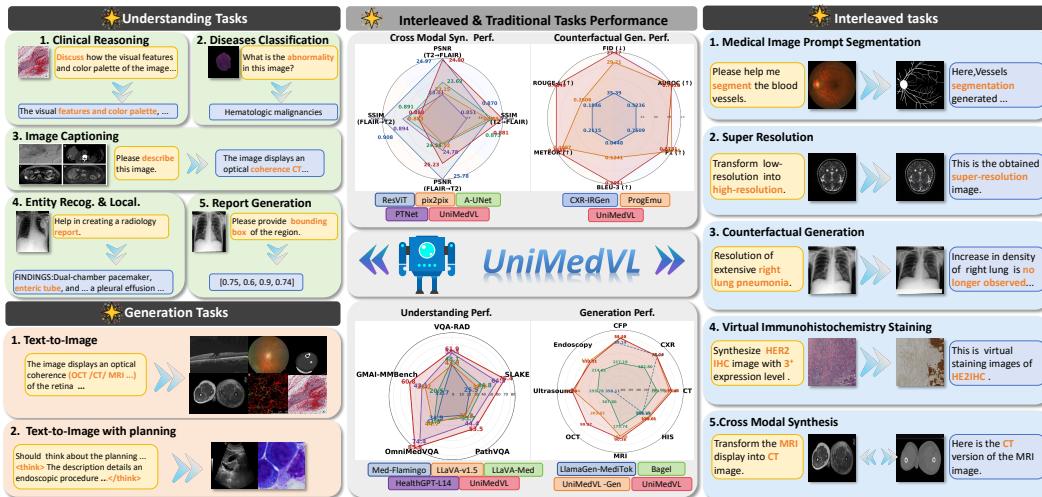


000 UNIMEDVL: UNIFYING MEDICAL MULTIMODAL 001 UNDERSTANDING AND GENERATION THROUGH 002 OBSERVATION-KNOWLEDGE-ANALYSIS 003 004

006 **Anonymous authors**
007 Paper under double-blind review

010 ABSTRACT

013 Clinical diagnosis demands models that can process multimodal medical inputs
014 (images, patient histories, lab results) and generate diverse outputs including both
015 textual reports and visual content (annotations, segmentation masks, and images).
016 Despite this need, existing medical AI systems disrupt this unified process: medical
017 image understanding models interpret images but cannot generate visual outputs,
018 while medical image generation models synthesize images but cannot provide
019 textual explanations. This leads to gaps in data representation, feature integration,
020 and task-level multimodal capabilities. To this end, we propose a multi-level frame-
021 work that mirrors clinical diagnosis through the Observation-Knowledge-Analysis
022 (OKA) paradigm. Specifically, at the observation level, we construct **UniMed-5M**,
023 a dataset comprising over 5.6M samples that reformat diverse unimodal data into
024 multimodal pairs for foundational observation. At the knowledge level, we pro-
025 pose **Progressive Curriculum Learning** that systematically introduce medical
026 multimodal knowledge. At the analysis level, we introduce **UniMedVL**, the first
027 medical unified multimodal model for the simultaneous analysis of image under-
028 standing and generation tasks within a single architecture. UniMedVL achieves
029 superior performance on five medical image understanding benchmarks, while
030 matching specialized models in generation quality across eight medical imaging
031 modalities. Crucially, our unified architecture enables bidirectional knowledge
032 sharing generation tasks enhance visual understanding features, demonstrating that
033 integrating traditionally separate capabilities within a single medical framework
034 unlocks improvements across diverse clinical scenarios. Code is available at [Link](#).
035



052 Figure 1: **Overview of UniMedVL unified framework.** Capabilities across medical image under-
053 standing and generation tasks and performance comparisons.

054
055
056
1 INTRODUCTION

057 Clinical diagnosis fundamentally follows a structured multi-level reasoning pipeline that is inherently
 058 multimodal in both inputs and outputs. Physicians systematically **observe** multimodal raw data
 059 (imaging patterns, patient histories, symptom descriptions (Huang et al., 2020)), integrate this with
 060 clinical **knowledge** (medical literature, domain expertise, cross-modal associations (Khader et al.,
 061 2023)), and **analyse** to produce diverse diagnostic outputs—textual reports explaining findings, visual
 062 annotations localizing abnormalities, and comparative imagery for treatment planning (Nguyen et al.,
 063 2023; Tanida et al., 2023).

064 Consider a radiologist examining suspected lung pathology: they process chest X-rays (visual), prior
 065 CT scans (cross-modal comparison), and patient history (textual) to generate multiple complementary
 066 outputs: detailed reports describing findings, visual annotations highlighting specific regions, and
 067 comparative visualizations for surgical planning. This exemplifies how clinical diagnosis requires
 068 unified processing of multimodal inputs to generate diverse multimodal outputs, where neither textual
 069 reports alone (lacking spatial localization) nor visual annotations alone (lacking reasoning context)
 070 suffice. While multimodal fusion has demonstrated substantial improvements in clinical decision-
 071 making (Benani et al., 2025; Soenksen et al., 2022), current medical AI remains fragmented, with
 072 state-of-the-art models achieving less than 60% accuracy compared to over 90% for human experts
 073 on diagnostic challenges (Kaczmarczyk et al., 2024). This fragmentation manifests at three critical
 074 levels: (i) **Data**: Medical datasets remain predominantly single-modal despite clear evidence that
 075 multimodal integration substantially improves diagnostic accuracy (Warner et al., 2024; Huang et al.,
 076 2023). (ii) **Features**: Current approaches lack systematic progressive training strategies that can
 077 effectively capture deep cross-modal relationships—most methods simply concatenate features rather
 078 than progressively building from basic pattern recognition to sophisticated multimodal reasoning (Haq
 079 et al., 2025). (iii) **Tasks**: While general-domain models have made progress in unified architectures,
 080 the medical domain still lacks truly unified models. For instance, although HealthGPT demonstrates
 081 both understanding and generation capabilities for medical tasks, it requires reloading different model
 082 checkpoints to switch between task types—a limitation that prevents seamless multi-task operation in
 083 clinical workflows (Lin et al., 2025).

084 To bridge this gap, we propose a workflow-guided framework that mirrors how physicians actually
 085 process medical information through the Observation-Knowledge-Analysis (OKA) paradigm: Ob-
 086 servation Level (Data): We construct UniMed-5M, a 5-million sample dataset that, unlike existing
 087 single-modal datasets, reformats medical data of various tasks into true multimodal input-output
 088 compatible pairs. Knowledge Integration Level (Features): We design Progressive Curriculum
 089 Learning that goes beyond naive concatenation. Through three carefully designed stages (alignment
 090 for medical data, fusion, and synthesis), our approach materialises models to discover cross-modal
 091 patterns better. Analysis Level (Tasks): We introduce UniMedVL, the first unified medical model
 092 capable of both understanding and generation within a single architecture at the same time. Our
 093 experiments validate two key insights: (1) Medical Representation learning requires framework, al-
 094 lowing the shared representations in medical domains necessitate the OKA framework with sufficient
 095 data quality and quantity, for unified medical multimodal learning; (2) Rapid adaptation is achievable,
 096 unified model architectures demonstrate the feasibility of quickly adapting to new medical tasks and
 097 datasets for scalable multimodal medical AI intelligence. Therefore, our contributions are as follows:
 098

- 099 • **Observation (Data-level):** We construct **UniMed-5M**, a large-scale dataset containing over 5.6M
 100 multimodal medical examples that reformat diverse unimodal datasets into uniform multimodal
 101 input-output pairs, and serve as the initial building blocks for unifying diverse medical tasks.
- 102 • **Knowledge integration (Feature-level):** We devise **Progressive Curriculum Learning**, a three-
 103 stage training paradigm that systematically builds medical multimodal capabilities: foundation
 104 training for basic pattern recognition, instruction tuning for cross-modal fusion, and unified multi-
 105 modal training for advanced synthesis.
- 106 • **Analysis (Task-level):** We introduce **UniMedVL**, a novel unified medical foundation model that
 107 provides multimodal capabilities—including understanding multimodal inputs and generating textual
 108 reports, image translation, segmentation masks, and synthetic medical images—within a single
 109 architecture without needing offline checkpoints once loaded.

108

2 RELATED WORK

109

110 2.1 MEDICAL MULTIMODAL LARGE LANGUAGE MODELS

111 Early medical MLLMs took an adapter-style approach, plugging medical vision encoders into general
 112 LLMs through lightweight projection or LoRA. Thawakar et al. (2024) aligned MedCLIP with Vicuna
 113 via a linear projector in XrayGPT. Li et al. (2023) bootstrapped instruction data from PubMed figures
 114 using GPT-4 in LLaVA-Med. These systems proved effective for VQA and reporting but kept fusion
 115 shallow and provided no unified route to image synthesis or editing. A second paradigm emphasizes
 116 data engineering. Chen et al. (2024b) leveraged GPT-4V to reformat noisy PubMed image–text pairs
 117 into the 1.3M-sample PubMedVision corpus in HuatuoGPT-Vision. While this approach mitigates
 118 data scarcity and noise, it remains primarily comprehension-oriented; unified, high-fidelity generation
 119 is still outside the model proper. Zhang et al. (2023a), which extends beyond radiology to molecules
 120 and proteins in a unified seq2seq architecture with BioMedGPT, and Singhal et al. (2025), which
 121 reaches expert-level exam performance via chain-of-thought and clinician-aligned instruction with
 122 Med-PaLM 2, strengthen biomedical reasoning but do not deliver a unified medical pipeline that
 123 natively spans image-level generation and text reasoning. More recently, Lin et al. (2025) emerged
 124 with HealthGPT as the first medical MLLM explicitly targeting unified multi-modal input and output.
 125 It introduces MOE LoRA to reduce task interference and to cover tasks. However, its unification
 126 relies on multiple task-specific models at inference time. As a result, different capabilities are not
 127 consolidated into a single, end-to-end model that uniformly expresses all tasks simultaneously.

128

129 2.2 UNIFIED MULTIMODAL UNDERSTANDING AND GENERATION MODELS

130 Outside the medical domain, unified multimodal research has developed along several paradigms.
 131 Autoregressive models (Team, 2024a; Wang et al., 2024; Lu et al.; 2024) treat images as discrete
 132 tokens within decoder-only Transformers, achieving architectural unity but limiting high-resolution
 133 synthesis due to long sequences and discrete reconstruction. Dual-encoder designs (Wu et al., 2025c;
 134 Ma et al., 2025d; Xu et al., 2025b) address the granularity conflict between semantic understanding
 135 and pixel-level generation through separate visual pathways, improving task-specific performance at
 136 increased inference cost. Hybrid objectives combine different generative paradigms: Zhou et al. (2024)
 137 jointly optimize text cross-entropy and image diffusion losses in Transfusion, Xie et al. (2024) embed
 138 discrete diffusion via Omni-Attention for faster synthesis in SHOW-O, while modular approaches (Wu
 139 et al., 2025e; ?; 2024a) bridge frozen MLLMs with diffusion models through learnable connectors.
 140 These solutions achieve cost-effectiveness but sacrifice end-to-end differentiability. Representation
 141 innovations target the semantics-fidelity gap through various strategies: multi-codebook quantization
 142 (Ma et al., 2025c), contrastive-aligned tokenization (Wu et al., 2024b), unified CLIP semantic spaces
 143 (Chen et al., 2025), and masked autoregressive encoders (Jiang et al., 2024). Advanced autoregressive
 144 methods (Liao et al., 2025; Zhang et al., 2025; Zhuang et al., 2025) enable high-fidelity interleaved
 145 generation through deep fusion, pre-filled tokens, and reinforcement learning from human feedback.
 146 Despite these advances, balancing semantic understanding with pixel-level reconstruction remains
 147 challenging, particularly for fine-grained medical localization and diagnostic-quality image synthesis
 required in clinical applications.

148

149 3 METHODOLOGY

150 Our workflow-guided multi-level framework systematically implements the clinical Observation-
 151 Knowledge-Analysis (OKA) paradigm through three corresponding stages: data-level observation
 152 for comprehensive multimodal dataset construction, feature-level knowledge integration through
 153 principled curriculum learning, and task-level analysis via unified model architecture. Each stage
 154 addresses specific computational challenges while maintaining clinical workflow coherence.

155

156 3.1 OBSERVATION LEVEL: UNIMED-5M DATASET CONSTRUCTION

157 At the observation level, comprehensive multimodal datasets are constructed to enable systematic
 158 processing of diverse medical inputs that mirror clinical practice. The dataset construction follows
 159 clinical workflow patterns where multiple data modalities are observed and initially processed before
 160 knowledge integration. The overall dataset curation pipeline is shown in Fig. 2.

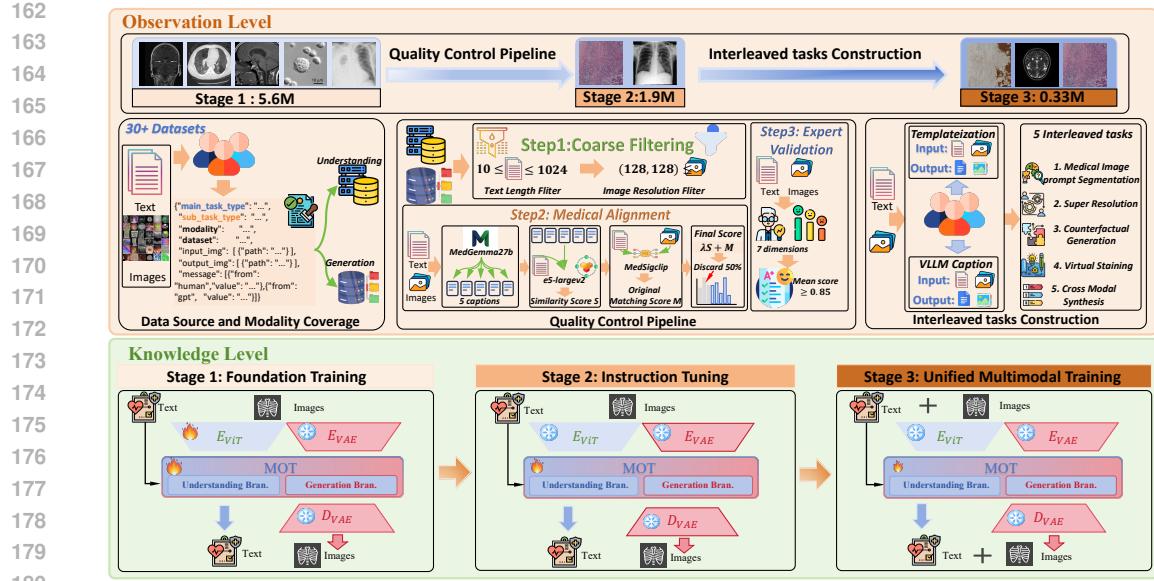


Figure 2: Overview of the proposed **Observation–Knowledge** framework. **Observation:** Covers data sources and modality coverage, quality control pipeline, and interleaved task construction for building training data across different model stages. **Knowledge:** Refers to the progressive curriculum training paradigm, consisting of three stages that gradually equip the model with generalized capabilities on interleaved tasks.

Data Source and Modality Coverage. A comprehensive medical dataset comprising 5.6M samples is assembled from diverse public repositories including PMC-OA, Quilt-1M, PubMedVision, GMAI-VL datasets, CheXpertPlus, PMC-VQA, Medical-Diff-VQA, SLAKE, PathVQA, and other specialized medical datasets through systematic data synthesis and augmentation methodologies detailed in Appendix A.2. The collection encompasses nine primary imaging modalities: chest X-rays (CXR), histopathology images (HIS), CT scans, MRI sequences, retinal fundus photography (CFP), optical coherence tomography (OCT), endoscopy, ultrasound, and fluorescence microscopy (FM). The dataset encompasses diverse medical AI task categories spanning understanding, generation, and multimodal input-output capabilities.

Quality Control Pipeline. We adopt a three-step pipeline that progressively increases fidelity while controlling cost:

- **Coarse Filtering.** Images are preprocessed through modality-specific normalization and resolution filtering ($\geq 128 \times 128$ pixels). Text undergoes specialized tokenization that preserves medical terminology, followed by length filtering (16–1024 characters).
- **Medical Alignment.** Because medical captions often emphasize specific pathological findings rather than exhaustive descriptions, we implement a dedicated verification pipeline. MedGemma-27b (Sellergren et al., 2025) generates five diverse captions per image; semantic similarity is computed with E5-large-v2 embeddings (Wang et al., 2022); and medical-specific alignment is assessed using MedSigCLIP (Sellergren et al., 2025). We then compute a combined alignment score $\text{score}_{\text{final}} = \lambda \cdot \text{similarity}_{\text{E5}} + \text{score}_{\text{MedSigCLIP}}$ with $\lambda = 0.5$, retaining the top 50% of pairs as high-quality training data.
- **Expert Validation.** Medical experts draw a 5% stratified subset for comprehensive auditing strictly along the seven dimensions: modality–image match, factual accuracy, information completeness, position/quantity accuracy, professionalism, planning coherence, and clinical reasoning quality. This stage serves as quality assurance rather than additional filtering: on the audited subset, expert ratings across all dimensions are observed to be ≥ 0.85 (with high inter-rater agreement), substantiating the dataset’s clinical reliability and mitigating potential LLM hallucinations.

Interleaved tasks Construction. This component encompasses five interleaved tasks: medical image prompt segmentation, super-resolution, counterfactual generation, virtual staining, and cross-modal

synthesis. We adopt two complementary construction strategies: Templateization and VLLM Caption. In Templateization, inputs and outputs are standardized into structured image–text pairs, where textual prompts explicitly guide the model beyond the provided image and outputs follow a templated format. In contrast, VLLM Captioning emphasizes generating semantically rich textual descriptions that interpret the corresponding images in medical contexts, including symptom analysis and clinical insights.

3.2 KNOWLEDGE LEVEL: PROGRESSIVE CURRICULUM LEARNING

At the knowledge integration level, deep cross-modal knowledge fusion is achieved through a principled curriculum learning paradigm that progressively builds from basic medical pattern recognition to sophisticated multimodal reasoning capabilities.

Progressive Curriculum Training Paradigm:

- **Stage 1: Foundation Training.** Foundational medical domain awareness is established through unsupervised exposure to comprehensive medical datasets. The foundation training stage prioritizes broad pattern recognition over task-specific performance, enabling robust medical concept acquisition through text-image paired learning and next-token prediction across diverse medical sources. Furthermore, the training emphasizes learning general medical visual-language alignments without task-specific constraints and overly curated datasets.
- **Stage 2: Instruction Tuning.** Medical expertise is systematically developed through fine-tuning on curated high-quality instruction data. The instruction-formatted medical tasks follow the format $(q, x_v, k) \rightarrow (a_t, a_v)$ where query q , visual input x_v , and knowledge context k generate textual a_t and visual a_v responses. We implement differentiated enhancement strategies for distinct task types: For medical understanding tasks such as VQA, we augment standard responses with existing Distilled Chain of Thought (DCOT) data that explicitly articulate the reasoning pathway from visual observation to medical conclusions. For generation tasks, we employ Caption Augmented Generation (CAG) pipeline to enhance caption quality, incorporating structured planning steps that guide the visual synthesis process. The details are provided in Appendix A.5.
- **Stage 3: Unified Multimodal Training.** Multimodal capabilities of generation and understanding are developed through sophisticated tasks requiring integrated visual-textual reasoning. This stage focuses on complex interleaved tasks that combine understanding and generation requirements within unified sequences. The training strategy maintains semantic stability from previous stages while enabling advanced synthesis capabilities through selective parameter optimization, preparing the model for comprehensive clinical analysis workflows.

3.3 ANALYSIS LEVEL: UNIMEDVL UNIFIED ARCHITECTURE

At the analysis level, comprehensive multimodal medical outputs are generated through a unified architecture that naturally simulates clinical analytical processes while maintaining cross-modal consistency. The UniMedVL architecture integrates the progressive curriculum learning paradigm into a cohesive system capable of both understanding and generation within a single model.

Task Organization. Model training is systematically organised into three primary tasks that reflect fundamental capabilities required for unified medical multimodal systems: (i) Understanding tasks encompassing medical image comprehension, VQA, diagnostic reasoning, image captioning, and clinical report generation; (ii) Generation tasks focusing on text-to-image synthesis with conditional medical image generation and planning-guided approaches; and (iii) Interleaved tasks combining visual-textual inputs and outputs requiring seamless multimodal integration. These interleaved tasks include sophisticated capabilities such as virtual immunohistochemistry staining, cross-modal synthesis (e.g., CT to MRI synthesis), counterfactual generation for treatment planning and development forecasting.

Model Architecture Overview. Following [Deng et al. \(2025\)](#), we adopt a unified architecture with dual visual encoders and mixture-of-transformer-experts (MoT). The understanding-oriented encoder E_{ViT} extracts semantic tokens $z_{ViT} = E_{ViT}(x_v)$ for multimodal comprehension tasks, while the generation-oriented encoder E_{VAE} produces latent representations $z_{VAE} = E_{VAE}(x_v)$ for visual synthesis tasks. The MoT module contains specialized transformer experts: an understanding expert

processes interleaved sequences of text and ViT tokens $[x_{text}, z_{ViT}]$ for vision-language understanding, while a generation expert handles sequences containing text and VAE tokens $[x_{text}, z_{VAE}]$ for image generation. Projection layers f_{ViT} and f_{VAE} bridge the visual encoders with the transformer experts, mapping encoded features to the shared hidden dimension. For generation outputs, the decoder D_{VAE} reconstructs visual content from the latent representations back to pixel space. A generalized causal attention mechanism enables both experts to operate on the same token sequence through shared self-attention operations for different types of tasks, integrating understanding and generation paradigms into a unified framework.

Training Objectives. The model is trained with a unified loss function combining understanding task loss and generation task loss. For understanding tasks, we employ next-token prediction:

$$\mathcal{L}_{NTP} = - \sum_{i=1}^n \log p(t_{i+1} | t_{\leq i}, z_{ViT}; \theta), \quad (1)$$

where t_i denotes the i -th text token and θ represents model parameters. For visual generation, rectified flow matching is applied on VAE latents:

$$\mathcal{L}_{flow} = \mathbb{E}_{t, \epsilon} [\|v_\theta(\tilde{z}_{VAE}, t, c) - v\|^2], \quad (2)$$

where v_θ is the velocity prediction network, \tilde{z}_{VAE} represents noisy latents processed by D_{VAE} , t is the time step, c denotes conditioning, and v is the target velocity field. The overall training loss is as follows:

$$\mathcal{L} = \mathcal{L}_{NTP}(z_{ViT}) + \alpha \cdot \mathcal{L}_{flow}(z_{VAE}), \quad (3)$$

where the coefficient α balances the contribution of generation tasks. We set $\alpha = 4$ empirically.

4 EXPERIMENTS

4.1 BENCHMARKS AND BASELINES

Evaluation Benchmarks. We evaluate UniMedVL across medical visual understanding and generation benchmarks. For **image understanding tasks**, we employ VQA-RAD (Lau et al., 2018), SLAKE (Liu et al., 2021), PathVQA (He et al., 2020), OmniMedVQA (Hu et al., 2024), and GMAI-MMBench (Ye et al., 2024), which cover diverse medical scenarios. For **image generation tasks**, we split the image–caption pairs in the proposed dataset into 80% for training and 20% for testing. We use the test set to evaluate UniMedVL’s text-to-image generation performance. For **interleaved tasks**, we utilize the BCI dataset (Liu et al., 2022b) for the virtual immunohistochemistry staining task. The IXI dataset (?) is leveraged to evaluate the super-resolution task, and the BraTS 2023 dataset (Adewole et al., 2023) is used for evaluating the cross-modal synthesis task. We use the ICG-CXR dataset (Ma et al., 2025b) to evaluate the counterfactual generation task.

Baseline Methods. These include two categories of methods: **specialized models** and **unified multimodal models**. For specialized models, we include understanding-only models such as Med-Flamingo (Moor et al., 2023), LLaVA-Med (Li et al., 2023), HuatuoGPT-Vision (Chen et al., 2024b), RadFM (Wu et al., 2025b), GMAI-VL (Li et al., 2024), LLaVA-v1.5 (Liu et al., 2024), and InternVL2 (Team, 2024b). We also compare with image translation models including CycleGAN (Zhu et al., 2017), pix2pix (Isola et al., 2017), pix2pixHD (Wang et al., 2018), pyramid pix2pix (Liu et al., 2022b), SRCNN (Dong et al., 2015), VDSR (Kim et al., 2016), SwinIR (Liang et al., 2021), Restormer (Zamir et al., 2022), AMIR (Yang et al., 2024), ResViT (Dalmaz et al., 2022), and TransUNet (Chen et al., 2021). Additionally, to determine the model performance of medical imaging generation capability, we finetuned LlamaGen-MediTok (Ma et al., 2025a). For unified multimodal models, we include general frameworks like Janus (Wu et al., 2025d) and Bagel (Deng et al., 2025), as well as medical unified models such as HealthGPT (Lin et al., 2025).

Evaluation Metrics. We employ task-specific metrics aligned with clinical relevance. For **image understanding tasks**, we utilize accuracy as the evaluation metric. For **image generation tasks**, we employ generation FID (gFID) and BioMedCLIP score to evaluate the quality of synthesized images. For **interleaved tasks**, we leverage PSNR and SSIM as evaluation metrics for virtual immunohistochemistry staining, super-resolution, and cross-modal synthesis tasks. For counterfactual generation, we follow the experimental setup of ProgEmu (Ma et al., 2025b), using gFID, AUC-ROC, and F1 to evaluate the quality of synthesized images, and BLEU-3, METEOR, and ROUGE-L to assess the quality of the explanatory text.

324
 325 **Table 1: Comparison of UniMedVL with other LVLMs and unified multi-modal models on**
 326 **medical visual understanding tasks. Bold and underlined text indicate the best performance and**
 327 **second-best performance, respectively.**

328 Model	328 Params	328 Medical	328 VQA-RAD	328 SLAKE	328 PathVQA	328 OmniMedVQA	328 GMAI-MMBench
Understanding Only							
330 LLaVA-v1.5	330 7B	330 \times	330 42.8	330 37.7	330 31.4	330 44.7	330 38.23
331 InternVL2	331 8B	331 \times	331 49.0	331 50.1	331 31.9	331 54.5	331 43.47
332 Med-Flamingo	332 8.3B	332 \checkmark	332 43.0	332 25.5	332 31.3	332 34.9	332 12.74
333 LLaVA-Med	333 7B	333 \checkmark	333 48.1	333 44.8	333 35.7	333 41.3	333 20.54
334 RadFM	334 14B	334 \checkmark	334 50.6	334 34.6	334 14.33	334 23.5	334 22.34
335 HuatuoGPT-Vision-7B	335 7B	335 \checkmark	335 53.0	335 49.1	335 32.0	335 50.0	335 50.22
336 MedGemma-4B	336 4B	336 \checkmark	336 67.6	336 71.2	336 33.7	336 68.4	336 44.0
337 Lingshu-7B	337 7B	337 \checkmark	337 62.7	337 77.0	337 59.6	337 82.0	337 52.3
338 Lingshu-32B	338 32B	338 \checkmark	338 71.4	338 84.7	338 61.3	338 80.4	338 52.7
339 GMAI-VL	339 7B	339 \checkmark	339 66.3	339 72.9	339 39.8	339 88.5	339 61.74
Unified Understanding and Generation							
340 Janus	340 1.3B	340 \times	340 52.8	340 26.9	340 27.9	340 45.7	340 39.30
341 Bagel	341 7B	341 \times	341 60.09	341 58.91	341 39.05	341 71.13	341 48.11
342 HealthGPT-M3	342 3.8B	342 \checkmark	342 55.9	342 56.4	342 39.7	342 68.5	342 42.08
343 HealthGPT-L14	343 14B	343 \checkmark	343 58.3	343 64.5	343 44.4	343 74.4	343 43.1
344 UniMedVL (Ours)	344 14B	344 \checkmark	344 61.9	344 75.4	344 53.5	344 85.8	344 60.75

344 Table 2: Performance comparison of our UniMedVL variants and other baseline models on the
 345 text-driven image generation task across different modalities. CS denotes BioMedCLIP Score. **Bold**
 346 and underlined text indicate the best performance and second-best performance, respectively.

347 Method	347 CFP		347 CXR		347 CT		347 HIS		347 MRI		347 OCT		347 Ultrasound		347 Endoscopy		347 Average	
	347 FID \downarrow	347 CS \uparrow																
349 LlamaGen-MediTok	89.14	-	68.16	-	-	-	198.63	-	-	-	-	-	358.11	-	-	-	171.85	-
350 Bagel	217.19	0.650	182.80	0.662	163.78	0.652	206.18	0.643	175.74	0.639	307.80	0.719	255.78	0.672	214.61	0.668	215.49	0.660
351 UniMedVL-Gen	<u>77.35</u>	0.699	190.38	0.672	<u>79.84</u>	0.694	107.20	0.699	82.99	0.699	107.06	0.721	100.44	0.700	121.89	0.704	108.40	0.699
352 UniMedVL	53.20	0.708	<u>73.04</u>	0.702	73.04	0.696	<u>149.01</u>	0.704	<u>90.36</u>	0.706	99.27	0.721	95.38	0.706	<u>133.11</u>	0.707	96.29	0.706

353 4.2 PERFORMANCE OF UNIMEDVL

354 **Medical Visual Understanding Performance.** We evaluate the understanding capabilities of
 355 UniMedVL across diverse medical VQA and image comprehension benchmarks. Table 1 presents
 356 comprehensive results comparing our model with existing medical VLLMs and unified multimodal
 357 models. Table 1 reveals a critical insight: unified architectures can achieve understanding perfor-
 358 mance comparable to specialized models without sacrificing generation capabilities. UniMedVL
 359 demonstrates this principle by maintaining competitive performance across diverse medical domains
 360 while supporting both understanding and generation within a single architecture. The key technical
 361 insight emerges from comparing unified models (bottom section) with understanding-only models
 362 (top section): HealthGPT requires separate model checkpoints for different tasks, while UniMedVL
 363 achieves superior performance (85.8% on OmniMedVQA vs. HealthGPT-L14’s 74.4%) with seam-
 364 less task switching. This validates our core contribution that progressive curriculum learning enables
 365 effective knowledge sharing between understanding and generation pathways in medical contexts.

366 **Medical Image Generation Performance.** Table 2 provides empirical evidence for cross-modal
 367 knowledge transfer in medical generation. A critical insight emerges from comparing UniMedVL-
 368 Gen (generation-only training) with UniMedVL: the average gFID improvement demonstrates that
 369 understanding tasks contribute semantic constraints to enhance generation quality. Furthermore,
 370 UniMedVL achieves BioMedCLIP scores of 0.706 on average across modalities, indicating strong
 371 semantic alignment between generated images and clinical text descriptions. This challenges the
 372 conventional assumption that joint training compromises individual task performance, instead showing
 373 that medical multimodal learning benefits from task synergy when the data bottleneck is relieved.

374 **Interleaved Multimodal Tasks Performance.** Table 3 demonstrates UniMedVL’s performance on
 375 interleaved multimodal tasks. For virtual immunohistochemistry staining (H&E \rightarrow IHC), UniMedVL
 376 achieves 20.27 PSNR, outperforming HealthGPT-M3 by 28%. In MRI super-resolution (4 \times), our
 377 model attains 27.29 PSNR/0.890 SSIM, showing substantial improvement over the unified baseline.
 For cross-modal synthesis (T2 \leftrightarrow FLAIR), UniMedVL reaches 25.07 average PSNR, approaching

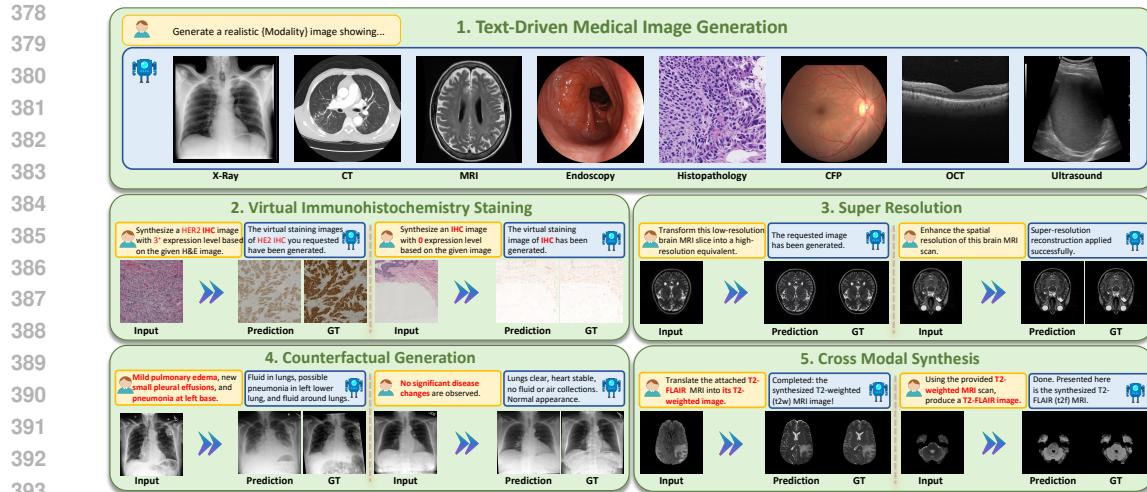


Figure 3: **Comprehensive visualization of UniMedVL’s multimodal capabilities.** Demonstration of diverse medical imaging tasks, including text-to-image generation, virtual staining, super resolution, counterfactual generation, and cross-modal synthesis.

specialized models while maintaining unified capabilities. Figure 3 provides qualitative visualization of these capabilities. The key insight emerges from comparing UniMedVL[†] (mixed data training without progressive stages) with full UniMedVL: consistent improvements across all tasks (e.g., 2.16 PSNR gain in virtual staining) demonstrate that our complete progressive training paradigm effectively learns cross-modal relationships that simple fine-tuning cannot capture.

Table 3: **Performance Comparison on specialised generation tasks.** histological staining transformation (H&E to IHC), MRI super-resolution (4 \times), and medical image translation ($T_2 \leftrightarrow FLAIR$). PSNR and SSIM are used in medical image translation. [†] indicates unified fine-tuning variant. **Bold** and underlined text indicate the best performance and second-best performance, respectively.

H&E \rightarrow IHC Staining		MRI Super-Resolution		Medical Image Translation		
Method	PSNR/SSIM	Method	PSNR/SSIM	Method	$T_2 \rightarrow FLAIR$	FLAIR $\rightarrow T_2$
CycleGAN	16.20/0.373	SRCCNN	28.81/0.892	ResViT	24.97 /0.870	25.78 / <u>0.908</u>
Pix2Pix	18.65/0.419	VDSR	30.04/0.914	pGAN	24.01/0.864	25.09/0.894
Pix2PixHD	19.63/0.471	SwinIR	31.55/0.933	pix2pix	23.15/0.869	24.52/0.883
Pyramid Pix2pix	21.16 / <u>0.477</u>	Restormer	<u>31.85</u> /0.938	A-UNet	23.69/0.873	23.84/0.876
		AMIR	31.99 / <u>0.939</u>	SAGAN	24.02/0.860	24.56/0.891
HealthGPT-M3	15.81/0.242	HealthGPT-M3	18.37/0.580	HealthGPT-M3	18.88/0.745	19.30/0.750
UniMedVL [†]	18.11/0.401	UniMedVL [†]	19.64/0.602	UniMedVL [†]	23.99/0.711	23.49/0.732
UniMedVL	<u>20.27</u> /0.456	UniMedVL	27.29/0.890	UniMedVL	24.90 / <u>0.881</u>	<u>25.23</u> /0.883
						25.07 / <u>0.882</u>

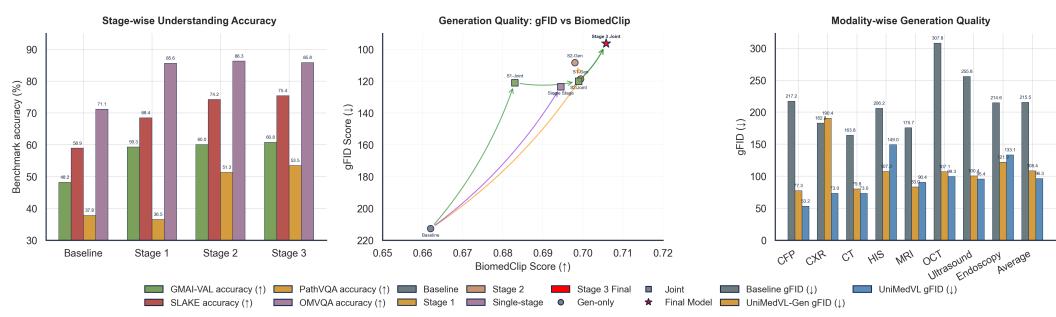
Table 4 evaluates counterfactual generation capabilities with explanatory text. UniMedVL[†] achieves 27.17 gFID and significantly higher text quality metrics (0.2641 BLEU-3, 0.4486 METEOR, 0.4649 ROUGE-L) compared to specialized baselines. The improved counterfactual check rate (0.797 AUROC) demonstrates that our unified training enables generation of clinically plausible scenarios with coherent textual explanations.

Table 4: Comparison of UniMedVL with baseline methods on medical counterfactual generation. **Bold** and underlined text indicate the best performance and second-best performance, respectively.

Method	Counterfactual Image			Explanatory Text		
	gFID \downarrow	AUROC \uparrow	F1 \uparrow	BLEU-3 \uparrow	METEOR \uparrow	ROUGE-L \uparrow
CXR-IRGen	35.39	0.5236	0.7609	0.0448	0.2115	0.1846
ProgEmu	<u>29.21</u>	<u>0.7921</u>	0.8914	<u>0.1241</u>	<u>0.4097</u>	<u>0.2606</u>
UniMedVL[†]	27.17	0.7970	<u>0.8731</u>	0.2641	0.4486	0.4649

432 **Table 5: Ablation study of the proposed progressive curriculum learning strategy.** UVE refers
 433 to the understanding-oriented vision encoder. G and U refer to the generation and understanding
 434 subsets of UniMed-5M, respectively. CAG: Caption Augmented Generation, DCOT: Distilled Chain
 435 of Thought. **Bold** indicates the best performance and underlined indicates second-best performance.
 436

437 Model	UVE	\mathcal{L}_{NTP}	$\mathcal{L}_{\text{flow}}$	Data Type	Understanding				Generation	
					GMAI-MMBench	SLAKE	PathVQA	OMVQA	<u>gFID</u> ↓	<u>BioMedCLIP</u> ↑
Baseline Comparison										
One-Stage-Joint-Base	✗	✓	✓	U+G	0.5354	0.6560	0.4946	0.7784	123.48	0.6945
Stage 1: Foundation Training										
F-Baseline	✗	✗	✗	-	0.481	0.589	0.390	0.7113	212.73	0.662
C-G-only	✗	✗	✓	G	-	-	-	-	118.5991	<u>0.6994</u>
B-U-only	✓	✓	✗	U	0.505	0.5476	0.3673	0.7723	-	-
H-Joint-Base	✓	✓	✓	U+G	0.593	0.6843	0.3649	0.8562	121.02	0.683
Stage 2: Instruction Tuning										
C-G-only	✗	✗	✓	CAG	-	-	-	-	108.40	0.698
B-U-only	✓	✓	✗	DCOT	0.5432	0.6032	0.4526	0.8167	-	-
H-Joint-Base	✓	✓	✓	High-quality U+G	<u>0.6004</u>	<u>0.7418</u>	<u>0.5130</u>	0.8626	120.036	0.6989
Stage 3: Unified Multimodal Training										
H-Joint-Base	✓	✓	✓	Interleaved tasks	0.6075	0.7540	0.5346	0.8584	96.287	0.7058



450 **Figure 4: Visual Comparison of Performance across different training stages and modalities.**
 451 **(Left:)** Stage-wise understanding accuracy performance. **(Center:)** Generation quality evolution
 452 with gFID reduction and BioMedCLIP score enhancement through different training stages. **(Right:)**
 453 Modality-specific generation performance comparison across UniMedVL variants.
 454

465 4.3 ABLATION STUDY

466 Table 5 and Figure 4 validate how our Knowledge level and Analysis level synergistically improve
 467 model capabilities. The critical finding is that joint training (H-Joint-Base) consistently outperforms
 468 single-task variants during the stage 1, indicating that UniMedVL learns fundamental unified multi-
 469 modal representations to effectively perform both understanding and generation tasks. The stage 2
 470 further improves the performance on both understanding and generation tasks. The understanding
 471 and generation capabilities are enhanced by the instructions with a reasoning process and high-quality
 472 image captions, respectively. The stage 3 brings significant improvement, demonstrating that the
 473 unified multimodal representations are further improved at this stage and support both understanding
 474 and generation tasks.
 475

476 5 CONCLUSION

477 We presented UniMedVL, a unified framework that simultaneously performs medical image under-
 478 standing and generation within a single model, validated through extensive experiments on over 5
 479 million medical samples demonstrating both state-of-the-art comprehension (75.40% SLAKE) and
 480 competitive generation quality (95.80 gFID). While our current work focuses on 2D medical imaging,
 481 the proposed OKA paradigm establishes foundations for exploring diverse medical AI tasks beyond
 482 those demonstrated, including 3D volumetric analysis, temporal reasoning, and multimodal clinical
 483 decision support. This work represents a critical step toward truly integrated medical AI systems
 484 where understanding and generation capabilities synergistically enhance clinical workflows.
 485

486 ETHICS STATEMENT
487

488 This work does not involve experiments on human subjects, patient interventions, or the collection of
489 private medical data. All datasets used in this study are publicly available, including CheXpertPlus,
490 SLAKE, PathVQA, OmniMedVQA, IXI, BraTS 2023, and other open-access medical datasets. Data
491 sources were used in accordance with their respective licenses, and all materials were de-identified
492 prior to use. The purpose of this research is to advance the scientific understanding of unified
493 multimodal modelling for healthcare data rather than to deploy clinical decision-support systems.
494 No patient-level decisions or clinical predictions were made based on model outputs. Expert audits
495 were limited to quality control of publicly available samples and did not involve identifiable patient
496 information.

497
498 REPRODUCIBILITY STATEMENT
499

500 To ensure reproducibility, all implementation details, model configurations, and training hyperpa-
501 rameters are provided in the Appendix. The full code and configuration files are available in an
502 anonymous repository. The UniMed-5M dataset construction process, including data sources, quality
503 control criteria, and interleaved task synthesis, is fully documented in the Appendix and illustrated
504 in Figure 2. All benchmarks used for evaluation, VQA-RAD, SLAKE, PathVQA, OmniMedVQA,
505 BraTS 2023, and IXI, are publicly available or derived from publicly available sources.

506 REFERENCES
507

509 Maruf Adewole, Jeffrey D Rudie, Anu Gbdamosi, Oluyemisi Toyobo, Confidence Raymond, Dong
510 Zhang, Olubukola Omidiji, Rachel Akinola, Mohammad Abba Suwaid, Adaobi Emegoakor, et al.
511 The brain tumor segmentation (brats) challenge 2023: Glioma segmentation in sub-saharan africa
512 patient population (brats-africa). [ArXiv](#), pp. arXiv-2305, 2023.

513 Ivo M Baltruschat, Parvaneh Janbakhshi, and Matthias Lenga. Brasyn 2023 challenge: Miss-
514 ing mri synthesis and the effect of different learning objectives. In [International Challenge on](#)
515 [Cross-Modality Domain Adaptation for Medical Image Segmentation](#), pp. 58–68. Springer, 2023.

516 Alaeddine Benani, Stéphane Ohayon, Fewa Laleye, Pierre Bauvin, Emmanuel Messas, Sylvain Bodard,
517 and Xavier Tannier. Is multimodal better? a systematic review of multimodal versus unimodal
518 machine learning in clinical decision-making. [medRxiv](#), pp. 2025–03, 2025.

520 Black Forest Labs. Flux, 2024. URL <https://github.com/black-forest-labs/flux>.
521 GitHub repository.

522 Pierre Chambon, Jean-Benoit Delbrouck, Thomas Sounack, Shih-Cheng Huang, Zhihong Chen, Maya
523 Varma, Steven QH Truong, Chu The Chuong, and Curtis P Langlotz. Chexpert plus: Augmenting
524 a large chest x-ray dataset with text radiology reports, patient demographics and additional image
525 formats. [arXiv preprint arXiv:2405.19538](#), 2024.

527 Jieneng Chen, Yongyi Lu, Qihang Yu, Xiangde Luo, Ehsan Adeli, Yan Wang, Le Lu, Alan L Yuille,
528 and Yuyin Zhou. Transunet: Transformers make strong encoders for medical image segmentation.
529 [arXiv preprint arXiv:2102.04306](#), 2021.

530 Jiuhai Chen, Zhiyang Xu, Xichen Pan, Yushi Hu, Can Qin, Tom Goldstein, Lifu Huang, Tianyi
531 Zhou, Saining Xie, Silvio Savarese, et al. Blip3-o: A family of fully open unified multimodal
532 models-architecture, training and dataset. [arXiv preprint arXiv:2505.09568](#), 2025.

533 Junying Chen, Chi Gui, Ruyi Ouyang, Anningzhe Gao, Shunian Chen, Guiming Chen, Xidong
534 Wang, Zhenyang Cai, Ke Ji, Xiang Wan, et al. Towards injecting medical visual knowledge into
535 multimodal llms at scale. In [Proceedings of the 2024 Conference on Empirical Methods in Natural](#)
536 [Language Processing](#), pp. 7346–7370, 2024a.

538 Junying Chen, Chi Gui, Ruyi Ouyang, Anningzhe Gao, Shunian Chen, Guiming Hardy Chen, Xidong
539 Wang, Ruifei Zhang, Zhenyang Cai, Ke Ji, et al. Huatuogpt-vision, towards injecting medical
visual knowledge into multimodal llms at scale. [arXiv preprint arXiv:2406.19280](#), 2024b.

540 Zhe Chen, Jiannan Wu, Wenhai Wang, Weijie Su, Guo Chen, Sen Xing, Muyan Zhong, Qinglong
 541 Zhang, Xizhou Zhu, Lewei Lu, et al. Internvl: Scaling up vision foundation models and aligning for
 542 generic visual-linguistic tasks. In *Proceedings of the IEEE/CVF conference on computer vision*
 543 and pattern recognition

544 pp. 24185–24198, 2024c.

545 Onat Dalmaz, Mahmut Yurt, and Tolga Çukur. Resvit: Residual vision transformers for multimodal
 546 medical image synthesis. *IEEE Transactions on Medical Imaging*, 41(10):2598–2614, 2022.

547 Chaorui Deng, Deyao Zhu, Kunchang Li, Chenhui Gou, Feng Li, Zeyu Wang, Shu Zhong, Weihao
 548 Yu, Xiaonan Nie, Ziang Song, et al. Emerging properties in unified multimodal pretraining. *arXiv*
 549 preprint arXiv:2505.14683, 2025.

550 Chao Dong, Chen Change Loy, Kaiming He, and Xiaoou Tang. Image super-resolution using deep
 551 convolutional networks. *IEEE transactions on pattern analysis and machine intelligence*, 38(2):
 552 295–307, 2015.

553 Jason Fries, Leon Weber, Natasha Seelam, Gabriel Altay, Debajyoti Datta, Samuele Garda, Sunny
 554 Kang, Rosaline Su, Wojciech Kusa, Samuel Cahyawijaya, et al. Bigbio: A framework for data-
 555 centric biomedical natural language processing. *Advances in Neural Information Processing*
 556 *Systems*, 35:25792–25806, 2022.

557 Imran Ul Haq, Mustafa Mhamed, Mohammed Al-Harbi, Hamid Osman, Zuhal Y Hamd, and Zhe
 558 Liu. Advancements in medical radiology through multimodal machine learning: A comprehensive
 559 overview. *Bioengineering*, 12(5):477, 2025.

560 Xuehai He, Yichen Zhang, Luntian Mou, Eric Xing, and Pengtao Xie. Pathvqa: 30000+ questions for
 561 medical visual question answering. *arXiv preprint arXiv:2003.10286*, 2020.

562 Xinyue Hu, L Gu, Q An, M Zhang, L Liu, K Kobayashi, T Harada, R Summers, and Y Zhu. Medical-
 563 diff-vqa: a large-scale medical dataset for difference visual question answering on chest x-ray
 564 images. *PhysioNet*, 12:13, 2023.

565 Yutao Hu, Tianbin Li, Quanfeng Lu, Wenqi Shao, Junjun He, Yu Qiao, and Ping Luo. Omnimedvqa:
 566 A new large-scale comprehensive evaluation benchmark for medical lvm. In *Proceedings of the*
 567 *IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 22170–22183, 2024.

568 Shih-Cheng Huang, Anuj Pareek, Saeed Seyyedi, Imon Banerjee, and Matthew P Lungren. Fusion
 569 of medical imaging and electronic health records using deep learning: a systematic review and
 570 implementation guidelines. *NPJ digital medicine*, 3(1):136, 2020.

571 Shih-Cheng Huang, Zepeng Huo, Ethan Steinberg, Chia-Chun Chiang, Curtis Langlotz, Matthew
 572 Lungren, Serena Yeung, Nigam Shah, and Jason Fries. Inspect: a multimodal dataset for pa-
 573 tient outcome prediction of pulmonary embolisms. *Advances in Neural Information Processing*
 574 *Systems*, 36:17742–17772, 2023.

575 Wisdom Ikezogwo, Saygin Seyfioglu, Fatemeh Ghezloo, Dylan Geva, Fatwir Sheikh Mohammed,
 576 Pavan Kumar Anand, Ranjay Krishna, and Linda Shapiro. Quilt-1m: One million image-text pairs
 577 for histopathology. *Advances in neural information processing systems*, 36:37995–38017, 2023.

578 Information eXtraction from Images (IXI) Project. Ixi dataset, 2024. URL
 579 <https://brain-development.org/ixi-dataset/>. Nearly 600 subjects with
 580 T1/T2/PD/MRA/DTI MRI acquired at three London hospitals.

581 Phillip Isola, Jun-Yan Zhu, Tinghui Zhou, and Alexei A Efros. Image-to-image translation with
 582 conditional adversarial networks. In *Proceedings of the IEEE conference on computer vision and*
 583 *pattern recognition*, pp. 1125–1134, 2017.

584 Zhenyu Jiang, Yuqi Xie, Jinhan Li, Ye Yuan, Yifeng Zhu, and Yuke Zhu. Harmon: Whole-body motion
 585 generation of humanoid robots from language descriptions. *arXiv preprint arXiv:2410.12773*,
 586 2024.

587 Robert Kaczmarczyk, Theresa Isabelle Wilhelm, Ron Martin, and Jonas Roos. Evaluating multimodal
 588 ai in medical diagnostics. *npj Digital Medicine*, 7(1):205, 2024.

594 Firas Khader, Gustav Müller-Franzes, Tianci Wang, Tianyu Han, Soroosh Tayebi Arasteh, Christoph
 595 Haarburger, Johannes Stegmaier, Keno Bressem, Christiane Kuhl, Sven Nebelung, Jakob Nikolas
 596 Kather, and Daniel Truhn. Multimodal deep learning for integrating chest radiographs and clinical
 597 parameters: A case for transformers. *Radiology*, 309(1):e230806, 2023. doi: 10.1148/radiol.
 598 230806.

599 Jiwon Kim, Jung Kwon Lee, and Kyoung Mu Lee. Accurate image super-resolution using very deep
 600 convolutional networks. In *Proceedings of the IEEE conference on computer vision and pattern*
 601 *recognition*, pp. 1646–1654, 2016.

602 Jason J Lau, Soumya Gayen, Asma Ben Abacha, and Dina Demner-Fushman. A dataset of clinically
 603 generated visual questions and answers about radiology images. *Scientific data*, 5(1):1–10, 2018.

604 Chunyuan Li, Cliff Wong, Sheng Zhang, Naoto Usuyama, Haotian Liu, Jianwei Yang, Tristan
 605 Naumann, Hoifung Poon, and Jianfeng Gao. Llava-med: Training a large language-and-vision
 606 assistant for biomedicine in one day. *Advances in Neural Information Processing Systems*, 36:
 607 28541–28564, 2023.

608 Tianbin Li, Yanzhou Su, Wei Li, Bin Fu, Zhe Chen, Ziyan Huang, Guoan Wang, Chenglong Ma,
 609 Ying Chen, Ming Hu, Yanjun Li, Pengcheng Chen, Xiaowei Hu, Zhongying Deng, Yuanfeng Ji,
 610 Jin Ye, Yu Qiao, and Junjun He. Gmai-vl & gmai-vl-5.5m: A large vision-language model and a
 611 comprehensive multimodal dataset towards general medical ai, 2024.

612 Jingyun Liang, Jiezhang Cao, Guolei Sun, Kai Zhang, Luc Van Gool, and Radu Timofte. Swinir: Im-
 613 age restoration using swin transformer. In *Proceedings of the IEEE/CVF international conference*
 614 *on computer vision*, pp. 1833–1844, 2021.

615 Chao Liao, Liyang Liu, Xun Wang, Zhengxiong Luo, Xinyu Zhang, Wenliang Zhao, Jie Wu, Liang
 616 Li, Zhi Tian, and Weilin Huang. Mogao: An omni foundation model for interleaved multi-modal
 617 generation. *arXiv preprint arXiv:2505.05472*, 2025.

618 Tianwei Lin, Wenqiao Zhang, Sijing Li, Yuqian Yuan, Binhe Yu, Haoyuan Li, Wanggui He, Hao
 619 Jiang, Mengze Li, Xiaohui Song, et al. Healthgpt: A medical large vision-language model for
 620 unifying comprehension and generation via heterogeneous knowledge adaptation. *arXiv preprint*
 621 *arXiv:2502.09838*, 2025.

622 Weixiong Lin, Ziheng Zhao, Xiaoman Zhang, Chaoyi Wu, Ya Zhang, Yanfeng Wang, and Weidi Xie.
 623 Pmc-clip: Contrastive language-image pre-training using biomedical documents. In *International*
 624 *Conference on Medical Image Computing and Computer-Assisted Intervention*, pp. 525–536,
 625 2023.

626 Bo Liu, Li-Ming Zhan, Li Xu, Lin Ma, Yan Yang, and Xiao-Ming Wu. Slake: A semantically-
 627 labeled knowledge-enhanced dataset for medical visual question answering. In *2021 IEEE 18th*
 628 *international symposium on biomedical imaging (ISBI)*, pp. 1650–1654. IEEE, 2021.

629 Haotian Liu, Chunyuan Li, Yuheng Li, and Yong Jae Lee. Improved baselines with visual instruction
 630 tuning. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*,
 631 pp. 26296–26306, 2024.

632 Shengjie Liu, Chuang Zhu, Feng Xu, Xinyu Jia, Zhongyue Shi, and Mulan Jin. Bci: Breast
 633 cancer immunohistochemical image generation through pyramid pix2pix. In *Proceedings of the*
 634 *IEEE/CVF Conference on Computer Vision and Pattern Recognition Workshops (CVPRW)*, pp.
 635 1815–1824, June 2022a.

636 Shengjie Liu, Chuang Zhu, Feng Xu, Xinyu Jia, Zhongyue Shi, and Mulan Jin. Bci: Breast
 637 cancer immunohistochemical image generation through pyramid pix2pix. In *Proceedings of the*
 638 *IEEE/CVF conference on computer vision and pattern recognition*, pp. 1815–1824, 2022b.

639 Jiasen Lu, Christopher Clark, Rowan Zellers, Roozbeh Mottaghi, and Aniruddha Kembhavi. Unified-
 640 io: A unified model for vision, language, and multi-modal tasks. In *The Eleventh International*
 641 *Conference on Learning Representations*.

648 Jiasen Lu, Christopher Clark, Sangho Lee, Zichen Zhang, Savya Khosla, Ryan Marten, Derek Hoiem,
 649 and Aniruddha Kembhavi. Unified-io 2: Scaling autoregressive multimodal models with vision
 650 language audio and action. In Proceedings of the IEEE/CVF Conference on Computer Vision and
 651 Pattern Recognition, pp. 26439–26455, 2024.

652 Chenglong Ma, Yuanfeng Ji, Jin Ye, Zilong Li, Chenhui Wang, Junzhi Ning, Wei Li, Lihao Liu,
 653 Qiushan Guo, Tianbin Li, et al. Meditok: A unified tokenizer for medical image synthesis and
 654 interpretation. arXiv preprint arXiv:2505.19225, 2025a.

655 Chenglong Ma, Yuanfeng Ji, Jin Ye, Lu Zhang, Ying Chen, Tianbin Li, Mingjie Li, Junjun He, and
 656 Hongming Shan. Towards interpretable counterfactual generation via multimodal autoregression.
 657 In International Conference on Medical Image Computing and Computer-Assisted Intervention,
 658 pp. 611–620. Springer, 2025b.

659 Chuofan Ma, Yi Jiang, Junfeng Wu, Jihan Yang, Xin Yu, Zehuan Yuan, Bingyue Peng, and Xiao-
 660 juan Qi. Unitok: A unified tokenizer for visual generation and understanding. arXiv preprint
 661 arXiv:2502.20321, 2025c.

662 Yiyang Ma, Xingchao Liu, Xiaokang Chen, Wen Liu, Chengyue Wu, Zhiyu Wu, Zizheng Pan, Zhenda
 663 Xie, Haowei Zhang, Xingkai Yu, et al. Janusflow: Harmonizing autoregression and rectified flow
 664 for unified multimodal understanding and generation. In Proceedings of the Computer Vision and
 665 Pattern Recognition Conference, pp. 7739–7751, 2025d.

666 Michael Moor, Qian Huang, Shirley Wu, Michihiro Yasunaga, Yash Dalmia, Jure Leskovec, Cyril
 667 Zakka, Eduardo Pontes Reis, and Pranav Rajpurkar. Med-flamingo: a multimodal medical few-shot
 668 learner. In Machine Learning for Health (ML4H), pp. 353–367. PMLR, 2023.

669 D. Nguyen, C. Chen, H. He, and C. Tan. Pragmatic radiology report generation. In Proceedings of
 670 Machine Learning for Health (ML4H), PMLR, 2023. doi: 10.48550/arXiv.2303.08715.

671 Andrew Sellergren, Sahar Kazemzadeh, Tiam Jaroensri, Atilla Kiraly, Madeleine Traverse, Timo
 672 Kohlberger, Shawn Xu, Fayaz Jamil, Cian Hughes, Charles Lau, et al. Medgemma technical report.
 673 arXiv preprint arXiv:2507.05201, 2025.

674 Karan Singhal, Tao Tu, Juraj Gottweis, Rory Sayres, Ellery Wulczyn, Mohamed Amin, Le Hou,
 675 Kevin Clark, Stephen R Pfahl, Heather Cole-Lewis, et al. Toward expert-level medical question
 676 answering with large language models. Nature Medicine, 31(3):943–950, 2025.

677 Luis R Soenksen, Yu Ma, Cynthia Zeng, Leonard Boussioux, Kimberly Villalobos Carballo,
 678 Liangyuan Na, Holly M Wiberg, Michael L Li, Ignacio Fuentes, and Dimitris Bertsimas. In-
 679 tegrated multimodal artificial intelligence framework for healthcare applications. NPJ digital
 680 medicine, 5(1):149, 2022.

681 Sanjay Subramanian, Lucy Lu Wang, Sachin Mehta, Ben Beglin, Madeleine Van Zuylen, Sravanthi
 682 Parasa, Sameer Singh, Matt Gardner, and Hannaneh Hajishirzi. Medicat: A dataset of medical
 683 images, captions, and textual references. arXiv preprint arXiv:2010.06000, 2020.

684 Tim Tanida, Philip Müller, Georgios Kaassis, and Daniel Rueckert. Interactive and explainable region-
 685 guided radiology report generation. In Proceedings of the IEEE/CVF Conference on Computer
 686 Vision and Pattern Recognition (CVPR), pp. 24101–24111, 2023. doi: 10.1109/CVPR52688.2023.
 687 02357.

688 Chameleon Team. Chameleon: Mixed-modal early-fusion foundation models. arXiv preprint
 689 arXiv:2405.09818, 2024a.

690 OpenGVLab Team. Internvl2: Better than the best—expanding performance boundaries of open-
 691 source multimodal models with the progressive scaling strategy, 2024b.

692 Omkar Chakradhar Thawakar, Abdelrahman M Shaker, Sahal Shaji Mullappilly, Hisham Cholakkal,
 693 Rao Muhammad Anwer, Salman Khan, Jorma Laaksonen, and Fahad Khan. Xraygpt: Chest
 694 radiographs summarization using large medical vision-language models. In Proceedings of the
 695 23rd workshop on biomedical natural language processing, pp. 440–448, 2024.

702 Adrian Thummerer, Erik van der Bijl, Arthur Jr Galapon, Florian Kamp, Mark Savenije, Christina
 703 Muijs, Shafak Aluwini, Roel JHM Steenbakkers, Stephanie Beuel, Martijn PW Intven, et al.
 704 Synthrad2025 grand challenge dataset: Generating synthetic cts for radiotherapy from head to
 705 abdomen. *Medical physics*, 52(7):e17981, 2025.

706 Liang Wang, Nan Yang, Xiaolong Huang, Binxing Jiao, Linjun Yang, Dixin Jiang, Rangan Majumder,
 707 and Furu Wei. Text embeddings by weakly-supervised contrastive pre-training. *arXiv preprint*
 708 *arXiv:2212.03533*, 2022.

709 Ting-Chun Wang, Ming-Yu Liu, Jun-Yan Zhu, Andrew Tao, Jan Kautz, and Bryan Catanzaro. High-
 710 resolution image synthesis and semantic manipulation with conditional gans. In *Proceedings of*
 711 *the IEEE conference on computer vision and pattern recognition*, pp. 8798–8807, 2018.

712 Xinlong Wang, Xiaosong Zhang, Zhengxiong Luo, Quan Sun, Yufeng Cui, Jinsheng Wang, Fan
 713 Zhang, Yueze Wang, Zhen Li, Qiying Yu, et al. Emu3: Next-token prediction is all you need.
 714 *arXiv preprint arXiv:2409.18869*, 2024.

715 Elisa Warner, Joonsang Lee, William Hsu, Tanveer Syeda-Mahmood, Charles E Kahn Jr, Olivier
 716 Gevaert, and Arvind Rao. Multimodal machine learning in image-based and clinical biomedicine:
 717 Survey and prospects. *International journal of computer vision*, 132(9):3753–3769, 2024.

718 Chaoyi Wu, Xiaoman Zhang, Ya Zhang, Hui Hui, Yanfeng Wang, and Weidi Xie. Towards gen-
 719 eralist foundation model for radiology by leveraging web-scale 2d&3d medical data. *Nature*
 720 *Communications*, 16(1):7866, 2025a.

721 Chaoyi Wu, Xiaoman Zhang, Ya Zhang, Hui Hui, Yanfeng Wang, and Weidi Xie. Towards gen-
 722 eralist foundation model for radiology by leveraging web-scale 2d&3d medical data. *Nature*
 723 *Communications*, 16(1):7866, 2025b.

724 Chengyue Wu, Xiaokang Chen, Zhiyu Wu, Yiyang Ma, Xingchao Liu, Zizheng Pan, Wen Liu, Zhenda
 725 Xie, Xingkai Yu, Chong Ruan, et al. Janus: Decoupling visual encoding for unified multimodal
 726 understanding and generation. In *Proceedings of the Computer Vision and Pattern Recognition*
 727 *Conference*, pp. 12966–12977, 2025c.

728 Chengyue Wu, Xiaokang Chen, Zhiyu Wu, Yiyang Ma, Xingchao Liu, Zizheng Pan, Wen Liu, Zhenda
 729 Xie, Xingkai Yu, Chong Ruan, et al. Janus: Decoupling visual encoding for unified multimodal
 730 understanding and generation. In *Proceedings of the Computer Vision and Pattern Recognition*
 731 *Conference*, pp. 12966–12977, 2025d.

732 Shengqiong Wu, Hao Fei, Leigang Qu, Wei Ji, and Tat-Seng Chua. Next-gpt: Any-to-any multimodal
 733 llm. In *Forty-first International Conference on Machine Learning*, 2024a.

734 Size Wu, Zhonghua Wu, Zerui Gong, Qingyi Tao, Sheng Jin, Qinyue Li, Wei Li, and Chen Change
 735 Loy. Openuni: A simple baseline for unified multimodal understanding and generation. *arXiv*
 736 *preprint arXiv:2505.23661*, 2025e.

737 Yecheng Wu, Zhuoyang Zhang, Junyu Chen, Haotian Tang, Dacheng Li, Yunhao Fang, Ligeng
 738 Zhu, Enze Xie, Hongxu Yin, Li Yi, et al. Vila-u: a unified foundation model integrating visual
 739 understanding and generation. *arXiv preprint arXiv:2409.04429*, 2024b.

740 Jinheng Xie, Weijia Mao, Zechen Bai, David Junhao Zhang, Weihao Wang, Kevin Qinghong Lin,
 741 Yuchao Gu, Zhijie Chen, Zhenheng Yang, and Mike Zheng Shou. Show-o: One single transformer
 742 to unify multimodal understanding and generation. *arXiv preprint arXiv:2408.12528*, 2024.

743 Weiwen Xu, Hou Pong Chan, Long Li, Mahani Aljunied, Rui Feng Yuan, Jianyu Wang, Chenghao
 744 Xiao, Guizhen Chen, Chaoqun Liu, Zhaodonghui Li, et al. Lingshu: A generalist foundation model
 745 for unified multimodal medical understanding and reasoning. *arXiv preprint arXiv:2506.07044*,
 746 2025a.

747 Zhiyang Xu, Juhai Chen, Zhaojiang Lin, Xichen Pan, Lifu Huang, Tianyi Zhou, Madiyan Khabsa,
 748 Qifan Wang, Di Jin, Michihiro Yasunaga, et al. Pisces: An auto-regressive foundation model for
 749 image understanding and generation. *arXiv preprint arXiv:2506.10395*, 2025b.

756 Zhiwen Yang, Haowei Chen, Ziniu Qian, Yang Yi, Hui Zhang, Dan Zhao, Bingzheng Wei, and Yan
 757 Xu. All-in-one medical image restoration via task-adaptive routing. In *International Conference*
 758 *on Medical Image Computing and Computer-Assisted Intervention*, pp. 67–77. Springer, 2024.

759

760 Jin Ye, Guoan Wang, Yanjun Li, Zhongying Deng, Wei Li, Tianbin Li, Haodong Duan, Ziyuan Huang,
 761 Yanzhou Su, Benyou Wang, et al. Gmai-mmbench: A comprehensive multimodal evaluation
 762 benchmark towards general medical ai. *Advances in Neural Information Processing Systems*, 37:
 763 94327–94427, 2024.

764

765 Syed Waqas Zamir, Aditya Arora, Salman Khan, Munawar Hayat, Fahad Shahbaz Khan, and
 766 Ming-Hsuan Yang. Restormer: Efficient transformer for high-resolution image restoration. In
 767 *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pp. 5728–
 768 5739, 2022.

769

770 Hong Zhang, Zhongjie Duan, Xingjun Wang, Yuze Zhao, Weiyi Lu, Zhipeng Di, Yixuan Xu, Yingda
 771 Chen, and Yu Zhang. Nexus-gen: A unified model for image understanding, generation, and
 772 editing. [arXiv preprint arXiv:2504.21356](https://arxiv.org/abs/2504.21356), 2025.

773

774 Kai Zhang, Jun Yu, Eashan Adhikarla, Rong Zhou, Zhiling Yan, Yixin Liu, Zhengliang Liu, Lifang
 775 He, Brian Davison, Xiang Li, et al. Biomedgpt: A unified and generalist biomedical generative
 776 pre-trained transformer for vision, language, and multimodal tasks. [arXiv e-prints](https://arxiv.org/abs/2305.2023a), pp. arXiv–2305,
 777 2023a.

778

779 Xiaoman Zhang, Chaoyi Wu, Ziheng Zhao, Weixiong Lin, Ya Zhang, Yanfeng Wang, and Weidi
 780 Xie. Pmc-vqa: Visual instruction tuning for medical visual question answering. [arXiv preprint](https://arxiv.org/abs/2305.10415)
 781 [arXiv:2305.10415](https://arxiv.org/abs/2305.10415), 2023b.

782

783 Chunting Zhou, Lili Yu, Arun Babu, Kushal Tirumala, Michihiro Yasunaga, Leonid Shamis, Jacob
 784 Kahn, Xuezhe Ma, Luke Zettlemoyer, and Omer Levy. Transfusion: Predict the next token and
 785 diffuse images with one multi-modal model. [arXiv preprint arXiv:2408.11039](https://arxiv.org/abs/2408.11039), 2024.

786

787 Jun-Yan Zhu, Taesung Park, Phillip Isola, and Alexei A Efros. Unpaired image-to-image translation
 788 using cycle-consistent adversarial networks. In *Proceedings of the IEEE international conference*
 789 *on computer vision*, pp. 2223–2232, 2017.

790

791

792

793

794

795

796

797

798

799

800

801

802

803

804

805

806

807

808

809

810 A APPENDIX
811812 APPENDIX TABLE OF CONTENTS
813

814 A.1 Implementation Details	17
815 A.1.1 Training Hyperparameters	17
816 A.1.2 Reconstruction Fidelity for Clinically Important Small Lesions	19
817 A.2 Dataset Statistics	20
818 A.2.1 Dataset Composition Details	20
819 A.2.2 Medical Domain and Modality Distribution	20
820 A.2.3 Modality and Anatomy Distribution	21
821 A.3 Data Enhancement Pipeline: CAG Implementation	22
822 A.3.1 Stage 1: Structured Description Generation	22
823 A.3.2 Stage 2: Caption Fusion Enhancement	22
824 A.3.3 Stage 3: Thinking-Enhanced Response Generation	23
825 A.4 Downstream Task Results	25
826 A.4.1 Medical Report Generation	25
827 A.4.2 Visual Question Answering	25
828 A.4.3 Medical Image Generation	26
829 A.4.4 Interleaved tasks	27
830 A.5 Expert Review Validation System	30
831 A.5.1 Expert Review Framework Overview	30
832 A.5.2 Evaluation Dimension Analysis	30
833 A.5.3 Dataset Quality Comparison Analysis	31
834 A.5.4 Medical Modality-Specific Analysis	31
835 A.6 Downstream Task Performance	32
836 A.6.1 Medical Report Generation	32
837 A.6.2 CXR Lung Opacity Image Translation	32
838 A.7 Failure Cases And Analysis	33
839 A.7.1 Medical Image Generation:	33
840 A.7.2 Medical Image Editing:	33
841 A.7.3 Medical Image Understanding:	35
842 A.8 Efficiency Analysis	42
843 A.9 Statement on the Use of Large Language Models	43

856
857
858
859
860
861
862
863

864 A.1 IMPLEMENTATION DETAILS
865866 A.1.1 TRAINING HYPERPARAMETERS
867868 Table 6: Training hyperparameters and configurations for the three-stage curriculum learning strategy
869 in UniMedVL. These stages collectively implement the Knowledge component of the OKA frame-
870 work.

	Stage 1 (Foundation)	Stage 2 (Instruction Tuning)	Stage 3 (Unified Multimodal)
Hyperparameters			
Learning rate	5×10^{-5}	2.5×10^{-5}	1.0×10^{-5}
Optimizer		AdamW	
Loss weight (CE : MSE)		0.25 : 1.0	
Training steps	85K	120K	70K
EMA ratio		0.995	
Image Resolution (VAE)	512-1024	512-1024	32-1024
Image Resolution (ViT)	378-980	224-518	378-980
Max tokens per sample	18.5K	20K	27K
Dropout		Text: 0.3, ViT/VAE: 0.05	
ViT training	Trainable	Frozen	Frozen
VAE training		Frozen	
Understanding branch		Trainable	
LLM training		Trainable	
Data Sampling Ratio (%)			
Text-Only	5	5	3
Text-to-Image (T2I)	25	45	35
Image-to-Text (I2T)	75	40	37
Interleaved	-	10	25

892 **Detailed Training Strategy Implementation.** Our training employs a three-stage curriculum
893 learning approach that implements the Knowledge component within the OKA framework. We use
894 the AdamW optimizer throughout all stages:

- 895 • Stage 1 (Foundation Training) establishes basic medical understanding over 85K steps with
896 a learning rate of 5×10^{-5} . The data composition prioritizes image-to-text tasks (75%),
897 complemented by text-to-image generation (25%) and pure text data (5%). This stage trains
898 both ViT and LLM components end-to-end while keeping the VAE frozen. The image
899 resolution is restricted with the range from 512-1024 pixels for the generation branch and
900 378-980 pixels for the understanding branch.
- 901 • Stage 2 (Instruction Tuning) extends training to 120K steps with a reduced learning rate of
902 2.5×10^{-5} . The data mixture evolves to balance text-to-image (45%) and image-to-text
903 (40%) tasks, while introducing interleaved multimodal datasets (10%). The ViT encoder is
904 frozen at this stage to preserve learned visual features. Token capacity increases to 20K per
905 sample.
- 906 • Stage 3 (Unified Multimodal Training) focuses on interleaved generation capabilities over
907 70K steps with a learning rate of 1.0×10^{-5} . This stage significantly increases interleaved
908 dataset usage (25%) while maintaining balanced generation (35%) and understanding (37%)
909 tasks. The expanded token budget (27K) and broader image resolution range (32-1024
910 pixels for generation) support interleaved tasks, including medical image super-resolution,
911 modality translation, and counterfactual generation.

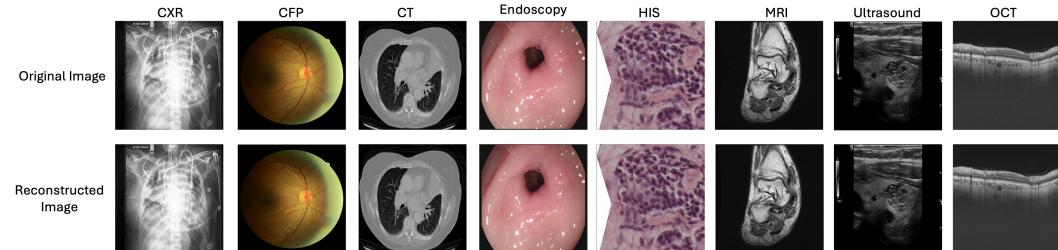
912 **Hardware Requirements and Training Infrastructure.** Our model training was conducted using
913 8x A800 GPUs (80GB memory each) for experimental validation. However, for optimal training
914 efficiency and to fully exploit the model’s capacity, we recommend a minimum configuration of 16x
915 A800 GPUs or equivalent hardware.916 **Technical Implementation Details.** The training employs a unified loss function that balances
917 understanding and generation objectives with a CE:MSE weight ratio of 0.25:1.0. We apply consistent

918 dropout rates across all stages (Text: 0.3, ViT/VAE: 0.05) to prevent overfitting. The EMA coefficient
 919 is set to 0.995 for stable model convergence. Throughout training, the VAE remains frozen to maintain
 920 stable latent representations.

921 **Rationale for Using Pretrained VAE without Fine-tuning.** Our approach leverages a general-
 922 purpose pretrained VAE model from FLUX (Black Forest Labs, 2024) without medical domain-
 923 specific fine-tuning. This design choice addresses two core questions: (1) the reconstruction capability
 924 of pretrained VAE on medical imaging modalities, and (2) the cost-benefit trade-off of fine-tuning
 925 versus preserving existing capabilities. Regarding the first question, we conducted comprehensive
 926 reconstruction experiments across eight medical imaging modalities to evaluate performance. For
 927 the second question, considering that our training data is not specifically designed for reconstruction
 928 optimization, we did not pursue domain-specific fine-tuning to avoid potential degradation of the
 929 model’s general-purpose capabilities while maintaining stable latent representations throughout our
 930 progressive training stages.

931
 932 Table 7: Reconstruction quality evaluation of pretrained VAE models on medical imaging modalities.
 933

Metric	Model	f_d	CFP	CT	CXR	Endoscopy	HIS	MRI	OCT	Ultrasound
rFID (Lower is Better)										
	VAE (FLUX)	8	13.22	5.81	5.42	11.77	10.00	10.58	13.23	9.64
	Direct End-to-end VAE (FLUX)	8	14.05	30.59	23.28	39.56	44.64	37.95	17.33	31.58
	VQGAN	8	27.22	15.97	33.57	27.73	21.33	67.68	29.48	18.66
	Emu3-VQ	8	16.27	11.83	27.91	20.83	13.52	69.89	25.43	11.99
	MedITok	16	14.39	7.88	22.27	10.66	6.32	46.54	17.64	6.55
PSNR (Higher is Better)										
	VAE (FLUX)	8	34.58	37.34	37.09	35.33	34.50	34.30	34.58	33.59
	Direct End-to-end VAE (FLUX)	8	35.11	34.43	31.28	31.98	29.69	34.82	30.83	35.17
	VQGAN	8	35.40	31.13	29.28	25.60	29.54	20.94	24.79	31.68
	Emu3-VQ	8	28.96	36.11	31.68	28.96	34.32	22.08	27.57	35.81
	MedITok	16	37.72	36.32	31.69	29.17	23.55	23.55	25.49	34.42
SSIM (Higher is Better)										
	VAE (FLUX)	8	0.892	0.951	0.973	0.934	0.922	0.921	0.892	0.938
	Direct End-to-end VAE (FLUX)	8	0.842	0.848	0.904	0.900	0.938	0.934	0.867	0.816
	VQGAN	8	0.923	0.885	0.753	0.768	0.844	0.484	0.248	0.317
	Emu3-VQ	8	0.943	0.928	0.793	0.847	0.957	0.547	0.751	0.955
	MedITok	16	0.953	0.937	0.855	0.890	0.972	0.660	0.935	0.883



960 Figure 5: **Qualitative comparison of VAE reconstruction quality across diverse medical imaging**
 961 **modalities.** Visual examples demonstrating reconstruction fidelity across eight medical imaging
 962 modalities (CFP, CT, CXR, Endoscopy, HIS, MRI, OCT, Ultrasound) using the pretrained FLUX
 963 VAE without domain-specific fine-tuning.

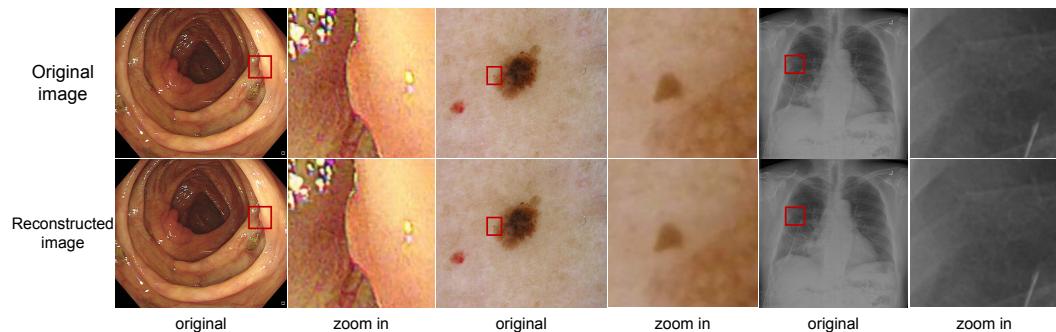
964
 965 The empirical evaluation demonstrates that the VAE (FLUX) achieves competitive reconstruction
 966 performance across eight distinct medical imaging modalities without requiring domain-specific
 967 fine-tuning. With a compression factor of $f_d = 8$, the model consistently delivers low rFID scores,
 968 competitive PSNR values, and robust SSIM scores. **To evaluate the necessity of domain-specific**
 969 **adaptation, we performed direct end-to-end fine-tuning of the FLUX VAE on medical imaging data**
 970 **(highlighted in red in Table 7).** The empirical results demonstrate that domain-specific fine-tuning
 971 yields negligible performance improvements for generation tasks across medical modalities, exhibiting
 972 inconsistent variations in rFID scores with marginal changes in corresponding metrics. Consequently,

972 these observations validate the deployment of pretrained VAE models without domain-specific
 973 fine-tuning for 2D medical imaging generation applications within our experimental framework.
 974

975 A.1.2 RECONSTRUCTION FIDELITY FOR CLINICALLY IMPORTANT SMALL LESIONS

977 While Table 7 demonstrates competitive aggregate reconstruction metrics (rFID, PSNR, SSIM) across
 978 diverse medical modalities, these metrics may not fully capture the preservation of small but clinically
 979 critical structures such as polyps in endoscopy, dermatoscopic features in skin lesions, or fractures
 980 in radiographs. To address this concern, we conducted targeted qualitative analysis focusing on
 981 the reconstruction fidelity of fine anatomical details and pathological findings that are essential for
 982 clinical diagnosis.

983 We selected representative cases from three imaging modalities where small lesion detection is
 984 clinically critical: (1) an endoscopy image containing polyp, (2) a dermoscopy image with skin lesion
 985 with globules, and (3) a x-ray image with fractures. For each case, we compared the original image
 986 with its VAE-reconstructed counterpart, examining both full-field views and magnified regions of
 987 interest (ROIs) centered on the lesions.



999
 1000 **Figure 6: Preservation of clinically important small lesions in VAE reconstruction.** Side-by-
 1001 side comparison of original images (top row) and VAE-reconstructed images (bottom row) across
 1002 three medical imaging scenarios. **Left:** Endoscopy image showing a polyp (highlighted in red
 1003 box). **Middle:** Dermoscopy image displaying skin lesion with globules (highlighted in red box).
 1004 **Right:** X-ray image with fractures (highlighted in red box). **Original** denotes the original image,
 1005 and **Zoom-in** denotes the zoomed-in view of the lesion within the red boxes. The magnified views
 1006 demonstrate that the FLUX VAE preserves fine structural details essential for clinical interpretation,
 1007 despite being a general-purpose encoder not specifically fine-tuned for medical imaging.

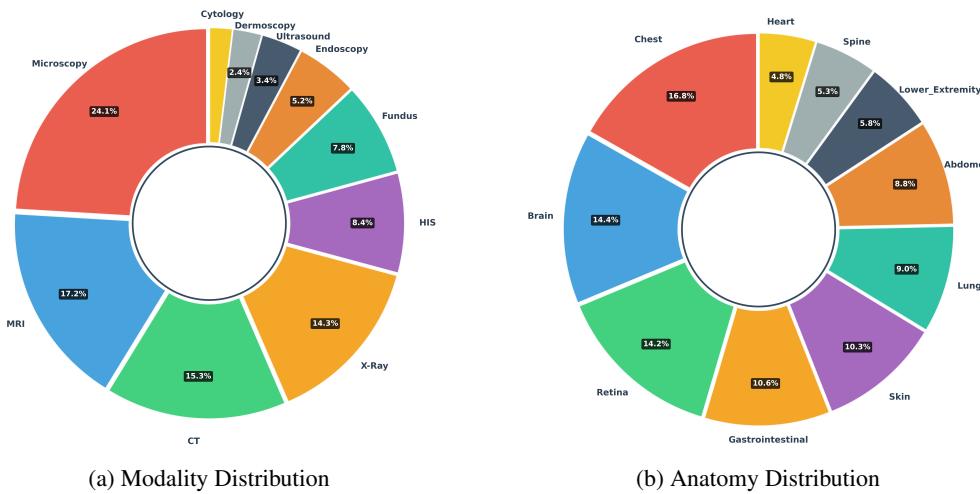
1008 Figure 6 illustrates that the pretrained FLUX VAE maintains visually discernible fidelity for small
 1009 pathological features. In the endoscopy image, the polyp’s morphology and surface texture remain
 1010 well preserved in the reconstruction, with boundary definition comparable to the original. For the
 1011 dermatoscopy image, the characteristic globular patterns, key diagnostic features for distinguishing
 1012 benign nevi from melanoma—are clearly visible in both the original and reconstructed versions. In
 1013 the X-ray image, the highlighted region and surrounding anatomical structures maintain structural
 1014 coherence post-reconstruction. These qualitative observations suggest that general-purpose VAE
 1015 also preserves clinically relevant fine-grained details that are critical for downstream diagnostic tasks
 1016 within our unified framework.

1026 **A.2 DATASET STATISTICS**
10271028 **A.2.1 DATASET COMPOSITION DETAILS**
10291030 Table 8: Overview of training stage data distribution, showing data composition, task types, and scale
1031 statistics across different stages. Stage 2 utilised the high-quality subset of stage 1 datasets.
1032

1033 Training Stage	1034 Total Entries	1035 Task Categories
Stage 1: Foundation Training		
1036 Understanding Tasks	4.0M	Image comprehension, VQA
1037 Generation Tasks	1.6M	Text-to-image, controllable generation
1038 <i>Stage 1 Subtotal</i>	<i>5.6M</i>	<i>Foundation capabilities</i>
Stage 2: Instruction Tuning		
1041 Understanding Tasks	698K	Image CoT, clinical reasoning
1042 Generation Tasks	668K	Enhanced T2I, medical translation
1043 CoT Understanding	317K	Chain-of-thought reasoning
1044 Text-only Tasks	230K	Medical QA, clinical dialogue
1045 <i>Stage 2 Subtotal</i>	<i>1.9M</i>	<i>Knowledge integration</i>
Stage 3: Unified Multimodal Training.		
1048 Interleaved Tasks	330K	5 interleaved tasks
1049 <i>Stage 3 Subtotal</i>	<i>0.33M</i>	<i>Unified capabilities</i>
1050 Total Dataset	5.6M	All medical tasks

1053 **A.2.2 MEDICAL DOMAIN AND MODALITY DISTRIBUTION**
10541055 Table 9: Major datasets detailed information, showing key dataset contributions sorted by data volume.
1056 For open-source datasets, the reported numbers indicate the actual subset sizes used in our training
1057 pipeline after filtering.
1058

1059 Dataset Name	1060 Total Entries	1061 Primary Tasks
1062 PMC-OA (Lin et al., 2023)	1.0M	Text-to-Image Generation
1063 Quilt-Im (Ikezogwo et al., 2023)	644K	Histopathology Understanding
1064 Healthqa (Lin et al., 2025)	638K	Clinical Reasoning, Image Caption
1065 PubMedVision (Chen et al., 2024a)	385K	Controllable T2I Generation
1066 Gmai-vl (Li et al., 2024)	288K	Enhanced T2I Generation
1067 Bigbio (Fries et al., 2022)	262K	Clinical Reasoning with CoT
1068 CheXpertPlus (Chambon et al., 2024)	223K	Medical Report Understanding
1069 PMC VQA (Zhang et al., 2023b)	204K	Image Caption
1070 Internvl (Chen et al., 2024c)	188K	Disease Classification, Clinical Reasoning
1071 Medicat (Subramanian et al., 2020)	132K	Controllable T2I Generation
1072 Medical-diff-vqa (Hu et al., 2023)	129K	Image Caption, Entity Recognition
1073 PMC-Infill (Wu et al., 2025a)	121K	Multi-image Understanding
1074 IXI T2/T1 SR 4x (Information eXtraction from Images (IXI) Project, 2024)	161K	Super resolution
1075 BraTS23 Modality Tran (Baltruschat et al., 2023)	52K	Cross modal synthesis
1076 SynthRAD Brain (MR to CT/CT to MR) (Thummerer et al., 2025)	66K	Cross modal synthesis
1077 SynthRAD Pelvis (MR to CT/CT to MR) (Thummerer et al., 2025)	42K	Cross modal synthesis
1078 ICG-CXR dataset (Ma et al., 2025b)	10K	Counterfactual generation
1079 BCI dataset (Liu et al., 2022a)	5K	Virtual immunohistochemistry staining
Total (Selected Datasets)	4.55M	—
Others Datasets	1.05M	—
Grand Total	5.6M	All Tasks

1080 A.2.3 MODALITY AND ANATOMY DISTRIBUTION
10811082 Figure 7 illustrates the comprehensive statistics of our curated medical datasets, showing both
1083 modality distribution and anatomical coverage. These statistics highlight the diversity and quality of
1084 data used across different training stages.1100 Figure 7: Comprehensive statistics of our curated medical datasets, showing modality distribution
1101 and anatomy distribution. These statistics demonstrate the diversity and quality of data used across
1102 different training stages.
1103
1104
1105
1106
1107
1108
1109
1110
1111
1112
1113
1114
1115
1116
1117
1118
1119
1120
1121
1122
1123
1124
1125
1126
1127
1128
1129
1130
1131
1132
1133

1134 A.3 DATA ENHANCEMENT PIPELINE: CAG IMPLEMENTATION
11351136 This section presents the complete prompt templates used in our Caption Augmented Generation
1137 (CAG) pipeline for image generation tasks, as described in Section 3. The CAG pipeline consists of
1138 two main stages: (1) structured medical description generation for quality control, and (2) caption
1139 fusion that combines original captions with generated descriptions.
11401141 A.3.1 STAGE 1: STRUCTURED DESCRIPTION GENERATION
11421143 Stage 1: Structured Description Generation Prompt
11441145 **Purpose:** Generate four-level structured medical image descriptions for quality control and
1146 similarity computation
11471148 You are a universally expert medical image analyst, proficient in all
1149 imaging modalities and anatomical systems.
1150 Your input is a single medical image, with no supplementary information.
1151 Your only task is to provide a comprehensive, objective, and structured
1152 description at four distinct levels, from the highest overview down to
1153 the most specific and exceptional findings.
1154 You must not offer any diagnostic, interpretive, or clinical advice.
11551156 ---
11571158 Output Structure (Four-Level, Top-to-Bottom -- definitions for your
1159 internal guidance; do NOT reproduce these headings in your answer)
1160

1161 LEVEL 1: IMAGE TYPE & GLOBAL CONTEXT

1162 • In one sentence, state the presumed imaging modality (if visually
1163 clear), main body region(s), and overall image category (e.g.,
1164 cross-sectional, projectional, histological).
1165 • Example: "This is an axial CT image of the abdomen and pelvis,
1166 showing cross-sectional anatomy at the level of the lower kidneys."
1167

1168 LEVEL 2: MACRO-ANATOMICAL OVERVIEW

1169 • In 2-4 concise lines, summarize the global distribution and layout
1170 of major anatomical regions, dominant structures, and any clearly
1171 visible large-scale abnormalities, masses, or disease patterns.
1172 • Describe anatomical orientation, symmetry, major organ relationships,
1173 and other visually prominent features.
1174

1175 LEVEL 3: ORGAN / SUBREGION DETAILS -- must be the most detailed section

1176 • In 6-12 lines (use complete sentences), describe the visual
1177 appearance of individual organs, vessels, bones, or other relevant
1178 subregions.
1179 • Provide precise, granular, reproducible details so that all main
1180 features can be reconstructed.
1181 • Maintain strict objectivity; do not include diagnostic language.
1182

1183 LEVEL 4: SPECIAL OR INCIDENTAL FINDINGS

1184 • List any unusual devices, postsurgical changes, image artifacts,
1185 rare morphologic features, or observations not already mentioned above.
1186 • If none are visible, explicitly state: "No distinct pathological
1187 or incidental findings are visible."
1188

1189 Writing Instructions

1190 1. Write the entire description as one continuous paragraph that
1191 implicitly follows the LEVEL 1 → LEVEL 4 order--do not include
1192 level headings, bullet points, or numbered lists in the paragraph.
1193 2. Do not use bullet points elsewhere (except within the examples).
1194 3. For more complex images, the portion corresponding to LEVEL 3 should
1195 naturally be longer; for simpler cases, keep it proportionally concise.
1196 4. Avoid any clinical judgement or speculation--describe only what is
1197 directly visible.
11981199 A.3.2 STAGE 2: CAPTION FUSION ENHANCEMENT
12001201 This stage fuses original captions with Stage 1 generated structured descriptions to create enhanced
1202 descriptions for image generation tasks.
1203

1188
1189

Stage 2: Caption Fusion Enhancement Prompt

1190
1191**Purpose:** Fuse original captions with structured descriptions for enhanced image generation prompts1192
1193

You are a universally expert medical image analyst, proficient in all imaging modalities and anatomical systems.

1194

CRITICAL CONSTRAINT: You must maintain absolute anatomical consistency. NEVER change, assume, or modify the anatomical location described in the original caption. Do not make assumptions about different anatomical locations or transfer descriptions between different body parts.

1195

Your input consists of:

1196

1. A structured, objective, four-level description derived from a locally deployed AI model (following a strict hierarchy from global overview to specific findings).
2. An original, data-derived textual description containing high-density, potentially diagnostic or interpretative information, which may lack structured clarity.

1197

Your task is to:

1198

- First, critically review and confirm the completeness of the structured description generated by the local model.
- Then, systematically extract and objectively incorporate relevant, visually verifiable details from the original data-derived description, enhancing information density without including diagnostic, interpretive, or clinical judgement.
- Clearly indicate and explicitly include visually evident anatomical abnormalities, structural deviations, or incidental observations present in the original data but omitted in the structured description.

1199

Output Structure (Four-Level, Top-to-Bottom)

1200

LEVEL 1: IMAGE TYPE & GLOBAL CONTEXT

1201

- In one sentence, state the presumed imaging modality, main body region(s), and overall image category.

1202

LEVEL 2: MACRO-ANATOMICAL OVERVIEW

1203

- In 2–4 concise lines, summarize global anatomical distribution, dominant structures, anatomical symmetry or deviations, and clearly visible large-scale abnormalities.

1204

LEVEL 3: ORGAN / SUBREGION DETAILS -- must be the most detailed section

1205

- In 6–12 complete sentences, describe individual organs, bones, vessels, and other relevant anatomical subregions in precise, reproducible detail.
- Objectively highlight visually confirmed abnormalities or structural deviations derived from the original data description.

1206

LEVEL 4: SPECIAL OR INCIDENTAL FINDINGS

1207

- Explicitly mention unusual devices, postsurgical changes, rare morphological features, or visually detectable anomalies present in the original description yet absent in the structured description.
- Clearly state the absence of commonly expected baseline anatomical or pathological features if definitively not observed in the image.

1208

Writing Instructions

1209

1. Write the final enhanced description as a single, continuous paragraph implicitly following LEVEL 1 → LEVEL 4 order--do not include explicit level headings, bullet points, or numbered lists.
2. Avoid any clinical judgement, diagnostic language, or speculative interpretation--include only details directly verifiable from visual inspection.
3. Start your output with "Please generate a realistic [modality] image showing" to make it a proper generation instruction.

1210

1211

1212

A.3.3 STAGE 3: THINKING-ENHANCED RESPONSE GENERATION

1213

1214

1215

1216

1217

1218

1219

1220

1221

1222

1223

1224

1225

1226

1227

1228

1229

1230

1231

1232

1233

1234

1235

1236

1237

1238

1239

1240

1241

This stage aims to elicit the reasoning process from the medical foundation model (MediGama-27B-IT) by prompting it to explicitly generate its internal thinking steps. We leverage this specialized medical model to simulate detailed reasoning processes through the structured prompt format. The resulting data, which includes both the explicit thinking traces and the final responses, is then used to train our model.

1242
1243

Stage 3: Thinking-Enhanced Response Generation Prompt (Revised v2)

1244
1245**Purpose:** Generate medical image responses with thinking tags for enhanced reasoning and quality control1246
1247
1248
1249

System: You are a medical image generator. You create [modality] images based on clinical descriptions. Your responses should describe what features you have generated in the image from the creator's perspective. Use bullet points to organize the anatomical structures and clinical features you have included in your generated image.

1250
1251

User: Based on this clinical description: "[clinical_description]"

1252
1253

You have been given the corresponding medical image. Please provide a response following this format:

1254
1255
1256
1257

Required format:
<think>Analyzing the clinical description, I need to generate an image that captures: 1) The key pathological process described, 2) The anatomical structures involved, 3) The specific imaging characteristics for [modality]. Based on the clinical presentation, I should include [key features reasoning]. [structured_caption if available]</think>

1258
1259
1260
1261

Here/This is the generated [modality] image that displays:
• [anatomical structure or clinical finding 1]
• [anatomical structure or clinical finding 2]
• [anatomical structure or clinical finding 3]

1262
1263

IMPORTANT:
1. In the <think> tag, reason through WHAT you need to generate and WHY based on medical knowledge
2. Respond from the GENERATOR perspective - describe what features you have CREATED/GENERATED in the image
3. Use the exact format above with bullet points (•) to list features
4. Start with 'Here is the generated [modality] image that displays:'
5. Each bullet point should describe a specific anatomical structure, clinical finding, or visual feature that you have included
6. Do NOT use observational language like 'shows', 'visible', 'can be seen' - instead use generative language like 'displays', 'includes', 'features', 'contains'

1264
1265
1266
1267
1268
1269
1270
1271
1272

Note: The thinking tag should reflect your decision-making process: "I need to generate X because Y", "The clinical description indicates I should include Z", etc.

1273
1274
1275
1276

The enhanced captions from Stage 2 (if the process "generating" is not generated successfully) and Stage 3 (if the process "thinking" is generated successfully) are sampled and then submitted to the Expert Review system (Section A.5) for final validation.

1277
1278
1279
1280
1281
1282
1283
1284
1285
1286
1287
1288
1289
1290
1291
1292
1293
1294
1295

1296 A.4 DOWNSTREAM TASK RESULTS
12971298 A.4.1 MEDICAL REPORT GENERATION
1299

1300

1301

1302

1303

1304

1305

1306

1307

1308

1309

1310

1311

1312

1313

1314

1315

1316

1317

1318

1319

1320

1321

1322

1323

1324

1325

1326

1327

1328 A.4.2 VISUAL QUESTION ANSWERING
1329

1330

1331

1332

1333

1334

1335

1336

1337

1338

1339

1340

1341

1342

1343

1344

1345

1346

1347

1348

1349

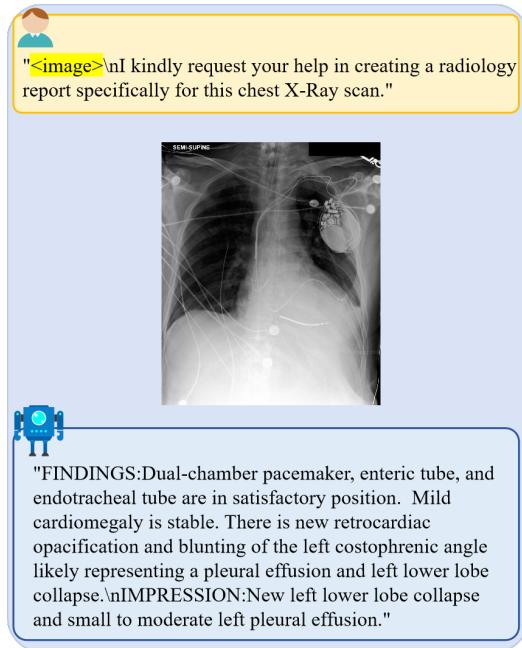


Figure 8: Medical report generation examples showing input images and generated reports.



Figure 9: Visual question answering examples across different medical imaging modalities.

1350
1351
1352
1353
1354

1355
1356
1357
1358
1359
1360
1361
1362
1363
1364
1365
1366
1367
1368

1369
1370

1371
1372
1373
1374

Figure 10: Visual question answering examples across different medical imaging modalities(part 2).

1375
1376
1377
1378
1379
1380

A.4.3 MEDICAL IMAGE GENERATION

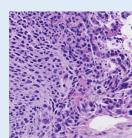
1381
1382
1383
1384
1385
1386
1387
1388
1389
1390

This is an axial CT image of the lower extremity, specifically depicting the distal portion of the femur and proximal portion of the tibia, showing cross-sectional anatomy. The image displays a cortical discontinuity through the distal femur, representing the right side, extending into the proximal tibial epiphysis. The fracture line appears transverse in orientation. The bone marrow space within the distal femur shows a trabecular pattern consistent with cancellous bone. The proximal tibia appears relatively normal, showing similar characteristics, with a dense cortex surrounding cancellous bone. The imaged portion of the distal femur on the left appears intact, with a normal trabecular pattern and cortical thickness. The lower part of the image shows a portion of the tibia, also appearing intact. The soft tissues surrounding the bone structures are not well visualized in detail due to the imaging plane and bone density, but no gross soft tissue abnormalities are apparent. The fracture is located in the immediate vicinity of the fracture. The distal femoral physes (growth plate) is not clearly delineated but is located adjacent to the fracture line. No significant displacement of the fracture fragments is apparent in this single axial plane. The image artifact of the gauntlet is visible along the lower border. The distal aspect of the right femur shows a clear break in the continuity of the cortical bone, separating the distal femur from the proximal tibia. The proximal tibial epiphysis is visible within both fragments. The proximal tibial epiphysis is also involved in the fracture line, showing a similar disruption of cortical continuity. There is no significant comminution visible in this plane. The absence of the contralateral distal femur is noted, as only the right side is fully visualized in the image. The absence of the distal fibula is also noted.



1391
1392
1393
1394
1395
1396
1397
1398
1399
1400
1401
1402

This is a histopathology image showing microscopic features of tissue architecture. The image displays a high-power view of neoplastic cells arranged in a disorganized pattern. The cells exhibit significant nuclear atypia, including pleomorphism (variation in size and shape), irregular nuclear contours, and prominent nucleoli. The nuclei also demonstrate hyperchromasia (dark staining). The cellular cytoplasm is scant. Mitotic figures are present. The tissue architecture is disrupted, with loss of normal glandular or ductal structures. There is invasion into the surrounding stroma by these atypical cells. Foci of necrosis, characterized by cellular debris and fragmented nuclear material, are observed. The extracellular matrix appears degraded in areas where tumor cells have infiltrated. The basal membrane is disrupted in regions of invasion. No distinct pathological or incidental findings are visible. No evidence of lymphocytes or inflammatory cells is present. No identifiable specific tissue of origin is definitely clear from this general view.



1403

Figure 11: Medical image generation examples with text prompts.

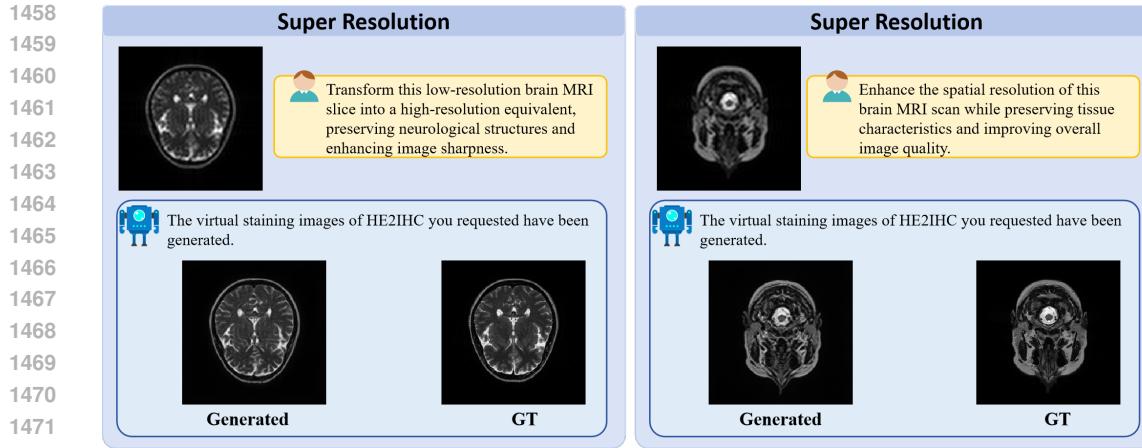


Figure 14: Super Resolution.

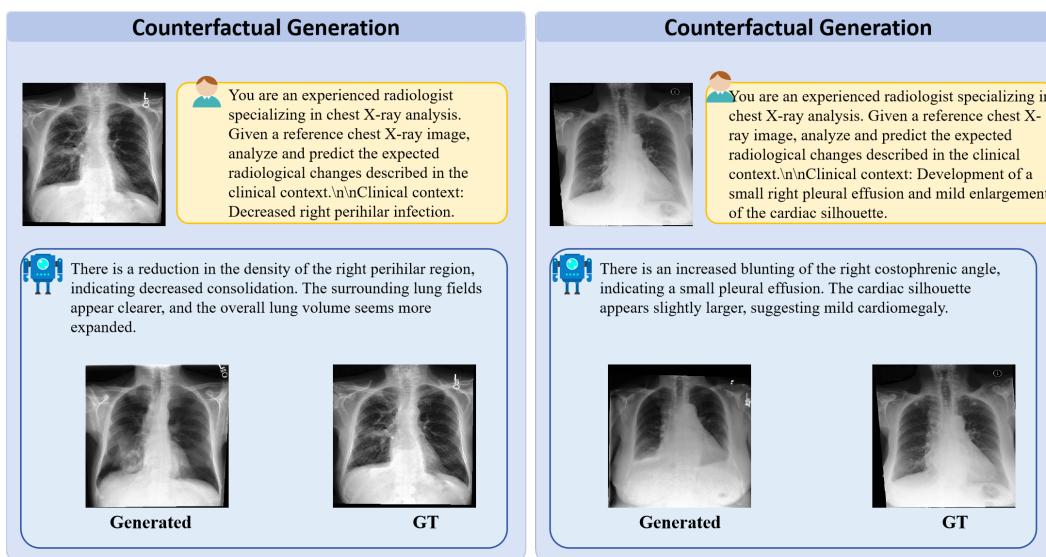


Figure 15: Counterfactual Generation.

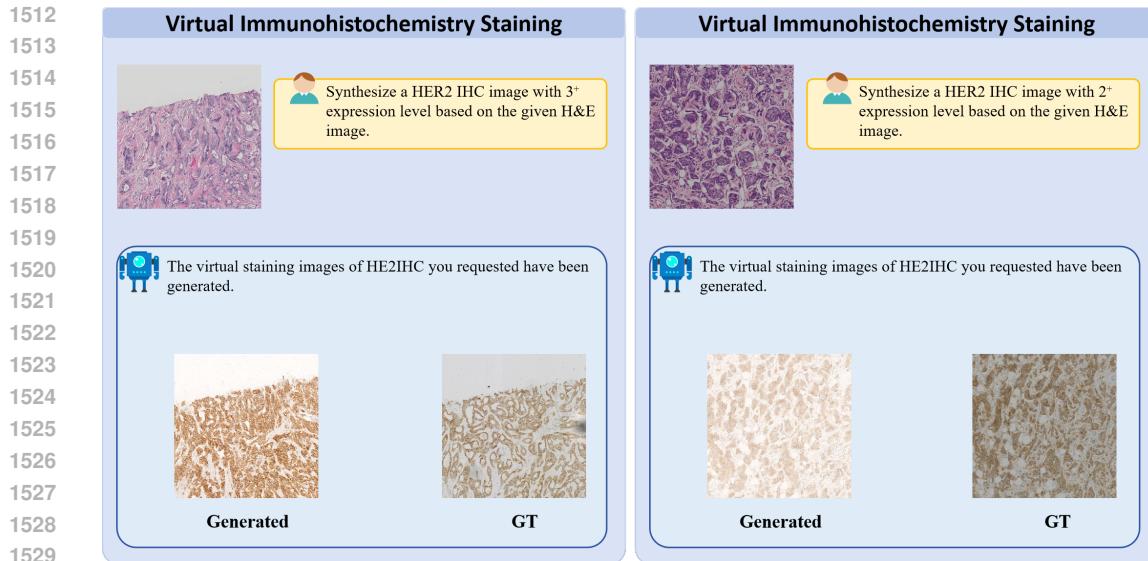


Figure 16: Virtual Immunohistochemistry Staining.

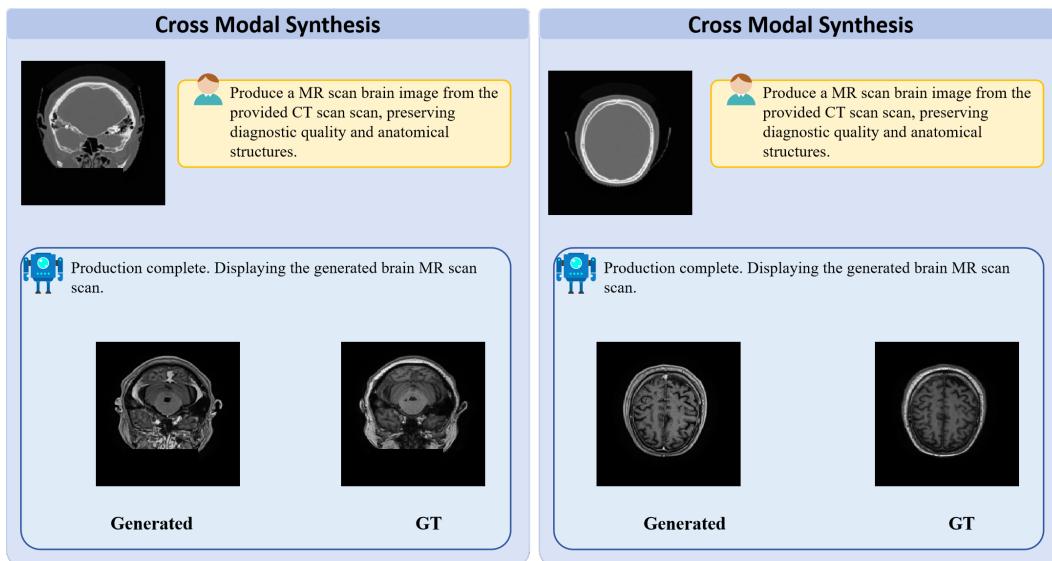
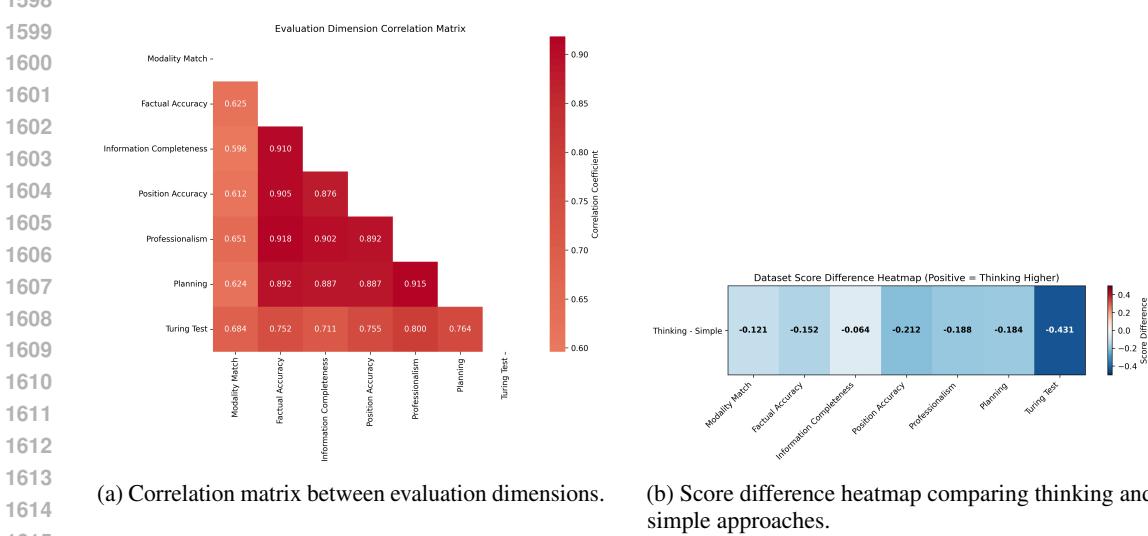
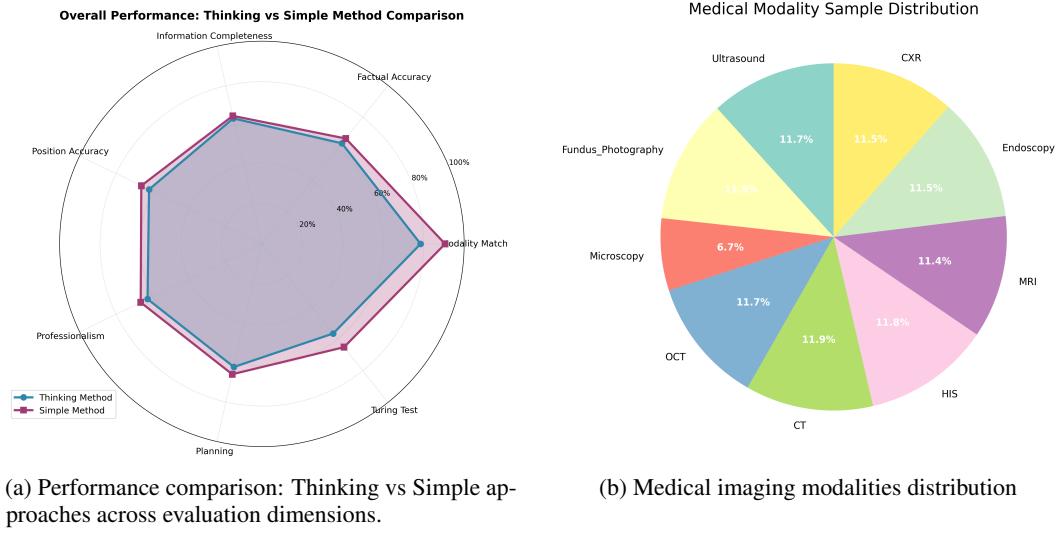


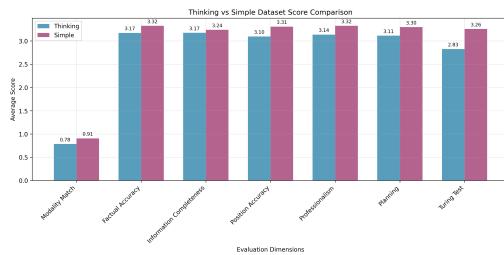
Figure 17: Cross Modal Synthesis.

1566 A.5 EXPERT REVIEW VALIDATION SYSTEM
15671568 This section presents an expert review validation system that evaluates the quality of our UniMed-5M
1569 dataset construction and two caption generation approaches described in the Data Enhancement
1570 Pipeline (Section A.3):1571 **Simple approach:** Caption fusion that combines structured descriptions from Stage 1 with original
1572 captions (Stage 2 of CAG pipeline).1573 **Thinking-enhanced approach:** Incorporates an additional planning process with <think> tags that
1574 integrates reasoning steps before medical image generation (Stage 3 of CAG pipeline). The validation
1575 system evaluates both data quality and methodological effectiveness.1577 A.5.1 EXPERT REVIEW FRAMEWORK OVERVIEW
15781579 Our expert review validation system is designed around a seven-dimensional medical evaluation
1580 framework that assesses medical AI performance.1581 Our evaluation framework encompasses seven dimensions that assess the synthetic quality of medical
1582 image captions. The framework begins with **Modality Match (0-1)**, which measures consistency
1583 between images and declared medical imaging modalities, followed by **Factual Accuracy (0-5)**
1584 that evaluates the precision of anatomical structure and pathological finding descriptions. **Info-
1585 rmation Completeness (0-5)** assesses coverage of diagnostically relevant key information, while
1586 **Position/Quantity Accuracy (0-5)** measures precision in anatomical localization and quantitative
1587 assessments. The framework also incorporates **Professionalism (0-5)** to evaluate adherence to
1588 medical reporting standards, **Planning Coherence (0-5)** to assess systematic thinking and logical
1589 organization quality, and finally **Clinical Reasoning (Turing Test) (0-5)** to measure approximation
1590 to human expert-level performance.1591 **Expert Validation Protocol:** Experts conducted audits of 200 samples across all seven dimensions.
1592 The evaluation process achieved inter-rater agreement exceeding 0.85 across all dimensions.
15931594 A.5.2 EVALUATION DIMENSION ANALYSIS
15951596 Figure 18 presents the correlation analysis and comparative results. Figure 18a shows inter-
1597 dimensional correlations, while Figure 18b compares the two generation approaches.1613 (a) Correlation matrix between evaluation dimensions. (b) Score difference heatmap comparing thinking and
1614 simple approaches.
1615
1616 **Figure 18: Expert evaluation analysis.** (a) Correlation matrix revealing inter-dimensional relation-
1617 ships (Pearson correlation coefficients ranging from 0.60 to 0.92). (b) Score difference heatmap
1618 comparing thinking and simple approaches (negative values indicate simple approach scores higher;
1619 all dimensions scored on 0-5 scale except Modality Match on 0-1 scale).

1620 A.5.3 DATASET QUALITY COMPARISON ANALYSIS
16211622 Figure 19 compares the two generation approaches across all evaluation dimensions. The radar chart
1623 (Figure 19a) shows closely aligned performance profiles.
1624

(a) Performance comparison: Thinking vs Simple approaches across evaluation dimensions.

(b) Medical imaging modalities distribution

1642 Figure 19: **Expert validation overview.** (a) Radar chart comparing performance profiles of thinking
1643 and simple approaches across all seven evaluation dimensions. (b) pie chart showing balanced
1644 representation across medical imaging modalities, ensuring comprehensive coverage.
16451646 A.5.4 MEDICAL MODALITY-SPECIFIC ANALYSIS
16471648 Figure 20 presents modality-specific performance across nine medical imaging modalities. Figure 20a
1649 shows statistical comparisons, and Figure 20b displays detailed performance metrics.
1650

(a) Statistical comparison between thinking and simple approaches.



(b) Modality-specific performance analysis.

1665 Figure 20: **Comprehensive performance analysis.** (a) Bar chart showing mean scores with confidence
1666 intervals. (b) heatmap displaying modality-specific performance scores.
16671668
1669
1670
1671
1672
1673

1674 A.6 DOWNSTREAM TASK PERFORMANCE
16751676 A.6.1 MEDICAL REPORT GENERATION
16771678 Table 10: **Medical report generation performance on MIMIC-CXR dataset.** Evaluation of
1679 automated radiology report generation using three metrics: ROUGE-L (lexical similarity), RaTE
1680 (radiology-specific terminology accuracy), and RadCliQ⁻¹ (clinical quality assessment). Higher
1681 scores indicate better performance for all metrics. Baseline results are sourced from Xu et al. (2025a).
1682 **Bold** indicates best performance and underlined indicates second-best performance.
1683

1684 Models	1685 MIMIC-CXR 1686 ROUGE-L	1687 MIMIC-CXR 1688 RaTE	1689 MIMIC-CXR 1690 RadCliQ ⁻¹
GPT-4.1	9.0	51.3	57.1
Claude Sonnet 4	20.0	45.6	53.4
Gemini-2.5-Flash	25.4	50.3	59.4
Med-R1-2B	19.3	40.6	42.4
MedLM-R1-2B	20.3	41.6	48.3
MedGemma-8B-IT	25.6	52.4	62.9
LLaVA-Med-7B	15.0	12.8	52.9
HuatuGPT-V-7B	23.4	48.9	48.2
BioMediX2-8B	20.0	44.4	53.0
Qwen2.5VL-7B	24.1	47.0	55.1
InternVL2-8B	23.2	47.0	56.2
InternVL3-8B	22.9	48.2	55.1
Lingshu-7B	30.8	<u>52.1</u>	69.2
HealthGPT-14B	21.4	48.4	52.7
HuatuGPT-V-34B	23.5	48.5	47.1
MedDr-40B	15.7	45.2	47.0
InternVL3-14B	22.0	48.6	46.5
Qwen2.5VL-32B	15.7	47.5	45.2
InternVL2.5-38B	22.7	47.5	54.9
InternVL3-38B	22.8	47.9	47.2
Lingshu-32B	<u>28.8</u>	50.8	<u>67.1</u>
UniMedVL	19.2	45.0	42.4

1703 A.6.2 CXR LUNG OPACITY IMAGE TRANSLATION
17041705 Table 11: **Unpaired chest X-ray opacity removal translation performance.** Task involves trans-
1706 forming CXRs with lung opacities to opacity-free counterparts while preserving anatomical structures.
1707 Evaluation metrics: FID (Fréchet Inception Distance) and KID (Kernel Inception Distance), where
1708 lower values indicate better image quality and distributional similarity. Baseline results are sourced
1709 from published literature. **Bold** indicates best performance and underlined indicates second-best
1710 performance.
1711

1712 Model	1713 FID ↓	1714 KID ↓
Baselines		
Original CXRs	81.80	0.043
Munit	109.4	0.073
Unit	103.2	0.061
CycleGAN	208.3	0.216
Uvrgan	210.4	0.225
Drit	117.6	0.087
AAMA-CDA	<u>67.18</u>	<u>0.016</u>
Unified Models		
HealthGPT-M3	62.19	0.031
UniMedVL[†]	35.1	0.008

1728
1729

A.7 FAILURE CASES AND ANALYSIS

1730
1731
1732

We present representative failure cases of UniMedVL organized by task type: medical image generation, medical image editing, and medical image understanding, with illustrative examples shown in Figures 21 and 22.

1733

A.7.1 MEDICAL IMAGE GENERATION:

1734
1735
1736
1737
1738
1739
1740
1741
1742
1743
1744

Global appearance and background artefacts. Although UniMedVL generally produces realistic images across modalities, a characteristic failure mode in text-to-image generation concerns embedded text and annotations (Figure 21). In some synthesised samples, the model hallucinates spurious on-image text or renders partially legible words, labels, or font styles that do not appear in the corresponding real clinical images or do not match typical acquisition overlays. These artefacts do not alter the main anatomical content but introduce visually unnatural patterns. As shown in Figure 21, the generated chest X-ray exhibits spurious text overlays (red boxes highlight the artefacts), the ultrasound image contains hallucinated text labels that deviate from standard clinical annotations, and the CT scan shows partially corrupted metadata text at the bottom that does not match typical DICOM overlay formatting.

1745
1746

A.7.2 MEDICAL IMAGE EDITING:

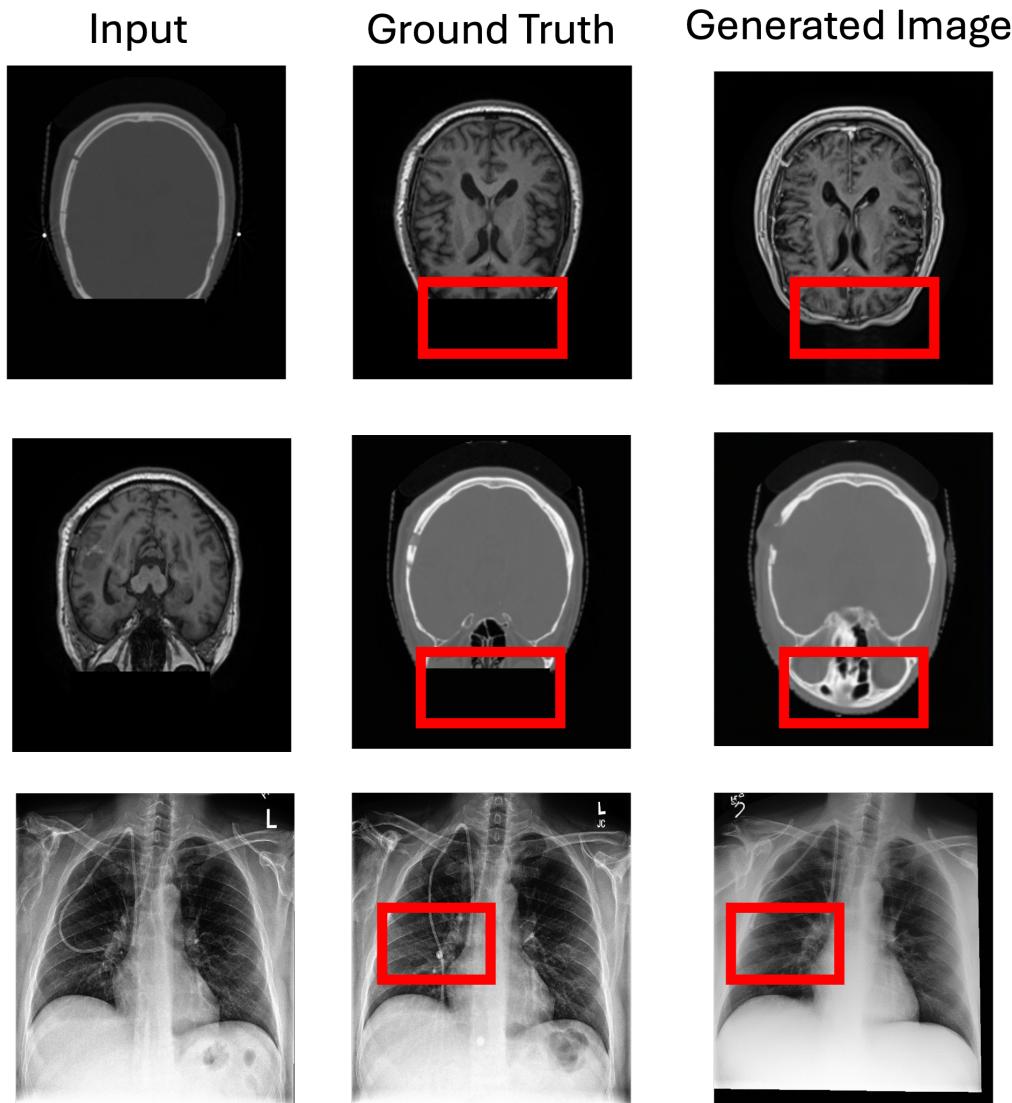
1747
1748
1749
1750
1751
1752
1753
1754
1755
1756
1757
17581759
1760
1761
1762
1763
1764
1765
1766

Figure 21: **Text and annotation artefacts in medical image generation.** Representative examples in the medical image generation task showing hallucinated or corrupted text elements across different imaging modalities. **Left:** Chest X-ray with spurious text overlay in the upper region (red box). **Center:** Ultrasound image displaying hallucinated text labels that do not conform to standard clinical annotation conventions (red box). **Right:** CT scan showing partially corrupted metadata text at the bottom edge that deviates from typical DICOM overlay formatting (red box).

1767
1768
1769
1770
1771
1772
1773
1774
1775
1776
1777
1778
1779
1780
1781

Structure preservation in generation and editing. For interleaved editing-style tasks (e.g., virtual staining, super-resolution, cross-modal synthesis, counterfactual generation), UniMedVL does not always perfectly preserve all spatial structures outside the region being semantically edited (Figure 22). For instance, in some counterfactual CXR generations, small devices or lines (e.g., catheters) can become slightly blurred or shifted, even when the main pathological change is correctly applied. Figure 22 illustrates these challenges across three representative cases: in the brain MRI cross-modal synthesis tasks (top two rows), the generated images show subtle structural discrepancies in the cerebellar and temporal regions compared to ground truth, and in the chest X-ray counterfactual generation (bottom row), while the model successfully modifies the target pathological region, occasionally minor shifts and blurring are observable in the cardiac silhouette boundaries.

1782
1783
1784
1785
1786
1787
1788
1789
1790
1791
1792
1793
1794
1795
1796
1797
1798
1799
1800



1823
1824
1825
1826
1827
1828
1829
1830
1831
1832
1833
1834
1835

Figure 22: Structural preservation challenges in interleaved generation and editing tasks. Comparative analysis in the medical image editing task across different medical image translation scenarios. Each row presents a triplet of **Input**, **Ground Truth**, and **Generated Image**. **Top row:** Brain MRI cross-modal synthesis, where the generated image exhibits subtle structural distortions in the cerebellar region compared to the ground truth. **Middle row:** Reverse brain MRI synthesis shows minor misalignment in the temporal lobe structures. **Bottom row:** Chest X-ray counterfactual generation task demonstrating reduction of pleural effusion; while the target pathological modification is applied, the generated image shows slight blurring and positional shifts in the cardiac silhouette and mediastinal borders.

1836 A.7.3 MEDICAL IMAGE UNDERSTANDING:

1837

1838 Table 12: **UniMedVL performance on GMAI-MMBench validation set.** Accuracy across 18
1839 clinical VQA sub-categories.

1840

1841 Task Category	1842 Accuracy
1843 Overall	1844 0.607
1845 Attribute Recognition	1846 0.659
1846 Blood Vessels Recognition	1847 0.593
1847 Bone	1848 0.623
1848 Cell Recognition	1849 0.513
1849 Counting	1850 0.457
1850 Disease Diagnosis	1851 0.669
1851 Image Quality Grading	1852 0.440
1852 Microorganism Recognition	1853 0.793
1853 Muscle	1854 0.580
1854 Nervous Tissue	1855 0.925
1855 Organ Recognition – Abdomen	1856 0.657
1856 Organ Recognition – Head and Neck	1857 0.845
1857 Organ Recognition – Pelvic	1858 0.560
1858 Organ Recognition – Thorax	1859 0.747
1859 Severity Grading	1860 0.372
1860 Surgeon Action Recognition	1861 0.287
1861 Surgical Instrument Recognition	1862 0.310
1862 Surgical Workflow Recognition	1863 0.471

1864

1865

1866

1867

1868

Medical demanding quantitative and fine-grained recognition scenarios. In medical demanding tasks that require precise quantitative reasoning or fine-grained visual recognition, the unified model shows clear limitations across multiple task categories. Table ?? presents UniMedVL’s performance on GMAI-MMBench validation set, revealing particularly low accuracy on surgical video recognition tasks: Surgeon Action Recognition, Surgical Instrument Recognition, and Surgical Workflow Recognition. The model also struggles with quantitative assessment tasks including Counting, Image Quality Grading, and Severity Grading.

1869

1870

1871

1872

1873

1874

1875

1876

1877

1878

1879

1880

1881

1882

1883

1884

1885

1886

1887

1888

1889

1890

1891

1892

1893

1894

1895

1896

1897

1898

1899

1900

1901

1902

1903

1904

1905

1906

1907

1908

1909

1910

1911

1912

1913

1914

1915

1916

1917

1918

1919

1920

1921

1922

1923

1924

1925

1926

1927

1928

1929

1930

1931

1932

1933

1934

1935

1936

1937

1938

1939

1940

1941

1942

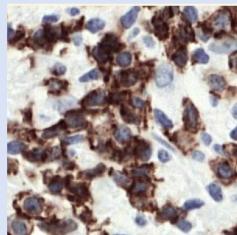
1943

Multimodal Medical Quantitative Assessment



From the image input, please count and tell me the number of lymphocytes shown.

- A: "the number of lymphocytes:30"
- B: "the number of lymphocytes:7"
- C: "the number of lymphocytes:24"**
- D: "the number of lymphocytes:25"



On this immunohistochemical slide, lymphocytes correspond to the brown-stained cells against blue nuclei. Visually estimating, I count ~30 such brown cells across the field. Therefore, among the options provided, I select **A: "the number of lymphocytes:30"** as the best match to the observed staining pattern.



The error arises from coarse visual estimation: clusters of overlapping brown staining are treated as separate lymphocytes, inflating the count. The model fails to delineate individual nuclei and to exclude weak or nonspecific staining, leading to overcounting versus the true value of 24 lymphocytes (option C).

Figure 23: **Counting failure in lymphocyte quantification. Green indicates the correct answer, and red indicates our model's prediction.**

1944

1945

1946

1947

1948

1949

1950

1951

1952

1953

1954

1955

1956

1957

1958

1959

1960

1961

1962

1963

1964

1965

1966

1967

1968

1969

1970

1971

1972

1973

1974

1975

1976

1977

1978

1979

1980

1981

1982

1983

1984

1985

1986

1987

Figure 24: **Surgeon Action Recognition failure.** **Green** indicates the correct answer, and **red** indicates our model’s prediction.

1988

1989

1990

1991

1992

1993

1994

1995

1996

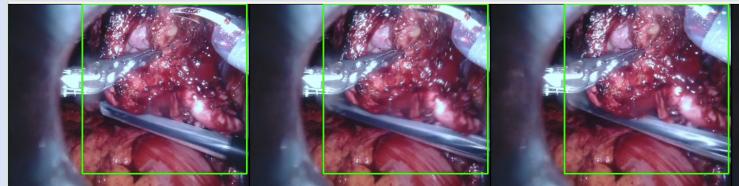
1997

Surgeon Action Recognition



Focusing on the boxed section of this endoscopy video, what surgical action is in progress?

- A: "bladder neck dissection"
- B: "pulling seminal vesicle"
- C: "bagging prostate"**
- D: "sucking smoke"



Based on the boxed endoscopic view, I compare visual cues with typical urologic steps. Bladder-neck dissection usually dominates such frames, whereas seminal-vesicle traction, prostate bagging, or smoke suction require more specific devices or anatomy not clearly visible. Given the limited context, I infer the action is bladder-neck dissection **(A)**.



This answer reflects procedural misclassification. The model relies on generic priors (“typical” frames of bladder-neck dissection) instead of carefully inspecting the boxed region for the retrieval bag and specimen, which are diagnostic of prostate bagging. It underutilizes the explicit prompt to recognize the specific surgical step in progress.

1998

1999

2000

2001

2002

2003

2004

2005

2006

2007

2008

2009

2010

2011

2012

2013

2014

2015

2016

2017

2018

2019

2020

2021

2022

2023

2024

2025

2026

2027

2028

2029

2030

2031

2032

2033

2034

2035

2036

2037

2038

2039

2040

2041

Figure 25: **Surgical Instrument Recognition failure.** Green indicates the correct answer, and red indicates our model's prediction.

2042

2043

2044

2045

2046

2047

2048

2049

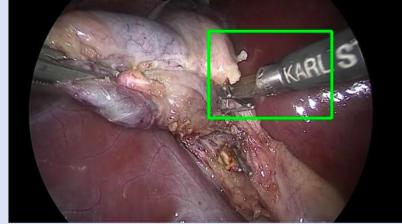
2050

2051

Surgical Instrument Recognition

Looking at the region delineated by a rectangle in the provided endoscopy image, which option would be the most target-like?

A: "bipolar"
B: "scissors"
C: "hook"
D: "irrigator"



Within the boxed endoscopic region, I compare the visible instrument with typical laparoscopic tools. The elongated tubular appearance suggests a fluid-delivery device rather than a cutting tool or curved hook. Thus, among the options "bipolar," "scissors," "hook," and "irrigator," I judge the instrument most consistent with an irrigator (D)

The mistake stems from coarse shape-based reasoning. The model focuses on the shaft's cylindrical form and ignores the distal jaws, which clearly resemble scissor blades rather than an irrigation nozzle. By not attending to fine-grained tip morphology and articulation, it misclassifies the laparoscopic scissors (correct answer B) as an irrigator.

2052
 2053
 2054
 2055
 2056
 2057
 2058
 2059
 2060
 2061
 2062
 2063
 2064
 2065
 2066
 2067
 2068
 2069
 2070
 2071
 2072
 2073
 2074
 2075
 2076
 2077
 2078
 2079
 2080
 2081
 2082
 2083
 2084
 2085
 2086
 2087
 2088
 2089
 2090
 2091
 2092
 2093
 2094
 2095
 2096
 2097
 2098
 2099
 2100
 2101
 2102
 2103
 2104
 2105

Surgical Workflow Recognition



Please identify which of the following options the surgical procedure shown in the image belongs to.
 A: "durotomy"
 B: "nasal corridor creation"
 C: "haemostasis"
 D: "anterior sphenoidotomy"



Endoscopic view shows nasal/sinus cavity with instrument working on bony/soft tissue surface. This appearance is more consistent with an endonasal approach to the sphenoid sinus than with durotomy or generic nasal corridor creation or simple haemostatic manoeuvres. Thus I classify it as anterior sphenoidotomy (D).



The model overemphasizes coarse anatomical context (endonasal access, apparent sinus wall) and underutilizes dynamic cues that the instrument is primarily controlling bleeding rather than creating a bony opening. Consequently, a supportive step (haemostasis, C) is misinterpreted as a definitive sinus procedure (anterior sphenoidotomy, D)

Figure 26: **Surgical Workflow Recognition failure.** Green indicates the correct answer, and red indicates our model's prediction.

2106
 2107
 2108
 2109
 2110
 2111
 2112
 2113
 2114
 2115
 2116
 2117
 2118
 2119
 2120
 2121
 2122
 2123
 2124
 2125
 2126
 2127
 2128
 2129
 2130
 2131
 2132
 2133
 2134
 2135
 2136
 2137
 2138
 2139
 2140
 2141
 2142
 2143
 2144
 2145
 2146
 2147
 2148
 2149
 2150
 2151
 2152
 2153
 2154
 2155
 2156
 2157
 2158
 2159

Image Quality Grading

 Assessing the boxed area in this endoscopy image, which artifact stands out?

A: "instrument artifacts"
B: "saturation"
 C: "blur"
 D: "low contrast"

 In the boxed region, a sharply bright, featureless patch is visible on the mucosal surface. I interpret this intense highlight as glare caused by the endoscope or another tool, i.e., an instrument-related artifact, rather than blur, low contrast, or general image-quality degradation. Therefore I select option A, "instrument artifacts."

 The model conflates extreme brightness with instrument glare, ignoring that no tool contour is visible and that the signal is locally clipped with complete texture loss. These properties are typical of sensor saturation rather than reflection from an instrument surface, so the correct dominant artifact is saturation (option B), not instrument artifacts.

Figure 27: **Image Quality Grading failure.** Green indicates the correct answer, and red indicates our model's prediction.

2160

2161

2162

2163

2164

2165

2166

2167

2168

2169

2170

2171

2172

2173

2174

2175

2176

2177

2178

2179

2180

2181

2182

2183

2184

2185

2186

2187

2188

2189

2190

2191

2192

2193

2194

2195

2196

2197

2198

2199

2200

2201

2202

2203

Figure 28: **Severity Grading failure.** Green indicates the correct answer, and red indicates our model's prediction.

2204

2205

2206

2207

2208

2209

2210

2211

2212

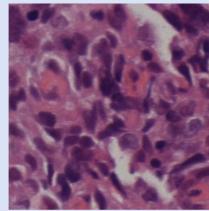
2213

Severity Grading



Please determine the severity of the provided image according to the core-needle biopsy grading protocol.

A: "axillary lymph node (aln) metastasis - n0"
 B: "axillary lymph node (aln) metastasis - n+(1-2)"
 C: "axillary lymph node (aln) metastasis - n+(>2)"



In this core-needle biopsy image, I see numerous metastatic tumor cells filling much of the field. According to the aln grading scheme (n0: none; n+(1-2): limited; n+(>2): extensive), this density appears compatible with involvement beyond two high-power fields, so I classify it as n+(>2) (C).



The error arises from misusing the HPF-based protocol: a single densely infiltrated patch is overinterpreted as implying involvement of >2 fields. The model extrapolates tumor extent beyond the visible image instead of restricting grading to sampled HPFs, so the burden actually corresponds to n+(1-2) metastasis (correct answer B).

2214 A.8 EFFICIENCY ANALYSIS
22152216 We measure inference efficiency with single expert activation. Comparison with BLIP3-o, Janus, and
2217 HealthGPT is provided in Tables 13.2218 **Image generation (Table 13(a)).** UniMedVL (14B) requires 40.03 TFLOPs per image and 28.39
2219 GB peak memory, whereas BLIP3-o 8B requires 142.59 TFLOPs per image and 25.55 GB. This
2220 corresponds to approximately 3.6x lower compute per image for UniMedVL with only about 11%
2221 higher peak memory (28.39 vs. 25.55 GB), even though UniMedVL uses a dual-encoder and MoT
2222 design and has almost twice the parameters. UniMedVL’s compute cost is also close to that of Janus
2223 7.42B (35.56 TFLOPs per image) while providing a larger unified model.2224 **VQA on GMAI-MMBench (Table 13(b)).** UniMedVL achieves 25.86 tokens/s with 2.256 TFLOPs
2225 per sample and 28.25 GB peak memory, compared with BLIP3-o’s 30.40 tokens/s, 9.307 TFLOPs
2226 per sample, and 18.21 GB. Thus, UniMedVL attains an approximately 4.1x reduction in FLOPs per
2227 sample with comparable throughput (about 85% of BLIP3-o’s tokens/s), at the cost of higher peak
2228 memory due to the dual-encoder design. Compared to another 14B unified model, HealthGPT-L14
2229 (12.57 tokens/s, 3.009 TFLOPs, 29.22 GB), UniMedVL is roughly twice as fast in throughput and
2230 more compute-efficient.2231 Table 13: **Efficiency Evaluation.** Comparison of throughput and computational costs across unified
2232 medical multimodal models with batch size 1.
22332234 (a) Image Generation Throughput
22352236 *Warm up with 10 images and measure efficiency over 20 images*

Model	Parameters	FLOPs/Image (TFLOPs)	Peak Mem (GB)
Janus	1B	10.01	5.19
HealthGPT-M3	3.8B	15.22	10.23
Janus	7.42B	35.56	17.10
BLIP3-o	8B	142.59	25.55
UniMedVL	14B	40.03	28.39

2246 (b) VQA Understanding Throughput (GMAI-MMBench validation set)

2247 *Warm up with 50 questions and measure efficiency over 150 VQA questions*

Model	Parameters	Tokens/s	FLOPs/Sample (TFLOPs)	Peak Mem (GB)
Janus	1B	70.11	0.498	4.46
HealthGPT-M3	3.8B	22.13	1.304	8.79
Janus	7.42B	52.94	1.894	14.59
BLIP3-o	8B	30.40	9.307	18.21
UniMedVL	14B	25.86	2.256	28.25
HealthGPT-L14	14B	12.57	3.009	29.22

2268
2269

A.9 STATEMENT ON THE USE OF LARGE LANGUAGE MODELS

2270
2271
2272
2273

Throughout the preparation of this manuscript, we utilized large language models to enhance the quality of the text. Specifically, we employed GPT and Claude for tasks related to language refinement, including correcting grammar and spelling, improving sentence clarity, and ensuring a consistent academic tone.

2274
2275
2276
2277
2278

The core scientific contributions, including the formulation of the problem, the proposed methodology, the design and execution of experiments, and the interpretation of results, are entirely the work of the authors. All text generated or modified by the LLM was meticulously reviewed, edited, and revised by the authors to ensure it accurately reflects our original ideas and findings. The authors bear full and final responsibility for all content presented in this paper.

2279
2280
2281
2282
2283
2284
2285
2286
2287
2288
2289
2290
2291
2292
2293
2294
2295
2296
2297
2298
2299
2300
2301
2302
2303
2304
2305
2306
2307
2308
2309
2310
2311
2312
2313
2314
2315
2316
2317
2318
2319
2320
2321