et al., 2025), and VividMed (Luo et al., 2025), evaluated with their official weights. **(4) General MLLMs:** Qwen3-VL-8B (zero-shot), Qwen3-VL-8B-Thinking (Bai et al., 2025) (zero-shot explicit CoT from the same backbone family, without fine-tuning on our 61K dataset), GPT-5.4 (OpenAI, 2026), and Claude-Opus-4.6 (Anthropic, 2026).

4.2 Main Results

VQA Performance. Table 1 presents comprehensive results across 7 benchmarks. A striking finding is that all re-implemented general latent reasoning methods (Coconut, CODI, SIM-CoT, MCOU, LVR) degrade the Qwen3-VL-8B backbone after fine-tuning on the same data (e.g., Coconut scores only 58.76 on VQA-RAD closed-ended, below the zero-shot backbone’s 67.63), indicating that naïve latent reasoning can harm multimodal performance. In contrast, VITAL consistently improves over the zero-shot backbone on every single benchmark, demonstrating that our visual-semantic dual supervision effectively prevents modality collapse during latent training. Notably, VITAL trained on only 61K samples outperforms medical MLLMs that are trained on datasets orders of magnitude larger, highlighting the data efficiency of well-supervised latent reasoning. Compared with Qwen3-VL-8B-Thinking (a zero-shot explicit-CoT model from the same backbone family), VITAL achieves substantially higher accuracy across all benchmarks (e.g., +40.56 on MSD, +45.49 on BiomedParse, +45.00 on GEMeX-RMCoT) while being $\sim 97\times$ faster at inference (see §4.3), confirming that latent reasoning in a well-supervised continuous space is far more effective than verbose token-level reasoning chains for medical VQA. On closed-ended metrics, VITAL surpasses all latent reasoning baselines and all medical MLLMs, improving over Qwen3-VL-8B by +47.52 on MSD, +46.67 on BiomedParse, +13.15 on PathVQA, and +7.05 on PMC-VQA; the only methods that outperform VITAL on individual splits are trillion-parameter proprietary models. On visual grounding (GEMeX-RMCoT), VITAL achieves 86.67, exceeding the best proprietary model by 28.34, strongly validating that latent states encode fine-grained spatial evidence. On the held-out in-house testset, VITAL reaches 80.50, outperforming GPT-5.4 (73.80) and all medical MLLMs, confirming robust generalization beyond the training distribution. On open-ended Token-F1, VITAL trails medical MLLMs optimized for verbose free-text generation (e.g., Fleming-VL), reflecting an inherent limitation of latent reasoning: it produces concise answers that score lower on token-overlap metrics designed to reward verbose outputs.

Comparison with Matched Verbalized-CoT Fine-Tuning. The Qwen3-VL-8B-Thinking baseline above evaluates zero-shot explicit CoT

Table 1: **Main results on medical VQA benchmarks.** Text latent reasoning baselines (Coconut, CODI, SIM-CoT) operate on text-only LLMs and thus have no original multimodal weights; we report their **re-implementations on our backbone (Qwen3-VL-8B + LoRA) with identical training data**. For visual latent reasoning baselines, we report both original weights and **fine-tuned** results. The top three results in each row are in **bold** and highlighted with an orange background, where the color intensity indicates the rank. The **+x.xx** on VITAL denotes the absolute improvement over Qwen3-VL-8B (zero-shot). All values are percentages (%).

Benchmarks	Methods	Metrics	General Text Latent			General Visual Latent		Medical MLLMs				General MLLMs				VITAL
			Coconut	CODI	SIM-CoT	MCOUT	LVR	LLaVA-Med	HuatuoGPT-Vision	Fleming-VL	VividMed	Qwen3-VL-8B (Zero-shot)	Qwen3-VL-8B -Thinking (Zero-shot)	GPT-5.4	Claude-Opus-4.6	
In-Domain Testsets																
MSD		Acc ↑	40.58	41.09	47.25	27.63 (38.54)	30.22 (48.63)	46.74	55.57	65.50	40.19	35.56	42.52	57.65	54.26	83.08 ^{+47.52}
BiomedParse		Acc ↑	38.15	36.77	43.37	22.88 (38.32)	28.85 (44.97)	45.31	50.27	61.49	37.74	32.01	33.19	53.52	55.12	78.68 ^{+46.67}
Medical VQA																
VQA-RAD	Acc _{closed} ↑	58.76	60.31	64.08	35.25 (59.87)	39.47 (63.19)	58.96	70.29	69.62	45.45	67.63	64.52	81.67	77.29	74.28 ^{+6.65}	
	Token-F1 _{open} ↑	33.23	28.14	25.93	5.62 (30.19)	3.29 (27.76)	27.45	34.35	37.94	7.29	31.50	26.77	16.73	3.07	32.59 ^{+1.09}	
PathVQA	Acc _{closed} ↑	56.14	56.05	62.57	41.14 (58.52)	45.23 (64.70)	57.56	65.38	64.10	46.90	60.80	64.86	72.06	72.85	73.95 ^{+13.15}	
	Token-F1 _{open} ↑	1.18	2.09	5.32	1.23 (5.14)	1.41 (4.69)	2.32	5.06	25.82	1.94	7.25	16.33	3.51	1.62	18.95 ^{+11.70}	
PMC-VQA		Acc ↑	50.20	48.60	51.30	25.60 (48.89)	30.40 (51.56)	13.85	53.20	61.30	31.15	54.65	55.75	65.80	64.80	61.70 ^{+7.05}
Visual Grounding QA																
GEMeX-RMCoT		Acc ↑	41.67	38.33	40.00	26.67 (45.00)	21.67 (38.33)	43.33	41.67	40.00	33.33	36.67	41.67	56.67	58.33	86.67 ^{+58.00}
Held-out In-house Testset																
		Acc ↑	41.50	48.40	52.90	38.20 (51.50)	39.70 (52.70)	45.70	66.10	73.60	47.10	52.60	55.30	73.80	72.50	80.50 ^{+27.90}

and is not fine-tuned on our data. To separately isolate the contribution of our latent reasoning framework from that of the training data, we train a matched verbalized-CoT baseline using the same Qwen3-VL-8B backbone, LoRA configuration, and 61K samples as VITAL, but generating the reasoning chain as explicit text tokens. VITAL outperforms this baseline by 13.21 points on the in-domain testsets, 11.75 points on VQA-RAD, and 26.40 points on the held-out in-house testset, while being $58.4\times$ faster (353 ms vs. 20.6 s). These controlled results demonstrate that VITAL’s gains arise from its latent reasoning framework with dual supervision rather than from data exposure alone. Full results are provided in §C.1 and Table 13.

Qualitative Analysis. Figure 3 illustrates VITAL’s dual interpretability on a colonoscopy case asking about polyp appearance and location. Through progressive latent reasoning ($z_1 \rightarrow z_3$), the text decoder reveals an increasingly refined chain: z_1 identifies a localized mucosal lesion with reddish color and slight elevation; z_2 characterizes its elongated, smooth-surfaced, polyp-like morphology; and z_3 pinpoints its position on the left mucosal fold. Simultaneously, the visual projector heatmaps progressively sharpen from a broad activation over the mucosal region to a tight focus on the polyp itself, confirming that latent states encode spatially grounded evidence. The final answer correctly synthesizes all four aspects (location, color, surface texture, and protrusion). In contrast, SIM-CoT hallucinates “multiple large polyps” with fabricated morphology in a wrong location; Coconut produces a vague, ungrounded description without

any spatial specificity; and LVR commits a severe factual error by diagnosing “malignant adenocarcinoma” in the “gastric wall”, misidentifying both the pathology and the organ. These failures highlight that, without visual-semantic dual supervision, latent reasoning can collapse into text-dominant hallucination or lose medical factual grounding.

Quantitative Interpretability. Beyond the qualitative cases in Figure 3, we quantitatively evaluate whether the re-attached auxiliary modules faithfully recover information from VITAL’s latent states. On 9,582 valid step-level pairs, the text decoder achieves an overall BERTScore-F1 of 85.14 and ROUGE-L of 52.02, indicating that the semantic content of the reasoning process remains recoverable. For visual interpretation, the final latent state achieves 74.2 pointing accuracy and 51.2 patch IoU against ground-truth masks, with localization improving consistently across reasoning steps. These results provide quantitative support for VITAL’s textual and visual interpretability. Full metric definitions and step-wise results are reported in §C.8.

4.3 Ablation Studies

Effect of dual supervision. We ablate the two auxiliary supervision signals by removing them individually and jointly. Table 2 reports average accuracy on three representative evaluation groups: in-domain testsets, VQA-RAD (closed-ended), and the held-out in-house testset. Task-only latent training (b) improves over the untuned backbone (a) on in-domain and in-house sets but drops on VQA-RAD (-9.32), indicating that unsupervised




Figure 3: **Qualitative comparison.** VITAL achieves precise visual grounding and accurate diagnosis via progressive latent reasoning ($Z_1 \rightarrow Z_3$). Conversely, base-lines struggle with hallucinations, ungrounded vagueness, or medical factual errors.

latent states can overfit to the training distribution. Adding either supervision alone restores and extends gains: $\mathcal{L}_{\text{visual}}$ (d) contributes more on in-domain (+22.32 over b) and VQA-RAD (+13.31), while $\mathcal{L}_{\text{text}}$ (c) provides moderate but consistent improvement across all sets. Crucially, the full dual supervision (e) outperforms both single-signal variants by a large margin, confirming that the two losses are complementary rather than redundant: semantic supervision prevents reasoning degradation while visual supervision anchors latent states to spatial evidence. Ablation on λ_1, λ_2 sensitivity are in §C.2.

Effect of reasoning depth K . We vary the number of latent reasoning steps K from 0 (no latent) to 4. As shown in Table 3, $K=0$ substantially underperforms all latent variants, confirming that the recurrent latent loop is essential. Performance improves monotonically with K : the largest single-step gain occurs from $K=1$ to $K=2$ (+16.04 on in-domain, +10.10 on in-house), after which returns diminish: $K=3$ to $K=4$ gains only

Table 2: **Ablation on dual supervision components.** Avg. accuracy (%) is reported.

Variant	$\mathcal{L}_{\text{task}}$	$\mathcal{L}_{\text{text}}$	$\mathcal{L}_{\text{visual}}$	In-D.	RAD	In-H.
(a) Backbone				33.97	67.63	52.60
(b) Task-Only	✓			45.47	58.31	56.60
(c) +Semantic	✓	✓		52.29	62.08	64.10
(d) +Visual	✓		✓	67.79	71.62	70.60
(e) VITAL	✓	✓	✓	81.08	74.28	80.50

Table 3: **Ablation on the number of latent steps K and curriculum learning strategy.** Avg. accuracy (%) is reported.

K	In-D.	RAD	In-H.	Strategy	In-D.	RAD	In-H.
0	41.52	69.62	63.80	(a) Full-mix	44.82	65.14	50.70
1	58.67	71.40	68.50	(b) 2-Phase	56.15	68.33	64.10
2	74.71	72.95	78.60	(c) 3-Phase	81.08	74.28	80.50
3	77.03	74.06	79.70	(d) Reverse	35.23	47.15	41.40
4	81.08	74.28	80.50				

+4.05/+0.22/+0.80 across the three sets. We use $K=4$ as the default since it achieves the best accuracy with still-acceptable latency (see §4.3).

Inference efficiency. Figure 4 compares VITAL against Direct Answer ($K=0$) and Qwen3-VL-8B-Thinking, a zero-shot explicit-CoT model from the same backbone family that is labeled “Explicit CoT” in the figure. This zero-shot Thinking-CoT baseline generates verbose reasoning chains, incurring 34.1s latency, which is **97** $\times$ slower than VITAL ($K=4$, 353ms), yet achieves only 38.35 accuracy, far below even the no-latent baseline. In contrast, VITAL at $K=4$ reaches 81.08 accuracy with latency comparable to direct answer (353ms vs. 325ms), demonstrating that latent reasoning delivers substantial quality gains at negligible cost. Notably, VITAL at $K=1$ is actually faster than direct answer (241ms) due to shorter output sequences, while already improving accuracy by +17.15.

Effect of curriculum strategy. We compare four training schedules in Table 3: (a) full-mix trains on all $K \in \{1, 2, 3, 4\}$ from the start; (b) 2-phase splits into $K \leq 2$ then all K ; (c) 3-phase (our default) progressively introduces $K=1$, $K \leq 2$, then all K ; (d) reverse starts from $K=4$ and decreases. Full-mix training struggles to optimize long chains from scratch, while the reverse curriculum performs worst overall (35.23 on in-domain), confirming that an easy-to-hard progression is critical for stable latent state learning. The 3-phase schedule outperforms 2-phase by +24.93/+5.95/+16.40, indicating that isolating single-step reasoning in




Figure 4: **Inference efficiency.** (a) Average output tokens on a logarithmic scale. (b) In-domain accuracy vs. inference latency on a logarithmic scale. “Zero-shot Explicit CoT” denotes Qwen3-VL-8B-Thinking without fine-tuning on our 61K dataset. VITAL achieves substantially higher accuracy with concise outputs; at $K=4$, it is $\sim 97\times$ faster than zero-shot Explicit CoT.

Phase 1 builds a stronger foundation before introducing multi-step chains. More ablation studies are presented in §C.

5 Conclusion

We presented VITAL, a latent-space reasoning framework for medical MLLMs that introduces visual-semantic dual supervision to prevent modality collapse while maintaining strict train–inference parity and zero inference overhead. The auxiliary text decoder and visual projector serve as training-time scaffolding that can be re-attached post-hoc for dual interpretability, revealing both the reasoning logic and spatial attention of latent states. VITAL consistently outperforms all latent reasoning baselines and medical MLLMs on 7 benchmarks, achieving performance competitive with trillion-parameter proprietary models and surpassing them on visual grounding. We hope VITAL demonstrates that effective latent reasoning in multimodal models requires explicit grounding in both semantic and visual spaces, opening a path toward efficient, interpretable clinical AI. Along this path, we envision models that learn when to stop thinking, adaptively allocating reasoning depth to question complexity, and that draw supervision from richer narrative sources, bridging the gap between pixel-level evidence and clinical language at scale.

6 Acknowledgements

This work is supported by the Zhejiang Key Laboratory Project (2024E10001).

7 Limitations

Data scale and modality coverage. Although our 61K five-tuple dataset is the largest medical visual latent reasoning dataset to date, it covers only 9 imaging modalities drawn from two public segmentation datasets (MSD and BiomedParse). Generalization to rare modalities (e.g., ultrasound elastography, OCT) and uncommon diseases has not been fully validated.

Dependence on teacher-generated reasoning chains. The training reasoning chains are distilled from a proprietary MLLM teacher. The quality ceiling of VITAL’s latent states is therefore bounded by the teacher’s medical reasoning ability, and potential errors or biases in the teacher may propagate through distillation.

Fixed reasoning depth. The current design uses a fixed maximum number of latent steps $K=4$ regardless of question difficulty. Simple questions may not require multi-step reasoning, while more complex cases could benefit from additional steps. Exploring adaptive early-exit mechanisms is a promising direction for future work.

Open-ended generation. Experiments show that VITAL’s open-ended Token-F1 trails medical MLLMs specifically optimized for verbose free-text generation. This reflects an inherent trade-off of the latent reasoning paradigm, which prioritizes answer correctness over output verbosity.

8 Ethical Considerations

Not intended for clinical deployment. VITAL is a research contribution and should not be used directly for clinical diagnostic decisions. Any real-world application would require rigorous clinical validation, regulatory approval, and deployment under physician supervision.

Data and privacy. All training data are derived from publicly released, de-identified medical datasets (MSD and BiomedParse). No patient-identifiable information is used during teacher distillation or model training.

Limitations of interpretability outputs. While VITAL provides dual interpretability through re-attachable auxiliary modules, the decoded reasoning chains and visual projector outputs are approximate visualizations of latent states rather than faithful representations of the model’s true decision process. Users should not treat these explanations as the sole basis for clinical judgment, nor should they over-rely on the completeness of the interpretability outputs.

References

- Anthropic. 2026. [Introducing claude opus 4.6](#).
- Michela Antonelli, Annika Reinke, Spyridon Bakas, Keyvan Farahani, Annette Kopp-Schneider, Bennett A Landman, Geert Litjens, Bjoern Menze, Olaf Ronneberger, Ronald M Summers, et al. 2022. The medical segmentation decathlon. *Nature communications*, 13(1):4128.
- Shuai Bai, Yuxuan Cai, Ruizhe Chen, Keqin Chen, Xionghui Chen, Zesen Cheng, Lianghao Deng, Wei Ding, Chang Gao, Chunjiang Ge, et al. 2025. Qwen3-vl technical report.
- Junying Chen, Chi Gui, Ruyi Ouyang, Anningzhe Gao, Shunian Chen, Guiming Hardy Chen, Xidong Wang, Ruifei Zhang, Zhenyang Cai, Ke Ji, Guangjun Yu, Xiang Wan, and Benyou Wang. 2024. [Huatuogpt-vision, towards injecting medical visual knowledge into multimodal llms at scale](#). *Preprint*, arXiv:2406.19280.
- Ankan Deria, Komal Kumar, Adinath Madhavrao Dukre, Eran Segal, Salman Khan, and Imran Razak. 2026. Medmo: Grounding and understanding multimodal large language model for medical images.
- Lin Fan, Yafei Ou, Zhipeng Deng, Pengyu Dai, Hou Chongxian, Jiale Yan, Yaqian Li, Kaiwen Long, Xun Gong, Masayuki Ikebe, et al. 2026. Step-cot: Step-wise visual chain-of-thought for medical visual question answering.
- Xiaotang Gai, Chenyi Zhou, Jiaxiang Liu, Yang Feng, Jian Wu, and Zuozhu Liu. 2025. [MedThink: A rationale-guided framework for explaining medical visual question answering](#). In *Findings of the Association for Computational Linguistics: NAACL 2025*, pages 7453–7465, Albuquerque, New Mexico. Association for Computational Linguistics.
- Google DeepMind. 2026. [Gemini 3.1 pro model card](#). Technical report, Google DeepMind.
- Shibo Hao, Sainbayar Sukhbaatar, DiJia Su, Xian Li, Zhiting Hu, Jason Weston, and Yuandong Tian. 2024. Training large language models to reason in a continuous latent space. *arXiv preprint arXiv:2412.06769*.
- Xuehai He, Yichen Zhang, Luntian Mou, Eric Xing, and Pengtao Xie. 2020. Pathvqa: 30000+ questions for medical visual question answering.
- Xiaoshuang Huang, Lingdong Shen, Jia Liu, Fangxin Shang, Hongxiang Li, Haifeng Huang, and Yehui Yang. 2025. Towards a multimodal large language model with pixel-level insight for biomedicine. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 39, pages 3779–3787.
- Juntao Jiang, Jiangning Zhang, Yali Bi, Jinsheng Bai, Weixuan Liu, Weiwei Jin, Zhucun Xue, Yong Liu, Xiaobin Hu, and Shuicheng Yan. 2026a. M3cotbench: Benchmark chain-of-thought of mllms in medical image understanding. *arXiv preprint arXiv:2601.08758*.
- Yankai Jiang, Qiaoru Li, Binlu Xu, Haoran Sun, Chao Ding, Junting Dong, Yuxiang Cai, Xuhong Zhang, and Jianwei Yin. 2026b. Ibisagent: Reinforcing pixel-level visual reasoning in mllms for universal biomedical object referring and segmentation.
- Xin Lai, Zhuotao Tian, Yukang Chen, Yanwei Li, Yuhui Yuan, Shu Liu, and Jiaya Jia. 2024. Lisa: Reasoning segmentation via large language model. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 9579–9589.
- Jason J Lau, Soumya Gayen, Asma Ben Abacha, and Dina Demner-Fushman. 2018. A dataset of clinically generated visual questions and answers about radiology images. *Scientific data*, 5(1):180251.
- Khai Le-Duc, Duy MH Nguyen, Phuong TH Trinh, Tien-Phat Nguyen, Nghiem T Diep, An Ngo, Tung Vu, Trinh Vuong, Anh-Tien Nguyen, Mau Nguyen, et al. 2025. S-chain: Structured visual chain-of-thought for medicine. *arXiv preprint arXiv:2510.22728*.
- Bangzheng Li, Ximeng Sun, Jiang Liu, Ze Wang, Jialian Wu, Xiaodong Yu, Hao Chen, Emad Barsoum, Muhao Chen, and Zicheng Liu. 2025. [Latent visual reasoning](#). *Preprint*, arXiv:2509.24251.

- Binxu Li, Tiankai Yan, Yuanling Pan, Jie Luo, Ruiyang Ji, Jiayuan Ding, Zhe Xu, Shilong Liu, Haoyu Dong, Zihao Lin, et al. 2024. Mmedagent: Learning to use medical tools with multi-modal agent.
- Chunyu Li, Cliff Wong, Sheng Zhang, Naoto Usuyama, Haotian Liu, Jianwei Yang, Tristan Naumann, Hoifung Poon, and Jianfeng Gao. 2023. [Llava-med: Training a large language-and-vision assistant for biomedicine in one day](#). In *Advances in Neural Information Processing Systems*, volume 36, pages 28541–28564. Curran Associates, Inc.
- Bo Liu, Xiangyu Zhao, Along He, Yidi Chen, Huazhu Fu, and Xiao-Ming Wu. 2025. Gemex-rmcot: An enhanced med-vqa dataset for region-aware multimodal chain-of-thought reasoning. In *Proceedings of the 33rd ACM International Conference on Multimedia*, pages 13213–13220.
- Guimeng Liu, Tianze Yu, Somayeh Ebrahimkhani, Lin Zhi Zheng Shawn, Kok Pin Ng, and Ngai-Man Cheung. 2026. How do medical mllms fail? a study on visual grounding in medical images.
- Lingxiao Luo, Bingda Tang, Xuanzhong Chen, Rong Han, and Ting Chen. 2025. Vividmed: Vision language model with versatile visual grounding for medicine.
- OpenAI. 2026. [Gpt-5.4 thinking system card](#). Technical report, OpenAI.
- Tan-Hanh Pham and Chris Ngo. 2025. Multimodal chain of continuous thought for latent-space reasoning in vision-language models. *arXiv preprint arXiv:2508.12587*.
- Jinhao Qiao, Sihan Li, Jiang Liu, Heng Yu, Yi Xiao, Hongshan Yu, and Yan Zheng. 2025. Med-scot: Structured chain-of-thought reasoning and evaluation for enhancing interpretability in medical visual question answering. *Computerized Medical Imaging and Graphics*, page 102659.
- Hanoona Rasheed, Muhammad Maaz, Sahal Shaji, Abdelrahman Shaker, Salman Khan, Hisham Cholakkal, Rao M. Anwer, Eric Xing, Ming-Hsuan Yang, and Fahad S. Khan. 2024. Glamm: Pixel grounding large multimodal model. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 13009–13018.
- Andrew Sellergren, Sahar Kazemzadeh, Tiam Jaroensri, Atilla Kiraly, Madeleine Traverse, Timo Kohlberger, Shawn Xu, Fayaz Jamil, Cían Hughes, Charles Lau, et al. 2025. Medgemma technical report. *arXiv preprint arXiv:2507.05201*.
- Xuan Shen, Yizhou Wang, Yufa Zhou, Xiangxi Shi, Pu Zhao, Yanzhi Wang, and Jiuxiang Gu. 2025a. Efficient reasoning with hidden thinking. *arXiv preprint arXiv:2501.19201*.
- Zhenyi Shen, Hanqi Yan, Linhai Zhang, Zhanghao Hu, Yali Du, and Yulan He. 2025b. Codi: Compressing chain-of-thought into continuous space via self-distillation. In *Proceedings of the 2025 Conference on Empirical Methods in Natural Language Processing*, pages 677–693.
- Yan Shu, Chi Liu, Robin Chen, Derek Li, and Bryan Dai. 2025. Fleming-vl: Towards universal medical visual reasoning with multimodal llms. *arXiv preprint arXiv:2511.00916*.
- Guoxin Wang, Jun Zhao, Xinyi Liu, Yanbo Liu, Xuyang Cao, Chao Li, Zhuoyun Liu, Qintian Sun, Fangru Zhou, Haoqiang Xing, et al. 2025a. Citrus-v: Advancing medical foundation models with unified medical image grounding for clinical reasoning.
- Yuan Wang, Jiaxiang Liu, Shujian Gao, Bin Feng, Zhihang Tang, Xiaotang Gai, Jian Wu, and Zuozhu Liu. 2025b. V2t-cot: From vision to text chain-of-thought for medical reasoning and diagnosis. In *International Conference on Medical Image Computing and Computer-Assisted Intervention*, pages 658–668. Springer.
- Xilin Wei, Xiaoran Liu, Yuhang Zang, Xiaoyi Dong, Yuhang Cao, Jiaqi Wang, Xipeng Qiu, and Dahua Lin. 2025. [Sim-cot: Supervised implicit chain-of-thought](#). *Preprint*, arXiv:2509.20317.
- Linshan Wu, Yuxiang Nie, Sunan He, Jiaxin Zhuang, Luyang Luo, Tao Li, Zhuoyao Xie, Dexuan Chen, Yinghua Zhao, Neeraj Mahboobani, et al. 2025. Unibiomed: A universal foundation model for grounded biomedical image interpretation.
- Suyang Xi, Songtao Hu, Yuxiang Lai, Wangyun Dan, Yaqi Liu, Shansong Wang, and Xiaofeng Yang. 2026. Medlvr: Latent visual reasoning for reliable medical visual question answering. *arXiv preprint arXiv:2604.09757*.
- Weiwen Xu, Hou Pong Chan, Long Li, Mahani Aljunied, Ruifeng Yuan, Jianyu Wang, Chenghao Xiao, Guizhen Chen, Chaoqun Liu, Zhaodonghui Li, et al. 2025. Lingshu: A generalist foundation model for unified multimodal medical understanding and reasoning.
- Sheng Zhang, Yanbo Xu, Naoto Usuyama, Hanwen Xu, Jaspreet Bagga, Robert Tinn, Sam Preston, Rajesh Rao, Mu Wei, Naveen Valluri, et al. 2023. Biomed-clip: a multimodal biomedical foundation model pretrained from fifteen million scientific image-text pairs. *arXiv preprint arXiv:2303.00915*.
- Xiaoman Zhang, Chaoyi Wu, Ziheng Zhao, Weixiong Lin, Ya Zhang, Yanfeng Wang, and Weidi Xie. 2024. Pmc-vqa: Visual instruction tuning for medical visual question answering, 2024.
- Theodore Zhao, Yu Gu, Jianwei Yang, Naoto Usuyama, Ho Hin Lee, Sid Kiblawi, Tristan Naumann, Jianfeng Gao, Angela Crabtree, Jacob Abel, et al. 2025. A foundation model for joint segmentation, detection

and recognition of biomedical objects across nine modalities. *Nature methods*, 22(1):166–176.

A Training Data Construction

VITAL requires each training sample to be a five-tuple $(\mathbf{I}, \mathbf{q}, \mathbf{a}, \{\mathbf{e}_k\}_{k=1}^K, \mathbf{f}_{\text{ROI}})$: a medical image, a question, the final answer, a K -step reasoning chain, and a pre-extracted ROI visual feature. We construct an automated, fully reproducible pipeline from two public medical segmentation datasets, producing approximately 61K high-quality samples in total. This appendix describes each stage in detail.

A.1 Source Datasets and Preprocessing

MSD. We use the Medical Segmentation Decathlon (MSD) (Antonelli et al., 2022) as our primary source of normal-anatomy data. MSD provides 3D CT/MRI volumes with voxel-level segmentation annotations for 10 organ or tumor classes: liver, lung, pancreas, hepatic vessels, spleen, colon, heart, brain tumor, prostate, and hippocampus. For each volume, we extract 2D axial slices and retain only those containing at least 100 pixels of the target organ mask, yielding high-quality image-mask pairs. Since MSD annotations cover only the target regions without pathological findings, we apply an additional pathology-term filter during teacher distillation (§A.3) to prevent hallucinated clinical descriptions.

BiomedParse. To cover pathological findings and diverse biomedical structures beyond normal anatomy, we incorporate the BiomedParse dataset (Zhao et al., 2025), a large-scale multi-modal biomedical segmentation corpus spanning multiple target types: organ, lesion, tumor, vessel, tissue, cell, and retinal structure. Unlike MSD, BiomedParse samples frequently contain abnormal findings, allowing us to train the model on clinically relevant reasoning such as lesion characterization and differential description.

From segmentation to QA tuples. Starting from the raw segmentation annotations, we first construct a four-tuple for each sample: (image_path, mask, target_name, target_type). Questions are then auto-generated using a template-based strategy conditioned on the target type and imaging modality. We define six question categories, namely *yes/no*, *identify*, *location-choice*, *location*, *describe*, and *reasoning*, each associated with a target reasoning depth $K \in \{1, \dots, 4\}$

(Table 4). The result is a preliminary dataset containing images, masks, user questions with placeholder answers, and metadata (question type, target name, target type, imaging modality).

Table 4: **Question types and associated reasoning depth.**

Type	K	Example Trigger
Yes/No	1	Is the liver visible in...
Identify	1	What organ is shown...
Loc.-Choice	1	Is X in left, right, or center...
Location	2	Where is X located...
Describe	2–3	Describe the appearance of...
Reasoning	3–4	Analyze the visual findings...

A.2 Teacher Distillation via Proprietary MLLM

We employ a state-of-the-art proprietary MLLM as the teacher model to generate high-quality reasoning chains and final answers for each sample.

Teacher model selection. The quality of distilled reasoning chains critically depends on the teacher’s medical visual reasoning capability. To select the most suitable teacher, we conduct a pilot study on a random subset of 2,000 samples from our preliminary dataset, evaluating three frontier MLLMs (Gemini-3.1-Pro-Preview (Google DeepMind, 2026), GPT-5.4 (OpenAI, 2026), and Claude-Opus-4.6 (Anthropic, 2026)) under identical prompting conditions (same system prompt, overlay images, and quality filters). We measure: (i) *reasoning accuracy*, defined as whether the final answer correctly identifies/locates the target, judged by a medical annotator; (ii) *leakage rate*, i.e., the fraction of outputs containing forbidden annotation-related terms; and (iii) *format compliance*, i.e., the fraction of outputs that are valid JSON with correct step counts on the first attempt.

Table 5: **Teacher model comparison on 2,000-sample pilot set.**

Model	Acc. (%) $\uparrow$	Leak. (%) $\downarrow$	Format (%) $\uparrow$
Gemini-3.1-Pro-Preview	37.85	35.10	44.10
GPT-5.4	89.35	9.45	97.20
Claude-Opus-4.6	66.60	12.15	78.50

As shown in Table 5, GPT-5.4 achieves the best overall performance across all three metrics, combining the highest reasoning accuracy with the lowest annotation leakage rate and near-perfect format compliance. We therefore adopt GPT-5.4 as our teacher model for the full-scale distillation pipeline.

Teacher-only overlay. A key challenge in teacher distillation is to generate reasoning chains that are both *spatially grounded* and *annotation-free*. If the teacher only observes the raw medical image, it may fail to localize the exact target region, especially for small lesions, subtle tissues, or ambiguous anatomical structures. In contrast, if the teacher explicitly refers to the segmentation mask or annotated region, the generated reasoning may leak supervision cues and become unsuitable for training a student model that only receives raw images.

To address this issue, we adopt a *teacher-only overlay* mechanism. For each training sample, the binary segmentation mask is rendered on top of the original image as a semi-transparent red overlay with transparency $\alpha = 0.4$. In our implementation, the overlay color is set to RGB (255, 0, 0), and the mask boundary is additionally highlighted with a yellow contour, RGB (255, 255, 0), to help the teacher localize the target region more reliably. The resulting overlay image is shown *only* to the teacher model during reasoning-chain distillation, while the student model is trained and evaluated using the original unannotated image.

This design creates two distinct visual views for the same sample. The *teacher view* contains privileged localization information through the semi-transparent target overlay, whereas the *student view* contains only the raw image without any annotation. During distillation, the teacher is explicitly instructed to use the overlay only for internal localization and to generate the final answer and reasoning chain *as if it only sees the raw, unannotated image*. Figure 5 illustrates the difference between the teacher view and the student view. As a result, the distilled reasoning can remain grounded in the true target location while avoiding direct leakage of mask-, overlay-, or annotation-related cues.

System prompt design. The teacher prompt is designed to generate reasoning chains that are correct, image-grounded, and free from teacher-view leakage. In each API call, the teacher model receives a teacher-only overlay image, the hidden target identity, the target type, and the student-visible question. The textual prompt is organized into the following components.

1. **Role definition.** The system prompt defines the model as a teacher for medical visual reasoning data generation. The teacher is instructed to generate concise supervision for




Figure 5: **Comparison between the teacher view and the student view.** The teacher view contains a semi-transparent target overlay for localization guidance, while the student view contains only the raw unannotated image.

latent reasoning training, including a final answer and a reasoning chain.

2. **Teacher-only information usage.** The teacher is allowed to use teacher-only information internally, including the overlay image and the hidden target identity, to localize the target region correctly. However, the output must be written as if it were produced by a student model that can only observe the original unannotated image.
3. **Annotation-leakage prevention.** The prompt explicitly forbids any mention or implication of hidden guidance, special markings, annotations, masks, overlays, highlighted areas, segmentation maps, ROI, labels, ground truth, teacher-only metadata, or extra visual cues. It also forbids statements suggesting that the correct target identity was provided beforehand.
4. **Image-grounded medical reasoning.** The teacher is required to base the answer and reasoning chain on visible image evidence, such as location, shape, boundary, extent, density or signal intensity, texture, and relationships to nearby structures. For lesion or finding targets, the teacher may name the given target itself, but should not introduce unsupported diagnosis, malignancy or benignity, severity, prognosis, treatment, or other clinical claims unless they are directly supported by visible evidence.
5. **Spatial convention.** For spatial descriptions, the teacher must use image-space wording only, such as left, right, center, upper,

and lower parts of the image. Patient-space or radiological-convention wording, such as patient-left, patient-right, anatomical-left, or anatomical-right, is explicitly disallowed.

6. **Output format.** The teacher must return strict JSON with exactly two fields: `final_answer` and `reasoning_chain`. Markdown, code fences, and additional explanations are not allowed.

The sample-specific user prompt further injects the target identity, target type, question type, and student-visible question. It also appends a question-type-specific guidance block that specifies the required reasoning-step range and provides an in-context JSON example. This design separates persistent global constraints from per-sample control signals: the system prompt enforces annotation-free medical reasoning behavior, while the user prompt controls the current target, question, and expected reasoning granularity. If a generated response fails format checking, step-count validation, or leakage filtering, a retry prompt is appended to require regeneration from scratch under stricter reminders.

Question-type-aware prompting. To control the granularity of distilled reasoning chains, we use question-type-aware prompting. Each sample is assigned a question type either from the pre-generated metadata or, if unavailable, by rule-based classification from the question text. The BiomedParse question-generation script uses target-type- and modality-aware templates to generate five teacher-compatible question categories: *yesno*, *identify*, *location*, *describe*, and *reasoning*. These templates are adapted to heterogeneous biomedical targets, including organs, anatomical sub-structures, vessels, tissues, lesions or findings, retinal structures, and microscopy-level structures. During teacher generation, an additional *location_choice* category is recognized for questions that ask the model to choose among coarse image regions such as left, right, or center.

For each question type, the prompt specifies an explicit reasoning-step target.

1. **yesno.** Simple visual confirmation questions use exactly one reasoning step. The reasoning chain only needs to provide direct visual evidence supporting the yes/no answer.

2. **identify.** Identification questions also use exactly one reasoning step. The teacher is expected to justify the recognized target based on its visible appearance and anatomical context.

3. **location_choice.** Coarse spatial-choice questions use exactly one reasoning step. The final answer must be a complete sentence rather than a bare word such as “left”, “right”, or “center”.

4. **location.** Spatial localization questions use exactly two reasoning steps. This allows the teacher to first identify the target region and then describe its image-space position relative to surrounding anatomy.

5. **describe.** Descriptive questions use two to three reasoning steps. The reasoning chain should cover the target’s appearance, shape, boundary, extent, and other visible characteristics.

6. **reasoning.** More complex visual interpretation questions use three to four reasoning steps. These questions require a fuller analysis of the target’s location, morphology, visual contrast, and relationship to nearby structures.

Each question-type guidance block contains three elements: the normalized question type, the required reasoning-step range, and an in-context JSON example showing the desired output style. This makes both the output format and the reasoning depth explicit for the teacher model. After generation, the number of produced reasoning steps is checked against the allowed range for the corresponding question type. The validated step count is stored as `latent_steps`, and the reasoning chain is stored as `text_reasoning_chain`. In this way, question-type-aware prompting directly determines the length of the latent reasoning supervision while preserving a consistent JSON format across all samples. Figure 6 illustrates a concrete prompt example used in our teacher-distillation pipeline.

Dataset-specific adaptations. For organs (normal anatomy), we additionally filter outputs against a pathology-style term list (e.g., *lesion*, *mass*, *tumor*, *calcification*, *inflammation*) to prevent the teacher from hallucinating clinical findings on healthy organ data. For other targets (which includes pathological objectives), pathology terms




Figure 6: **Prompt design example for teacher distillation.** The teacher receives three sources of information in a single API call: a global system prompt that specifies the medical reasoning role, annotation-leakage constraints, image-space spatial convention, and strict JSON output format; a question-type-aware user prompt that provides the hidden target identity, target type, reasoning-step target, student-visible question, and an in-context JSON example; and a teacher-only overlay image that is used only for internal localization. The teacher MLLM then generates a student-view JSON output containing final_answer and reasoning_chain, written as if only the raw unannotated image were visible.

are permitted and expected. Both variants share the same annotation-leakage filter and spatial-consistency checks.

A.3 Multi-Round Quality Control

Each teacher-generated sample undergoes a six-round automated verification pipeline. Samples failing any round are automatically retried with progressively stricter prompt constraints (up to 4 retries).

- 1. JSON parse validation.** Verify that the teacher output is valid JSON conforming to the required schema (final_answer: string, reasoning_chain: list of strings).
- 2. Forbidden-term filtering.** Scan both final_answer and all entries of reasoning_chain for any of the 30 annotation-leakage terms. Any hit triggers rejection and retry.
- 3. Pathology-style filtering (Organ only).** For the normal-anatomy subset, check for inappropriate use of pathology terms (e.g., describing a healthy liver as having a “lesion” or “mass”). This filter is disabled for BiomedParse where pathological descriptions are valid.
- 4. Location-mixed filtering.** Detect mixing of patient-space and image-space spatial references (e.g., “patient’s right” or “anatomical left”), which violates our image-space-only convention.
- 5. Step-count validation.** Verify that |reasoning_chain| falls within the target range [min_steps, max_steps] defined by the question type.
- 6. Answer normalization.** Standardize final_answer format according to question type: identification answers shorter than 4

words are expanded into complete sentences (e.g., “pancreas” → “The main organ shown is the pancreas.”); location answers are reformatted with explicit image-space phrasing; all answers are ensured to end with a period.

Retry mechanism. Upon rejection, the sample is re-submitted with an appended “strict rules” prompt block that explicitly lists the violation and demands regeneration from scratch. After 4 failed attempts, the sample is logged and excluded from the final dataset.

Human audit. After the automated pipeline completes, we conduct a 5% random audit of the final dataset. The observed annotation leakage rate is $< 0.3\%$, confirming the effectiveness of our multi-round filtering.

A.4 Dataset Statistics

The final VITAL dataset comprises **61,411** samples drawn from two public sources spanning 9 imaging modalities: CT, MRI, X-ray, ultrasound, endoscope, fundus photography, dermoscopy, OCT, and pathology. Figure 7 and Tables 6–7 provide a comprehensive breakdown by source, sub-set, train/val/test partition, and imaging modality.

Table 6: **Per-sub-set sample counts by reasoning depth K .**

Sub-set	$K=1$	$K=2$	$K=3$	$K=4$	Total
MSD	14,168	12,256	1,023	6,516	33,963
Radiography	1,320	1,698	188	950	4,156
amos22	1,124	1,439	132	730	3,425
CAMUS	647	834	27	488	1,996
COVID-QU-Ex	590	776	54	434	1,854
ACDC	580	747	26	393	1,746
MMs	580	746	29	388	1,743
CXR_Masks_and_Labels	497	599	52	361	1,509
NeoPolyp	444	609	52	368	1,473
kits23	469	551	69	338	1,427
ISIC	358	339	117	261	1,075
PanNuke	352	445	12	263	1,072
PolypGen	316	390	32	220	958
LGG	293	278	107	232	910
FH-PS-AOP	274	368	21	215	878
LIDC-IDRI	230	342	42	177	791
BreastUS	187	236	42	131	596
QaTa-COV19	97	133	11	73	314
G1020	73	112	12	68	265
siim-acr-pneumothorax	91	81	22	58	252
UWaterlooSkinCancer	69	81	24	53	227
COVID-19_CT	61	101	7	55	224
OCT-CME	74	83	2	46	205
REFUGE	61	79	5	55	200
GlaS	45	44	5	29	123
LiverUS	7	12	0	10	29
Total	23,007	23,379	2,113	12,912	61,411

Source diversity. The MSD portion (33,963 samples) is derived from the Medical Segmentation Decathlon, covering 10 organ classes from 3D

Table 7: **Sample distribution by imaging modality and reasoning depth K .**

Modality	$K=1$	$K=2$	$K=3$	$K=4$	Total
CT	13,622	12,470	1,071	6,618	33,781
MRI	3,883	3,990	364	2,211	10,448
X-Ray	2,595	3,287	327	1,876	8,085
Ultrasound	1,115	1,450	90	844	3,499
Endoscope	760	999	84	588	2,431
Dermoscopy	427	420	141	314	1,302
Pathology	397	489	17	292	1,195
Fundus	134	191	17	123	465
OCT	74	83	2	46	205
Total	23,007	23,379	2,113	12,912	61,411

CT/MRI volumes with dense axial-slice sampling. The BiomedParse portion (27,448 samples, 25 sub-sets) extends coverage to pathological findings and diverse biomedical structures across 9 imaging modalities (Table 7). CT and MRI together account for 72.0% of all samples, providing a strong foundation in cross-sectional imaging; X-ray (13.2%) and ultrasound (5.7%) offer complementary image projection and real-time modalities; and the remaining modalities (endoscope, dermoscopy, pathology, fundus, and OCT) collectively contribute 9.1%, ensuring the model is exposed to the full breadth of clinical imaging.

Reasoning depth distribution. Single-step samples ($K=1$, 37.5%) and two-step samples ($K=2$, 38.1%) together account for the majority, reflecting the prevalence of identification and localization questions. Multi-step reasoning ($K=3-4$, 24.5%) is concentrated in descriptive and analytical questions requiring richer clinical rationale. This distribution naturally aligns with our curriculum learning strategy (§3.5): Phase 1 and 2 can leverage the abundant $K \leq 2$ samples, while Phase 3 trains on the full spectrum.

Comparison with prior datasets. To our knowledge, this is the largest dataset specifically designed for training latent/chain-of-thought reasoning in the medical vision domain. Existing medical reasoning datasets are significantly smaller: MedThink (Gai et al., 2025) provides rationale annotations for $\sim 9.5\text{K}$ QA pairs (R-RAD: 3,515 + R-SLAKE: 5,980), and M3CoTBench (Jiang et al., 2026a) offers only 1,079 samples as an evaluation benchmark. Our dataset exceeds these efforts by $6 \times -57 \times$ in scale, while additionally providing aligned ROI visual features for every sample, a unique annotation absent from all prior medical




Figure 7: **Dataset statistics.** **Left:** Sample distribution by source and reasoning depth K . **Right:** Train / validation / test split by reasoning depth K .

reasoning datasets.

A.5 Adaptive ROI Feature Extraction

For each training sample, we extract a d_v -dimensional ROI feature from a frozen medical vision encoder as the supervision target for $\mathcal{L}_{\text{visual}}$. This section describes the encoder selection, the adaptive extraction strategy, and the hyperparameter tuning that together ensure high-quality visual supervision signals.

A.5.1 Vision Encoder Selection

We evaluate two biomedical vision encoders as candidates for ROI feature extraction: **BiomedCLIP** (Zhang et al., 2023), a CLIP-based model pretrained on 15M PubMed image–text pairs, and **MedSigLIP** (Sellersgren et al., 2025), a SigLIP variant medically tuned by Google as part of the MedGemma family. Table 8 summarizes the key architectural differences.

Table 8: **Comparison of candidate medical vision encoders.**

Property	BiomedCLIP	MedSigLIP
Input resolution	224×224	448×448
Patch size	16	14
Patch grid	14×14 = 196	32×32 = 1,024
Hidden dim (d_v)	768	1,152
CLS token	Yes	No (all spatial)

The critical advantage of MedSigLIP is its 5.2× higher spatial resolution: 1,024 patches versus 196. For a typical medical image (e.g., 1,024×1,024

pixels), each BiomedCLIP patch covers $\sim 73 \times 73$ pixels, whereas each MedSigLIP patch covers only $\sim 14 \times 14$ pixels. This dramatically improves the fidelity of mask-aligned patch selection: a small lesion occupying 2% of the image area would align with ~ 4 BiomedCLIP patches but ~ 20 MedSigLIP patches, yielding a far richer and more discriminative ROI feature after mean pooling. Additionally, SigLIP’s architecture produces only spatial patch tokens (no CLS token), which simplifies the mask-to-patch alignment pipeline. We therefore adopt **MedSigLIP** (SigLIP-SO400M/14-448) as our frozen vision encoder throughout all experiments. This choice is further validated by end-to-end ablation (§C.4), where MedSigLIP outperforms both BiomedCLIP and the backbone’s own encoder by a large margin on downstream accuracy.

Figure 8 provides a qualitative comparison on representative samples spanning CT, X-ray, endoscopy, and dermoscopy modalities. For each sample, we visualize the patch-level activation heatmap (cosine similarity between each patch feature and the mean ROI feature, overlaid on the original image). MedSigLIP’s activations are sharply concentrated on the target anatomical structure with minimal background noise, whereas BiomedCLIP’s lower-resolution grid produces diffuse, spatially imprecise activations that often spread into surrounding tissue. This contrast is especially pronounced for small or thin structures (e.g., vessels, polyps), where BiomedCLIP’s 196-patch grid lacks the granularity to isolate the target from its context.




Figure 8: **Qualitative comparison of patch-level activation heatmaps between MedSigLIP and BiomedCLIP.** MedSigLIP (**top row**) produces sharply localized activations concentrated on the target region, while BiomedCLIP (**bottom row**) exhibits diffuse, spatially imprecise responses due to its coarser 14×14 patch grid.

A.5.2 Extraction Pipeline

We extract patch-level features by hooking the output of the last vision encoder block, yielding a $1,024 \times 1,152$ feature map per image. The extraction follows an *adaptive strategy* governed by two hyperparameters: the **area threshold** T and the **crop margin ratio** P .

Full-image pathway (mask ratio $\geq T$). When the target structure is large relative to the image, we feed the entire image to MedSigLIP. The segmentation mask is downsampled to the 32×32 patch grid via area interpolation (binarized at 0.3), and the ROI feature is computed as the mean of all mask-aligned patch features.

Tight-crop pathway (mask ratio $< T$). When the target is small, a full-image approach would dilute the ROI signal among overwhelming background patches. Instead, we tightly crop the ROI using the mask bounding box expanded by a margin ratio P on each side, pad to a square with zero-fill, and re-extract patch features from the cropped image. A validity mask distinguishes real content from padding regions, and only valid patches contribute to the mean-pooled ROI feature.

Post-processing. All ROI features are ℓ_2 -normalized to the unit hypersphere. For each sample, we additionally store the top-5 patches (by cosine similarity with the ROI feature) along with their grid positions, enabling downstream visualization of which spatial locations contribute most to the feature representation. Each sample’s output is saved as a .pt file.

A.5.3 Hyperparameter Grid Search

The two hyperparameters, area threshold T and crop margin ratio P , jointly determine how and when the extraction strategy switches. We conduct a grid search over $T \in \{0.05, 0.10, 0.20\}$ and $P \in$

$\{0.0, 0.03, 0.05, 0.10\}$ on a random subset of 1,000 samples, evaluating three complementary metrics:

- **Coverage (%)**. Fraction of the 32×32 patch grid covered by the ROI mask after downsampling. Higher values indicate denser spatial supervision.
- **Intensity**. Mean ℓ_2 -norm of patch activations within the ROI region, reflecting feature activation strength.
- **SNR**. Ratio of mean ROI activation to mean background activation, measuring how well the ROI stands out from the background.

Figure 9 presents the results as three heatmaps. Several trends emerge:

- **Coverage increases monotonically with T .** A higher threshold routes more samples to the full-image pathway, which naturally produces denser patch grids. However, excessively high coverage (e.g., $T=0.20$, $P=0.0$: 45.1%) comes at the cost of including many background patches.
- **Intensity peaks at moderate T .** The tight-crop pathway ($T=0.05$) yields the highest raw activation (2,767.7–2,771.9) because cropped images concentrate the target, but this advantage diminishes as background patches are included.
- **SNR favors higher T with small P .** The SNR peaks around $T=0.20$ (1.032–1.049), indicating that the full-image pathway preserves contextual contrast between ROI and background better than aggressive cropping.

We select **$T=0.20$** and **$P=0.05$** as the final configuration, as it achieves the best trade-off: competitive intensity (2,752.7), the highest SNR among




Figure 9: **Grid search over area threshold T and crop margin ratio P for adaptive ROI feature extraction.** **Left:** Coverage, where higher T triggers full-image extraction for more samples, increasing patch coverage. **Center:** Feature activation intensity within the ROI. **Right:** Signal-to-noise ratio (ROI vs. background activation). The selected configuration $T=0.20$, $P=0.05$ (highlighted) achieves the best balance: moderate coverage (39.4%), near-peak intensity (2,752.7), and the highest SNR (1.040) among configurations with $T \geq 0.10$.

non-trivial coverage levels (1.040), and sufficient coverage (39.4%) for robust mean pooling. This configuration is used for all 61,411 samples in the final dataset.

B Implementation Details

This appendix provides the complete training configuration, per-module parameter breakdown, and evaluation protocol details for reproducibility.

B.1 Model Architecture

VITAL is built on top of **Qwen3-VL-8B-Instruct** (Bai et al., 2025), a multimodal large language model comprising a vision encoder and a language backbone. After accounting for weight tying between `embed_tokens` and `lm_head`, the de-duplicated parameter count is 8.14B. Table 9 summarizes the architectural configuration.

Table 9: **Backbone architecture configuration.**

Property	LLM	Vision Encoder
Hidden size	4,096	1,152
Layers	36	27
Attention heads	32	16
KV heads (GQA)	8	—
Intermediate size	12,288	—
Patch size	—	16
Vocab size	—	151,936
Parameters	7,568M	576M

B.2 Trainable Modules

Table 10 gives the full parameter breakdown. Both the vision encoder and the LLM backbone are *frozen*; trainability is introduced through LoRA adapters and three auxiliary modules that serve as training-time scaffolding.

Table 10: **Parameter statistics of VITAL.** *Scaffolding modules* (marked with †) are discarded at inference. **Highlighted** rows indicate trainable modules.

Module	#Params	#Trainable	Status
Vision Encoder	576.4M	0	Frozen
LLM Backbone	7,568.4M	0	Frozen
+ LoRA ($r=64$, $\alpha=128$)	174.6M	174.6M	Trained
Visual Projector†	21.5M	21.5M	Trained
Projection $\pi_{in}^\dagger$	8.4M	8.4M	Trained
Text Decoder (Qwen3-1.7B)†	1,720.6M	1,720.6M	Trained
Total (de-duplicated)	10,069.9M	1,925.1M (19.1%)	

Visual Projector. The visual projector maps the LLM hidden state ($d=4,096$) to the MedSigLIP feature space ($d_v=1,152$) using a norm-stabilized residual MLP (21.51M parameters):

$$VP(z) = \text{LN}_{d_v} \left(W_2 \left(\text{GELU}(W_1 \text{LN}_d(z)) + z \right) \right),$$

where $W_1 \in \mathbb{R}^{4096 \times 4096}$ (16.78M), $W_2 \in \mathbb{R}^{1152 \times 4096}$ (4.72M), and a dropout of 0.1 is applied after GELU. The input LayerNorm eliminates the norm imbalance between LLM hidden states and vision features; the residual connection preserves representational richness; and the output LayerNorm re-scales to the ℓ_2 -normalized feature space.

Projection layer π_{in} . A single linear layer $\pi_{in} \in \mathbb{R}^{2048 \times 4096}$ with bias projects the LLM hidden state to the text decoder’s input dimension (8.39M parameters).

Auxiliary Text Decoder. We use Qwen3-1.7B as the auxiliary text decoder, fully fine-tuned during training. It is an independent causal language model with its own embedding layer and language model head (weight-tied), sharing no parameters with the backbone. Table 11 summarizes its configuration.

Table 11: **Auxiliary text decoder (Qwen3-1.7B) configuration.**

Property	Value
Hidden size	2,048
Layers	28
Attention heads	16
KV heads (GQA)	8
Intermediate size	6,144
Vocab size	151,936
embed_tokens	311.16M
Transformer layers	1,409.41M
lm_head	weight-tied
Total	1,720.57M

External vision encoder for ROI features. ROI visual features are extracted by a frozen MedSigLIP encoder (SigLIP-SO400M/14-448) (Seller-gren et al., 2025), which accepts 448×448 images and produces a $32 \times 32 = 1,024$ patch grid with hidden dimension $d_v = 1,152$. This encoder is used *only* for pre-extracting supervision targets and is not part of the VITAL model at either training or inference time. The adaptive extraction strategy is detailed in §A.5.

B.3 Training Configuration

Table 12 lists the complete set of training hyperparameters.

Table 12: **Training hyperparameters.**

Hyperparameter	Value
Optimizer	AdamW
β_1, β_2	0.9, 0.999
Learning rate	2×10^{-5}
Warmup ratio	0.05
Weight decay	0.01
Max gradient norm	1.0
Precision	bf16
λ_{text}	1.0
λ_{visual}	0.1
LoRA dropout	0.05
Curriculum schedule	
Phase 1 ($K=1$)	5 epochs, $\sim 19.6\text{K}$ samples
Phase 2 ($K \leq 2$)	5 epochs, $\sim 39.4\text{K}$ samples
Phase 3 (all K)	10 epochs, $\sim 52.2\text{K}$ samples
Hardware	6 $\times$ NVIDIA RTX PRO 6000
Total training time	~ 96 hours

Each curriculum phase warm-starts from the previous phase’s checkpoint. The loss weights $\lambda_{\text{text}}=1.0$ and $\lambda_{\text{visual}}=0.1$ are kept constant across all phases; sensitivity analysis is provided in §C.2.

B.4 Evaluation Protocol

In-domain testsets (MSD, BiomedParse). We use stratified held-out test splits (4,607 samples total; see Figure 7) and report overall accuracy. In evaluation, all the question-answer pairs are turned into closed-ended questions. The predicted answer is matched against the ground truth after lowercasing and whitespace normalization.

VQA-RAD, PathVQA. We follow the standard evaluation protocols established by each benchmark. For VQA-RAD (Lau et al., 2018) and PathVQA (He et al., 2020), we report accuracy for closed-ended subsets and **token-level F1** for open-ended subsets. Token-level F1 is computed by treating the predicted and ground-truth answers as bags of tokens (after lowercasing and whitespace tokenization), then calculating the harmonic mean of token-level precision and recall.

PMC-VQA. PMC-VQA (Zhang et al., 2024) is a multiple-choice benchmark; we report accuracy on the test split.

GEMeX-RMCoT. GEMeX-RMCoT (Liu et al., 2025) evaluates medical visual grounding through reasoning-based multiple-choice questions that require attending to specific anatomical regions. We report accuracy following the official evaluation script.

Held-out in-house testset. To mitigate evaluation bias caused by training-set contamination leading to inflated accuracy and token-F1 scores, we introduce a private held-out evaluation set. This test-set consists of 1,000 VQA pairs derived from brain neuroimaging data (MRI and PET scans), covering questions about anatomical identification, lesion localization, and clinical reasoning. All questions are formulated as closed-ended (multiple-choice or yes/no); we report accuracy. None of the evaluated baselines have been exposed to this data during training, ensuring a fair comparison.

C More Experiments

C.1 VITAL vs. Matched Verbalized-CoT Fine-Tuning

To distinguish the contribution of VITAL’s latent reasoning framework from that of the training data, we construct a matched verbalized-CoT baseline. Specifically, we fine-tune Qwen3-VL-8B with LoRA using the same 61K training samples and

the same LoRA configuration as VITAL, but require the model to generate each reasoning chain explicitly as text tokens before producing the final answer. This matched baseline should not be confused with Qwen3-VL-8B-Thinking in the main comparison and efficiency analysis: the latter is a zero-shot explicit-CoT model that is not fine-tuned on our 61K dataset. We additionally include a direct-answer LoRA baseline, the zero-shot Qwen3-VL-8B backbone, and the zero-shot Qwen3-VL-8B-Thinking model. This comparison controls for both the backbone and training data, isolating the effect of representing the reasoning process in explicit text versus continuous latent states with visual-semantic dual supervision.

Table 13: **Comparison with matched verbalized-CoT fine-tuning.** The matched verbalized-CoT baseline and VITAL use the same Qwen3-VL-8B backbone, LoRA configuration, and 61K dataset. Qwen3-VL-8B-Thinking instead performs zero-shot explicit CoT without training on this dataset. In-D. denotes average accuracy on the two in-domain testsets, and RAD denotes closed-ended accuracy on VQA-RAD. Accuracy values are percentages.

Method	Data	In-D.	RAD	In-H.	Latency
Qwen3-VL-8B (zero-shot)	–	33.97	67.63	52.60	286 ms
+ LoRA (direct answer)	61K	45.47	58.31	56.60	325 ms
+ LoRA (verbalized CoT, matched)	61K	67.87	62.53	54.10	20.6 s
Qwen3-VL-8B-Thinking (zero-shot)	–	38.35	64.52	55.30	34.1 s
VITAL	61K	81.08	74.28	80.50	353 ms

As shown in Table 13, VITAL substantially outperforms the matched verbalized-CoT baseline on all three representative evaluations, with absolute gains of 13.21 points on the in-domain testsets, 11.75 points on VQA-RAD, and 26.40 points on the held-out in-house testset. Because the two methods share the same backbone, LoRA configuration, and training samples, these results show that the gains cannot be attributed to the 61K dataset alone; the latent reasoning mechanism and its dual supervision provide a clear advantage over verbalizing the same reasoning chains.

The verbalized-CoT baseline improves in-domain accuracy over direct-answer LoRA (67.87 vs. 45.47), but does not transfer as reliably beyond the training distribution. On VQA-RAD, it scores 62.53, below the zero-shot backbone’s 67.63; on the held-out in-house testset, it also underperforms direct-answer LoRA (54.10 vs. 56.60). This pattern suggests that explicit rationale fine-tuning may overfit to the linguistic form of the training reasoning chains, whereas VITAL’s visual-semantic

constraints preserve stronger visual grounding and generalization.

VITAL is also substantially more efficient. The verbalized-CoT model requires 20.6 seconds to generate its reasoning tokens before completing an answer, while VITAL performs four single-state recurrent reasoning steps and completes inference in 353 ms, corresponding to a $58.4\times$ speedup. The zero-shot Qwen3-VL-8B-Thinking model is even slower at 34.1 seconds and remains less accurate than the matched verbalized-CoT model on the in-domain and in-house evaluations. Together, these results demonstrate that VITAL’s advantage arises from the reasoning framework rather than data exposure alone: continuous reasoning with visual-semantic dual supervision improves both predictive performance and inference efficiency over both zero-shot Thinking CoT and matched verbalized-CoT fine-tuning.

C.2 Ablation on Loss Weight Sensitivity

As a complement to the dual supervision ablation in §4.3 (Table 2), we perform a grid search over the loss weights λ_1 (semantic) and λ_2 (visual) to assess sensitivity. Figure 10 reports average accuracy across in-domain testsets for each (λ_1, λ_2) combination. The optimal setting is $\lambda_1=1.0$, $\lambda_2=0.1$ (81.08%). Performance is more sensitive to λ_2 : moderate values (0.05–0.3) consistently outperform extremes, as too-large λ_2 forces latent states toward the visual subspace at the expense of reasoning capacity. For λ_1 , values in $[0.3, 1.0]$ are relatively stable, but excessive semantic weight ($\lambda_1 \geq 2.0$) degrades performance, likely due to the auxiliary decoder dominating gradient flow and hindering task-oriented learning. We also observe that $\lambda_1=5.0$ collapses performance regardless of λ_2 , confirming that balanced dual supervision is critical.

C.3 Ablation on Visual Projector Architecture

We compare five architectural choices for the visual projector (VP) that maps latent states to the medical vision encoder’s feature space. Results are reported in Table 14. A plain linear layer (a) performs poorly (27.72%), confirming that bridging the LLM hidden space to the vision feature space requires non-linear capacity. Adding an MLP (b) doubles accuracy, and LayerNorm (c) provides a further +18.65 gain by resolving the norm imbalance between the two spaces, consistent with our motivation in §3.3. The residual connection (d)




Figure 10: **Loss weight sensitivity heatmap.** Each cell shows in-domain avg. accuracy (%). The best setting ($\lambda_1=1.0$, $\lambda_2=0.1$) is used as default.

yields the largest single-component boost (+12.25 over c), as it preserves the representational richness of latent states while the non-linear pathway learns a modality-bridging correction. Surprisingly, the Transformer Block (e) collapses to 46.36%, likely because the self-attention over a single token adds no useful inductive bias while introducing optimization instability.

Table 14: **Ablation on visual projector architecture.** In-domain avg. accuracy (%) is reported.

Variant	MLP	LN	Res.	Attn.	In-D.
(a) Linear					27.72
(b) MLP	✓				50.18
(c) MLP + LN	✓	✓			68.83
(d) Res. MLP + LN (ours)	✓	✓	✓		81.08
(e) Transformer Block	✓	✓	✓	✓	46.36

C.4 Ablation on Vision Encoder for Visual Supervision

A central claim of VITAL is that visual supervision should come from an *independent, external* medical vision encoder rather than the model’s own visual representations. To validate this, we compare three choices for the ROI feature source used in $\mathcal{L}_{\text{visual}}$: **(a) Qwen3-VL-8B** (self-referential): ROI features are extracted from the backbone’s own vision encoder by pooling over mask-aligned patch tokens. This mirrors the design of MedLVR (Xi et al., 2026), where the model aligns its latent states to its own visual representations. **(b) BiomedCLIP** (Zhang et al., 2023): an external CLIP-based encoder pretrained on 15M PubMed image-text pairs (224×224 , $14 \times 14 = 196$ patches, $d_v=768$). **(c) MedSigLIP** (Sellergren et al., 2025): an external SigLIP variant medically tuned as part of MedGemma (448×448 , $32 \times 32 = 1,024$ patches,

$d_v=1,152$). All three variants use the same training data, model architecture, and hyperparameters; only the frozen encoder providing the visual supervision target differs.

Table 15: **Ablation on vision encoder for visual supervision.** Self-referential supervision (the model’s own encoder) underperforms independent external encoders. Avg. accuracy (%) is reported.

Encoder	Type	In-D.	RAD	In-H.
(a) Qwen3-VL-8B	Self-referential	64.73	69.18	60.90
(b) BiomedCLIP	External	55.72	59.42	54.60
(c) MedSigLIP	External	81.08	74.28	80.50

Analysis. The results reveal a nuanced picture. MedSigLIP (c) achieves the best performance across all three evaluation sets by a large margin (+16.35 over Qwen3-VL-8B on in-domain, +19.60 on in-house), confirming the value of high-resolution external visual supervision. Interestingly, the self-referential variant (a) outperforms BiomedCLIP (b) despite providing a theoretically weaker supervision signal. We attribute this to BiomedCLIP’s low spatial resolution ($14 \times 14 = 196$ patches at 224×224 input): its coarse patch grid produces spatially imprecise ROI features that are too blurry to provide meaningful fine-grained guidance, effectively giving the model a noisy target that harms optimization. The self-referential approach, while limited by its circular nature, at least operates at the backbone’s native resolution and produces coherent patch features. MedSigLIP combines the best of both worlds: it is an independent external encoder (avoiding the trivial-satisfaction problem of self-referential alignment) with $5.2 \times$ higher spatial resolution than BiomedCLIP ($1,024$ patches at 448×448), yielding discriminative and precisely localized ROI features (see qualitative comparison in Figure 8). This result validates our design choice and provides direct empirical evidence that (i) external supervision is superior to self-referential alignment when the external encoder has sufficient spatial resolution, and (ii) spatial resolution of the supervision source is a critical factor for effective visual grounding in latent reasoning.

C.5 Ablation on ROI Extraction Strategy

See §A.5.3 for hyperparameter grid search results.

C.6 Ablation on Shared vs. Per-Step Visual Target

A natural question is whether providing step-specific visual supervision at different granularities could outperform a single shared ROI target. To investigate this, we design a *per-step visual target* variant that provides progressively finer-grained supervision to each latent step, and compare it against our default shared-target design.

Per-step target construction. For each training sample, we construct a sequence of $K=4$ visual targets via linear interpolation between a global image feature $\mathbf{f}_{\text{global}}$ and the ROI feature $\mathbf{f}_{\text{ROI}}$:

$$\mathbf{f}_k = \text{Normalize}(\alpha_k \cdot \mathbf{f}_{\text{ROI}} + (1 - \alpha_k) \cdot \mathbf{f}_{\text{global}}), \quad (8)$$

where $\alpha_k \in \{0.0, 0.33, 0.67, 1.0\}$ for $k = 1, \dots, 4$. The global feature $\mathbf{f}_{\text{global}} \in \mathbb{R}^{d_v}$ is computed as the ℓ_2 -normalized mean of all 1,024 patch features from the frozen MedSigLIP encoder (without any mask guidance). This creates a coarse-to-fine curriculum: z_1 is supervised toward a holistic image-level representation, while z_4 is supervised toward the precise ROI feature.

Training protocol. The per-step variant modifies only the visual loss computation: at each step k , the visual projector output is regressed against $\mathbf{f}_k$ instead of the shared $\mathbf{f}_{\text{ROI}}$. All other settings (model architecture, LoRA configuration, semantic supervision, curriculum schedule, and hyperparameters) remain identical to the default VITAL configuration.

Table 16: **Ablation on shared vs. per-step visual target.** Avg. accuracy (%) is reported.

Visual Target Strategy	In-D.	RAD	In-H.
(a) Shared (all steps $\rightarrow \mathbf{f}_{\text{ROI}}$)	81.08	74.28	80.50
(b) Per-step (interpolated α_k)	72.02	64.75	67.80

Analysis. The shared-target strategy consistently outperforms the per-step variant by a substantial margin: +9.06 on in-domain, +9.53 on VQA-RAD, and +12.70 on the held-out in-house testset. We attribute this to two factors. First, the per-step design imposes an artificial coarse-to-fine curriculum on the visual pathway, which conflicts with the reasoning trajectory that the model naturally learns under semantic supervision: each z_k reconstructs a distinct reasoning sentence whose content is not

necessarily ordered from global to local, so forcing the visual target to follow a fixed interpolation schedule creates gradient conflicts between the two supervision signals. Second, the shared target provides a consistent gradient direction at every step. Since $\mathcal{L}_{\text{text}}$ already differentiates the steps (each z_k must reconstruct a different reasoning sentence), the shared visual anchor serves as a stable spatial attractor that complements rather than competes with the per-step semantic signal. The modality collapse analysis (§C.7) further confirms that a shared target suffices to produce progressively differentiated latent states without any explicit step-wise visual curriculum.

C.7 Modality Collapse Analysis

To quantify whether latent states degenerate into a homogeneous representation across reasoning steps (i.e., modality collapse), we compute the pairwise cosine similarity between all latent states $\{z_1, \dots, z_K\}$ within each sample and average over the in-domain test set. Formally, for a given model variant, let x denote a test sample drawn from the in-domain test set $\mathcal{D}_{\text{test}}$, and let $z_i(x) \in \mathbb{R}^d$ be the i -th latent state produced by the recurrent loop on input x . We define the inter-step similarity matrix $\mathbf{S} \in \mathbb{R}^{K \times K}$ as:

$$S_{ij} = \mathbb{E}_{x \sim \mathcal{D}_{\text{test}}} \left[\frac{z_i(x)^T z_j(x)}{\|z_i(x)\| \cdot \|z_j(x)\|} \right]. \quad (9)$$

A collapsed model produces $S_{ij} \approx 1$ for all i, j , meaning that successive reasoning steps encode nearly identical information and the recurrent loop is effectively a no-op. A healthy model should exhibit lower off-diagonal values, with similarity decreasing as the step distance $|i - j|$ increases, indicating that each step contributes distinct, progressively refined information.

Figure 11 visualizes $\mathbf{S}$ for three variants: (a) Task-Only, (b) +Visual, and (c) VITAL. The Task-Only variant shows severe collapse: all off-diagonal entries exceed 0.97, confirming that without auxiliary supervision the latent loop degenerates into near-identical copies of the initial state. Adding visual supervision alone (b) partially alleviates collapse (S_{14} drops from 0.974 to 0.887), but the middle steps remain highly correlated ($S_{23} = 0.957$), suggesting that a single supervision signal cannot uniformly prevent stagnation across all steps. With full dual supervision (c), VITAL achieves substantially lower inter-step similarity: $S_{14} = 0.776$ and even adjacent steps

maintain meaningful differentiation ($S_{12} = 0.872$, $S_{34} = 0.903$). This confirms that the combination of semantic and visual supervision forces each latent state to encode genuinely distinct information, with the semantic branch ensuring progressive reasoning refinement and the visual branch anchoring each step to spatial evidence.

C.8 Quantitative Reliability of the Auxiliary Interpretability Modules

The preceding qualitative visualizations show that the auxiliary text decoder and visual projector provide complementary views of VITAL’s latent reasoning process. To quantify the reliability of these post-hoc interpretations, we re-attach both modules and evaluate them on the full in-domain test set, comprising 4,607 samples and 9,582 valid step-level latent–reasoning pairs. We evaluate whether the text decoder recovers the semantic content associated with each latent state and whether the visual projector localizes the corresponding image evidence.

Text decoder reconstruction quality. For every valid latent state z_k , we auto-regressively decode a textual explanation and compare it with the ground-truth reasoning sentence at step k . We report ROUGE-L to measure lexical and sequence-level overlap and BERTScore-F1 to measure semantic similarity.

Table 17: **Step-level reconstruction quality of the auxiliary text decoder.** Scores are computed between the decoded explanation from z_k and its ground-truth reasoning sentence. The average is taken over all 9,582 valid step-level pairs. All scores are percentages.

Step k	Pairs	ROUGE-L $\uparrow$	BERTScore-F1 $\uparrow$
1	4,607	54.86	86.19
2	2,880	50.93	84.32
3	1,127	47.82	84.11
4	968	46.65	83.76
Avg.	9,582	52.02	85.14

As shown in Table 17, the decoded explanations retain high semantic similarity at every reasoning depth: BERTScore-F1 remains above 83%, with an overall score of 85.14%. ROUGE-L remains above 46% at all steps and averages 52.02%, showing that the decoder also recovers substantial lexical and structural content. Both metrics decrease moderately with reasoning depth. Later steps generally synthesize more complex descriptions and spatial relationships, making exact reconstruction

more difficult; nevertheless, the relatively stable BERTScore-F1 indicates that semantic content is largely preserved even when the surface wording differs from the reference. These results quantitatively support the use of the decoder as a textual window into the information encoded by the latent states.

Visual projector localization quality. For each latent state z_k , we compute the cosine similarity between $VP(z_k)$ and every patch feature in MedSigLIP’s 32×32 feature grid, producing a patch-level activation map. We compare this map with the ground-truth segmentation mask using three complementary metrics: (i) *pointing accuracy*, which measures whether the maximally activated patch lies inside the ground-truth region; (ii) *patch IoU*, computed between the top 50% activated patches and the mask-aligned ground-truth patches; and (iii) *center distance*, the normalized Euclidean distance between the activation centroid and the ground-truth mask centroid, for which lower values are better. We compare VITAL with the Task-Only latent variant to determine whether the localization behavior arises from the auxiliary supervision.

Table 18: **Step-level localization accuracy of the visual projector.** Pointing accuracy and patch IoU are percentages; center distance is normalized by image size. The overall results are micro-averaged over all valid step-level samples rather than equally averaging the four step-wise scores, as the number of valid samples differs substantially across reasoning depths. Bold values indicate the better result at each step.

Step k	Method	Pointing Acc. $\uparrow$	Patch IoU $\uparrow$	Center Dist. $\downarrow$
1	Task-Only	31.2	14.8	0.2847
	VITAL	57.3	28.4	0.1842
2	Task-Only	33.5	15.6	0.2791
	VITAL	68.1	37.6	0.1253
3	Task-Only	30.8	14.3	0.2903
	VITAL	71.4	44.9	0.0874
4	Task-Only	32.1	15.1	0.2824
	VITAL	74.2	51.2	0.0651
Overall	Task-Only	32.3	15.2	0.2823
	VITAL	65.4	37.3	0.1334

Table 18 shows a clear separation between VITAL and Task-Only latent training. At the final reasoning step, VITAL achieves 74.2% pointing accuracy and 51.2% patch IoU, while reducing normalized center distance to 0.0651. Moreover, all three localization metrics improve monotonically from z_1 to z_4 : pointing accuracy increases by 16.9 points, patch IoU by 22.8 points, and center dis-




Figure 11: **Inter-step cosine similarity matrices averaged over the in-domain test set.** Lower off-diagonal values (darker cells) indicate more differentiated latent states across reasoning steps. (a) Task-Only exhibits near-total collapse; (c) VITAL maintains healthy inter-step diversity.

tance decreases by 0.1191. This progression quantitatively corroborates the coarse-to-fine spatial refinement observed in Figures 3, 12, 13, and 14.

In contrast, Task-Only produces weak and nearly unchanged localization across reasoning steps, with pointing accuracy remaining between 30.8% and 33.5% and patch IoU between 14.3% and 15.6%. The comparison demonstrates that task supervision alone does not yield the progressive grounding exhibited by VITAL, and that the auxiliary supervision is critical for aligning latent reasoning states with spatial evidence. Together with the text reconstruction results, these measurements provide quantitative evidence that the two re-attached modules recover complementary and meaningful information from VITAL’s latent states: semantic reasoning content through the text decoder and localized visual evidence through the visual projector.

C.9 More Case Studies

To further validate that VITAL’s latent states progressively converge toward the target region throughout the reasoning process, we visualize the visual projector output at every latent step for representative samples spanning five imaging modalities: CT, X-ray, pathology, dermoscopy, and ultrasound. For each sample ($K=4$), we project every latent state $z_1, \dots, z_4$ through the trained visual projector and compute patch-level cosine similarity with the resulting feature vector in MedSigLIP’s feature space, producing a per-step activation heatmap overlaid on the original image.

Figure 12 presents 8 cases arranged in a 4×10 grid (two cases per row). For each case, the first four images show the heatmap evolution from z_1 to z_4 , and the fifth image displays the ground-truth

segmentation mask (green overlay) for reference. Several consistent patterns emerge across modalities: (i) **progressive spatial refinement**: early latent states (z_1, z_2) produce broad, diffuse activations covering a coarse anatomical neighborhood, while later states (z_3, z_4) sharpen into tight, well-localized focus on the target structure, mirroring a coarse-to-fine reasoning process; (ii) **convergence to ground truth**: the final-step heatmap (z_4) closely aligns with the green GT mask across all modalities, confirming that the recurrent loop successfully drives latent states toward spatially precise visual grounding; (iii) **cross-modality robustness**: the progressive refinement pattern is consistent regardless of imaging modality, target size, or tissue contrast, demonstrating that dual supervision learns a modality-agnostic latent-to-visual mapping. These qualitative results complement the inter-step similarity analysis (§C.7) by showing that the decreasing cosine similarity between successive steps corresponds to meaningful spatial refinement rather than random drift.

We now present two detailed case studies that illustrate VITAL’s dual interpretability and compare its reasoning quality against representative baselines.

Case 1: CT liver analysis (Figure 13). The question asks the model to analyze the liver in an abdominal CT image. VITAL’s latent reasoning chain reveals a structured four-step process: z_1 identifies the liver’s large, solid parenchymal structure and its contrast against the surrounding intestinal tract and blood vessels; z_2 localizes it to the right part of the image extending into the middle; z_3 characterizes its capsule-like outline with smooth, gentle curves;




Figure 12: **Progressive visual projector activations across latent reasoning steps ($z_1 \rightarrow z_4$).** The figure shows 8 cases (two per row) spanning CT, X-ray, pathology, dermoscopy, and ultrasound. For each case, the first four images display the patch-level activation heatmap at each latent step, and the fifth image shows the ground-truth region (green overlay). Warmer colors indicate higher cosine similarity with the projected latent feature. Across all modalities, activations progressively sharpen from diffuse early-step responses to precise final-step localization that closely matches the ground truth.

z_4 assesses the tissue density as relatively uniform. The accompanying visual projector heatmaps progressively concentrate from a broad abdominal activation to a tight focus on the liver parenchyma, closely matching the GT ROI. The final answer correctly synthesizes all observations: the liver is in the right part of the image as a broad, solid organ with smooth edges and relatively uniform soft tissue attenuation.

In contrast, the baselines exhibit distinct failure modes. SIM-CoT suffers from *limited medical knowledge*: it incorrectly describes the liver as located “at the center of the scan” and provides only generic observations without precise spatial grounding. LVR produces outright *hallucinations*, fabricating findings of “heterogeneous attenuation,” “a small hypodense focal lesion in the inferior right hepatic lobe,” and “early fatty infiltration with a focal cystic change” that are entirely absent from the image. Claude-Opus-4.6 commits a *grounding error* by placing the liver in the “middle portion” of the image, and HuatuoGPT-Vision similarly mislocalizes it to the “upper-central portion.” These errors highlight the difficulty of accurate spatial reasoning without explicit visual grounding, which VITAL’s dual supervision directly addresses.

Case 2: Ultrasound breast tumor characterization (Figure 14). The question asks the model to reason about the location, shape, and visible characteristics of a breast tumor in an ultrasound image.

VITAL’s latent chain proceeds as follows: z_1 identifies a region of notably darker echotexture compared to the surrounding breast tissue, recognizing it as a distinct mass-like feature; z_2 characterizes the shape as roughly elliptical with width greater than height, noting the absence of spiky protrusions into adjacent tissues; z_3 assesses the boundary quality as smooth and clearly defined rather than irregular; z_4 localizes the tumor to the central upper part of the image, immediately beneath the superficial tissue layer. The visual projector heatmaps sharpen from a diffuse response across the breast parenchyma to a precise localization of the hypoechoic nodule, aligning with the GT ROI. The final answer accurately describes a small, hypoechoic oval mass in the upper central area with a clear boundary and uniform internal structure.

The baselines again demonstrate characteristic failures. SIM-CoT shows *limited medical knowledge* by misidentifying the tumor as a “fluid-filled pocket in the breast tissue rather than a solid mass,” confusing a solid hypoechoic nodule with a cyst. LVR commits a *grounding error*, placing the mass in “the lower right portion” when it is clearly in the upper center. Claude-Opus-4.6 *hallucinates* characteristics: it describes the nodule as located “in the center portion” with “mildly indistinct margins and a not uniform internal echotexture,” contradicting the clearly defined boundary and homogeneous interior visible in the image. HuatuoGPT-Vision similarly mislocalizes the tumor to “the lower right



Figure 13: **Case Study 1: CT liver analysis.** **Right:** VITAL’s latent reasoning chain ($z_1 \rightarrow z_4$) with decoded text and visual projector heatmaps showing progressive spatial refinement toward the liver. **Left:** Baseline outputs exhibiting limited medical knowledge (SIM-CoT), hallucination (LVR), and grounding errors (Claude-Opus-4.6, HuatuoGPT-V).

portion.” These cases demonstrate that VITAL’s visual-semantic dual supervision enables both spatially accurate grounding and clinically precise characterization that baselines consistently fail to achieve.

D Future Work

We identify several promising directions for extending VITAL. **(1) Adaptive reasoning depth.** The current design fixes $K=4$ for all samples. Introducing a learned early-exit mechanism (e.g., confidence-based halting or a termination token) would allow the model to allocate fewer steps to simple questions and more to complex clinical reasoning. **(2) Broader data coverage.** Our 61K dataset spans 9 modalities from segmentation annotations. Incorporating radiology reports as an additional supervision source could scale training data by orders of magnitude without requiring pixel-level masks for every sample. **(3) Reinforcement learning.** VITAL currently relies on teacher distillation, bounding its quality by the teacher’s reasoning ability. Outcome-based or process-based RL rewards could push latent reasoning beyond this ceiling, particularly for complex multi-step clinical scenarios.



Figure 14: **Case Study 2: Ultrasound breast tumor characterization.** **Right:** VITAL’s latent reasoning chain ($z_1 \rightarrow z_4$) progressively identifies echotexture contrast, elliptical shape, smooth boundaries, and precise localization. **Left:** Baselines exhibit limited medical knowledge (SIM-CoT), grounding errors (LVR, HuatuoGPT-V), and hallucination (Claude-Opus-4.6).