

Clinical Cognition Alignment for Gastrointestinal Diagnosis with Multimodal LLMs

Huan Zheng^{1,*}, Yucheng Zhou^{1,*}, Tianyi Yan¹, Dubing Chen¹,
Hongbo Lu^{2,3}, Wenlong Liao², Tao He², Pai Peng², Jianbing Shen^{1,✉}

¹SKL-IOTSC, CIS, University of Macau

²Shanghai Jiao Tong University

³COWARobot Co. Ltd.

Abstract. Multimodal Large Language Models (MLLMs) have demonstrated remarkable potential in medical image analysis. However, their application in gastrointestinal endoscopy is currently hindered by two critical limitations: the misalignment between general model reasoning and standardized clinical cognitive pathways, and the lack of causal association between visual features and diagnostic outcomes. In this paper, we propose a novel Clinical-Cognitive-Aligned (CogAlign) framework to address these challenges. First, we endow the model with rigorous clinical analytical capabilities by constructing the hierarchical clinical cognition dataset and employing Supervised Fine-Tuning (SFT). Unlike conventional approaches, this strategy internalizes the hierarchical diagnostic logic of experts, ranging from anatomical localization and morphological evaluation to microvascular analysis, directly into the model. Second, to eliminate visual bias, we provide a theoretical analysis demonstrating that standard supervised tuning inevitably converges to spurious background correlations. Guided by this insight, we propose a counterfactual-driven reinforcement learning strategy to enforce causal rectification. By generating counterfactual normal samples via lesion masking and optimizing through clinical-cognition-centric rewards, we constrain the model to strictly ground its diagnosis in causal lesion features. Extensive experiments demonstrate that our approach achieves State-of-the-Art (SoTA) performance across multiple benchmarks, significantly enhancing diagnostic accuracy in complex clinical scenarios.

Keywords: Multimodal Large Language Models · Gastrointestinal Diagnosis · Clinical Cognition Alignment · Counterfactual-Driven GRPO

1 Introduction

Gastrointestinal (GI) malignancies constitute a substantial portion of the global cancer burden, establishing endoscopic screening as the gold standard for early

*Equal contribution. ✉ Corresponding author: *Jianbing Shen*. This work was supported by the National Natural Science Foundation of China (No. 624B2002) and the Jiangyin Hi-tech Industrial Development Zone under the Taihu Innovation Scheme (EF2025-00003-SKL-IOTSC).

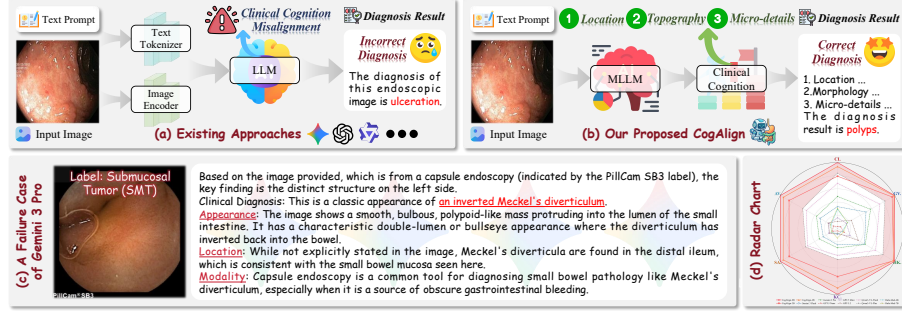


Fig. 1: Illustration of the motivation. (a) Existing methods suffer from clinical cognition misalignment. (b) Our CogAlign framework enforces a strict clinical cognitive flow. (c) A representative failure case generated by Gemini 3 Pro. (d) A radar chart highlighting the superior accuracy of CogAlign across diverse benchmarks.

detection and intervention [25]. Given the high dependence on operator experience and the inherent inter-observer variability in clinical practice, computer-aided diagnosis systems have emerged as a critical support tool to mitigate miss rates [8, 47]. Over the past decade, data-driven deep learning approaches, particularly Convolutional Neural Networks (CNNs) [13] and Vision Transformers (ViTs) [9, 41], have demonstrated expert-level proficiency in specialized tasks such as polyp detection [37] and lesion classification [39, 51]. Conventional paradigms are fundamentally restricted by their closed-set nature and opaque decision-making processes. These discriminative models typically function as silent classifiers that output rigid categorical labels without providing the underlying diagnostic rationale [42, 57]. Such opacity precludes clinical validation, undermining the diagnostic reliability required for high-stakes medical environments.

The recent advent of MLLMs marks a transformative shift from specialized discriminative models to generalized reasoning agents [34, 44, 45]. By synergizing the perceptual capabilities of advanced visual encoders with the extensive knowledge and inferential power of Large Language Models (LLMs), MLLMs introduce a versatile framework for endoscopic analysis [18, 23]. Unlike their predecessors, these foundation models possess the unique capacity to process visual information and generate coherent linguistic descriptions simultaneously [46]. This paradigm offers the potential to mimic the workflow of an endoscopist by identifying pathological features and providing comprehensive reports [6].

Despite this promise, the direct deployment of general MLLMs [3, 18] in gastrointestinal endoscopy is hindered by two critical limitations, as illustrated in Fig. 1 (a) and (b). The first is the misalignment between general model reasoning and standardized clinical cognitive pathways. In clinical practice, an endoscopist’s diagnosis follows a rigorous, hierarchical cognitive flow: initially localizing the anatomical site, subsequently evaluating morphological features, analyzing micro-details, and finally concluding with a diagnosis. In contrast, general MLLMs often exhibit scattered reasoning, skipping critical analytical steps or hallucinating non-existent features. This cognitive gap renders their outputs unreliable for high-stakes medical decisions. The second limitation is the lack of

causal association between visual features and diagnostic outcomes. MLLMs are susceptible to confounding visual factors, frequently relying on spurious correlations in the background, rather than characterizing the pathological lesion itself. As shown in the failure case in Fig. 1(c), even advanced models like Gemini 3 Pro can be misled by environmental artifacts, causing them to hallucinate a diagnosis based on the capsule modality context rather than the actual submucosal tumor features. This absence of causal grounding makes the models brittle and prone to failure when deployed in diverse clinical environments where such artifacts vary. As shown in Fig. 1(d), these deficiencies collectively constrain the diagnostic capability of existing models, resulting in suboptimal accuracy.

To address these challenges, we propose CogAlign for gastrointestinal diagnosis. Our approach is designed to bridge the gap between general reasoning and expert clinical protocols while ensuring diagnoses are strictly grounded in medical visual features. First, to tackle the clinical cognition misalignment, we construct a hierarchical clinical cognition dataset that encapsulates the step-by-step diagnostic logic of experts. We internalize this structured assessment process into the model, enforcing a diagnostic trajectory that moves strictly from anatomical localization and morphological evaluation to micro-details analysis.

Second, to resolve the issue of visual bias, we provide a theoretical analysis demonstrating that standard supervised tuning inevitably converges to spurious background shortcuts. Guided by this insight, we introduce a counterfactual-driven Group Relative Policy Optimization (GRPO) strategy for causal rectification. By masking lesion areas to generate counterfactual normal samples, we construct a counterfactual reference to isolate lesion-specific features. We then optimize the model using clinical-cognition-centric rewards, constraining the outcomes to be causally grounded in visual evidence of the lesion rather than background. Our contributions are summarized as follows:

- We propose CogAlign, a novel framework bridging the gap between general model capabilities and specialized clinical requirements. It integrates hierarchical cognitive tuning with counterfactual-driven reinforcement learning to ensure reliable gastrointestinal diagnosis.
- We construct a new dataset and apply SFT to instill rigorous analytical capabilities. This allows the model to emulate expert logic, progressing systematically from anatomical localization to microscopic detail analysis.
- We theoretically demonstrate that standard tuning relies on spurious background shortcuts and introduce a counterfactual-driven GRPO strategy to rectify this bias. Using counterfactual normal samples and clinical-cognition-centric rewards, we enforce strict causal grounding in pathological features.
- Extensive evaluations confirm that our approach achieves SoTA performance.

2 Related Work

2.1 Medical Multimodal Large Language Models

The rapid evolution of general Multimodal Large Language Models (MLLMs) has sparked significant interest in adapting these foundation models for the medical

domain [5, 27, 49, 52, 53]. By aligning powerful vision encoders with autoregressive language models, researchers have developed systems capable of interpreting complex clinical imagery and generating coherent text [22, 24]. Early pioneering models such as LLaVAMed [21] demonstrated the feasibility of adapting general visual instruction tuning to biomedicine [20]. These systems rely on vast datasets of image and text pairs to achieve proficiency in tasks like medical visual question answering, radiology report generation, and broad clinical reasoning [29, 38, 50].

Recent progress has focused on improving domain specific accuracy through parameter efficient fine tuning techniques [16] and specialized medical instruction datasets [30]. Researchers have successfully scaled these architectures to handle diverse modalities including X-rays, magnetic resonance imaging, and histopathology slides [43, 54–56]. Despite these impressive capabilities [1], current medical foundation models frequently struggle with diagnostic reliability in high stakes environments. They are prone to visual hallucinations and often act as superficial pattern matchers rather than genuine reasoning agents. Furthermore, standard training paradigms fail to enforce structured clinical logic, causing these models to skip critical analytical steps. They also exhibit severe vulnerability to visual bias, frequently grounding their textual outputs in spurious background correlations rather than genuine pathological evidence. Overcoming these fundamental limitations remains a primary hurdle for deploying multimodal models in reliable clinical assistance.

2.2 Gastrointestinal Disease Diagnosis

Computer Aided Diagnosis systems have become an integral component of modern gastroenterology, designed to assist clinicians in mitigating interobserver variability and reducing lesion miss rates during endoscopic screening [31]. Over the past decade, the field has been dominated by discriminative deep learning paradigms [19]. Convolutional Neural Networks and Vision Transformers have been extensively engineered to tackle specific gastrointestinal tasks, achieving expert level accuracy in polyp detection, anatomical landmark recognition, and ulcer classification [10, 32]. Advanced segmentation and object detection methods have been tailored to address the unique visual challenges of endoscopy, such as varying illumination, diverse organ topologies, and specular reflections [15, 37].

However, the clinical utility of these conventional methods is inherently restricted by their closed set nature and opaque decision making processes [2]. Traditional models function as silent classifiers that output rigid categorical predictions without providing the underlying diagnostic rationale [14]. To address the need for interpretability, recent literature has begun exploring report generation for endoscopy using vision language frameworks [28, 35]. While these preliminary multimodal approaches can produce descriptive text, they generally treat endoscopic analysis as a standard image captioning problem [7]. They fail to reflect the rigorous cognitive workflow of a senior endoscopist, which sequentially progresses from spatial anatomical localization to morphological assessment and finally to microscopic detail analysis [26]. Consequently, current




Fig. 2: Overview of the dataset curation pipeline. (a) shows the collection and filtering of diverse endoscopic images. (b) shows the generation of hierarchical clinical cognition reasoning chains. (c) shows the human expert refinement process to eliminate hallucinations. (d) shows a generated sample example.

models lack causal diagnostic grounding and remain highly susceptible to environmental noise such as surgical instrument artifacts and mucosal bubbles.

3 Hierarchical Clinical Cognition Dataset

Current public datasets for gastrointestinal endoscopy primarily consist of image-label pairs, lacking the intermediate reasoning steps required for transparent diagnosis [4, 17, 40]. Training on such data encourages models to learn shortcut features rather than clinical logic. Hence, we construct a hierarchical clinical cognition dataset designed to instill expert-level cognitive patterns into the MLLM.

3.1 Clinical Cognitive Hierarchy Definition

We define a standardized diagnostic protocol derived from the cognitive workflows of expert gastroenterologists. Unlike general image captioning, our annotation schema enforces a strict coarse to fine reasoning flow comprising three distinct stages prior to the final diagnosis. This hierarchical structure accurately mirrors the cognitive process of medical experts:

1. **Anatomical Localization:** Identification of the specific organ segment to provide essential spatial context and document the imaging conditions.
2. **Morphological Evaluation:** Assessment of macroscopic features, encompassing lesion shape, elevation, size, color, and boundaries.
3. **Micro-detail Analysis:** Scrutiny of fine grained surface patterns, such as villous structures, alongside vascular configurations.

3.2 Human-in-the-Loop Curation Pipeline

Manually annotating reasoning chains for large scale medical data is prohibitively expensive and time consuming. Therefore, we design a semi-automated curation pipeline incorporating a rigorous human in the loop mechanism.

During the data collection phase shown in Fig. 2(a), we aggregate diverse endoscopic images from public websites. A dedicated filtering process ensures the diversity and quality. Then, in the clinical cognition generation phase depicted in Fig. 2(b), we leverage Gemini 3 Pro [11], to act as a teacher model. By utilizing a specific prompt that explicitly outlines the three stage hierarchy defined above, we query the teacher model to generate structured reasoning descriptions for each input image. Finally, to eliminate potential hallucinations inherent to general multimodal models, we implement a data refinement phase detailed in Fig. 2(c). Human experts meticulously review the generated annotations. Annotations that pass the review are saved automatically, whereas samples containing factual errors fail the initial inspection and undergo manual revision by the experts.

3.3 Dataset Overview

We construct a comprehensive endoscopic dataset from five prominent public repositories, namely CrohnIPI [40], GastroVision [17], HyperKvasir [4], Kvasir-Capsule [36], and The SEE-AI Project [48], we assemble a total corpus of 24,515 samples. We establish a stratified split comprising 19,736 samples for training and 4,779 samples for testing. Specifically, the dataset encompasses 23 distinct single-label categories and 49 complex multi-label pathology combinations. As demonstrated by the shown example in Fig. 2(d), this curation process yields a high-quality dataset denoted as $\mathcal{D} = \{(\mathbf{x}_i, \mathbf{q}_i, \mathbf{r}_i, \mathbf{l}_i)\}_{i=1}^N$. In this formulation, $\mathbf{x}_i$ represents the image, q_i is the diagnostic query, $\mathbf{r}_i$ signifies the verified hierarchical clinical cognition reasoning chain, and $\mathbf{l}_i$ denotes the diagnostic label.

4 Methodology

4.1 Problem Definition

As illustrated in Fig. 3, the proposed CogAlign framework is designed to enforce a dual alignment: (1) aligning the model’s reasoning process with the standardized hierarchical cognitive pathways of clinical experts, and (2) grounding with causal pathological features rather than spurious background correlations.

Formally, given an image $\mathbf{x}$ and a diagnostic instruction $\mathbf{q}$, our goal is to generate a response $\mathbf{y}$ that not only provides the correct diagnostic label $\mathbf{l}$ but also produces a structured reasoning chain $\mathbf{r}$ that mirrors clinical standards:

$$\mathbf{y} = \mathbf{r} \oplus \mathbf{l} = \arg \max_{\mathbf{y}} P(\mathbf{y}|\mathbf{x}, \mathbf{q}; \theta), \quad (1)$$

where θ represents the trainable parameters of the MLLM, and $\oplus$ denotes the sequential concatenation of the reasoning process and the conclusion.

4.2 Clinical-Cognitive Reasoning Alignment

General MLLMs [3, 11], while possessing broad semantic knowledge, operate within an unconstrained generative space that often diverges from the disciplined sequential logic of expert endoscopists. To bridge this gap, we implement

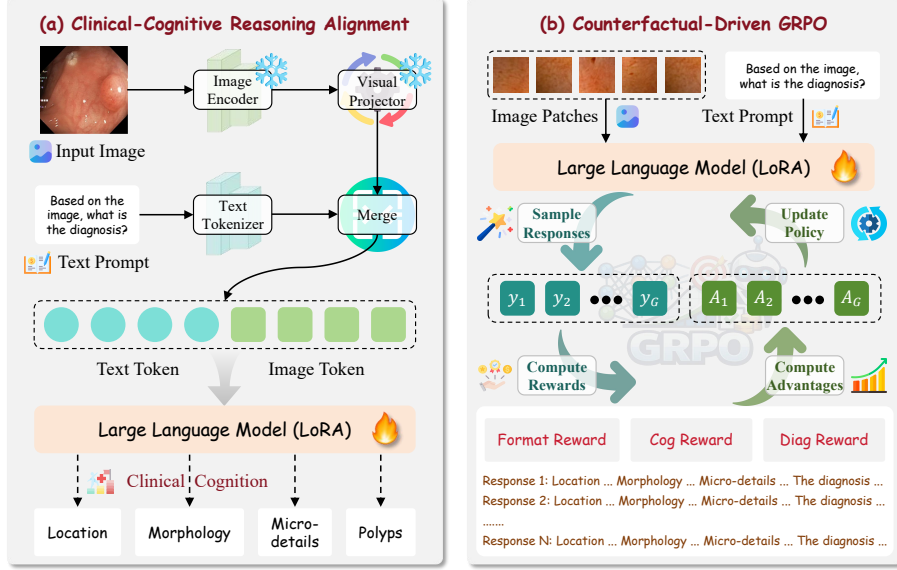


Fig. 3: Overview of the proposed CogAlign framework. The pipeline consists of two fundamental stages. Left panel demonstrates the clinical cognitive reasoning alignment phase, where the multimodal large language model undergoes supervised fine tuning. Right panel details the reinforcement learning phase guided by counterfactuals.

a Clinical-Cognitive Reasoning Alignment phase via SFT. The primary objective of this stage is to constrain the model’s generation manifold, forcing it to internalize the hierarchical reasoning chain $\mathbf{r}$, from anatomical localization to micro-detail analysis, before yielding a final diagnosis.

Formally, we utilize the hierarchical dataset $\mathcal{D} = \{(\mathbf{x}_i, \mathbf{q}_i, \mathbf{y}_i)\}_{i=1}^N$ constructed in Sec.3, where $\mathbf{y}_i = \mathbf{r}_i \oplus \mathbf{l}_i$ represents the target sequence concatenating the reasoning steps and the diagnostic conclusion. We employ a visual encoder to extract feature embeddings from the endoscopic image $\mathbf{x}_i$, which are projected into the LLM’s embedding space. The model is then optimized to generate the target sequence $\mathbf{y}_i$ in an autoregressive manner, effectively modeling the joint probability of the reasoning rationale and the diagnostic outcome. The optimization objective is defined as minimizing the negative log-likelihood of the next token:

$$\mathcal{L}_{\text{SFT}}(\theta) = -\frac{1}{N} \sum_{i=1}^N \sum_{t=1}^{L_i} \log P(\mathbf{y}_{i,t} | \mathbf{x}_i, \mathbf{q}_i, \mathbf{y}_{i,<t}; \theta), \quad (2)$$

where L_i denotes the length of the sequence $\mathbf{y}_i$, and θ represents the trainable parameters. Crucially, this objective enforces a strong statistical dependency: the final diagnosis $\mathbf{l}$ becomes a conditional consequence of the preceding morphological and micro-detail analysis contained within $\mathbf{r}$, rather than a direct, opaque classification from visual features.

4.3 Visual-Cognitive Misalignment and Causal Rectification

We provide a formal derivation of why SFT converges to a biased shortcut and how counterfactual intervention mathematically enforces causal grounding.

Definition 1 (Latent Factor Model). *An image X is generated by $\Psi : \mathcal{Z}_c \times \mathcal{Z}_e \rightarrow \mathcal{X}$, where Z_c and Z_e are causal and spurious latents. The diagnostic model is $f_\theta(X) = \sigma(\mathbf{w}^\top \phi(X))$, where $\phi(X) = [\phi_c(Z_c); \phi_e(Z_e)]$ is the feature.*

Definition 2 (Effective Feature Sensitivity). *The sensitivity of f to factor Z_i is defined as the norm of the Jacobian:*

$$\mathcal{S}_i = \|\nabla_{Z_i} f_\theta(\Psi(Z_c, Z_e))\|_2, \quad i \in \{c, e\}. \quad (3)$$

Theorem 1 (Shortcut Convergence in SFT). *Let $K(Z_e) < K(Z_c)$. Under gradient descent optimization of the SFT loss $\mathcal{L}$, the model parameters $\mathbf{w} = [\mathbf{w}_c; \mathbf{w}_e]$ satisfy $\|\mathbf{w}_e\| > \|\mathbf{w}_c\|$, leading to $\mathcal{S}_e > \mathcal{S}_c$.*

Proof. Consider the gradient flow of the SFT objective $\mathcal{L} = -\mathbb{E}[Y \log f + (1 - Y) \log(1 - f)]$. The dynamics of the weights for each feature are:

$$\frac{d\mathbf{w}_c}{dt} = -\eta \frac{\partial \mathcal{L}}{\partial \mathbf{w}_c} = \eta \mathbb{E}[(Y - f) \cdot \phi_c(Z_c)] \quad (4)$$

$$\frac{d\mathbf{w}_e}{dt} = -\eta \frac{\partial \mathcal{L}}{\partial \mathbf{w}_e} = \eta \mathbb{E}[(Y - f) \cdot \phi_e(Z_e)] \quad (5)$$

According to the Simplicity Bias principle [33], for low-complexity features Z_e , the spectral norm of the corresponding feature mapping ϕ_e is larger and converges faster in the early stages of gradient descent:

$$\|\phi_e(Z_e)\| \gg \|\phi_c(Z_c)\| \implies \left\| \frac{d\mathbf{w}_e}{dt} \right\| > \left\| \frac{d\mathbf{w}_c}{dt} \right\|. \quad (6)$$

As $t \rightarrow \infty$, $(Y - f) \rightarrow 0$. Since $\mathbf{w}_e$ captured the majority of the variance early on, the optimization stagnates before $\mathbf{w}_c$ is fully learned, yielding $\mathcal{S}_e > \mathcal{S}_c$.

Theorem 2 (Causal Rectification via Counterfactual Penalty). *Let $\mathcal{R}_{cf} = \mathbb{E}[f(\Psi(\mathbf{0}, Z_e))^2]$ be the counterfactual penalty. Minimizing the total objective $\mathcal{J} = \mathcal{L} + \lambda \mathcal{R}_{cf}$ as $\lambda \rightarrow \infty$ ensures $\mathcal{S}_e \rightarrow 0$.*

Proof. The optimal parameters θ^* must satisfy the stationary condition $\nabla_\theta \mathcal{J} = 0$:

$$\nabla_\theta \mathcal{L} + \lambda \nabla_\theta \mathcal{R}_{cf} = 0. \quad (7)$$

Substituting the gradient of the penalty term $\mathcal{R}_{cf}$:

$$\nabla_\theta \mathcal{L} + 2\lambda \mathbb{E} \left[f(X_{cf}) \cdot \frac{\partial f}{\partial \theta} \right] = 0. \quad (8)$$

As $\lambda \rightarrow \infty$, for the equation to hold, the model must satisfy $f(X_{cf}) \rightarrow 0$. Given $X_{cf} = \Psi(\mathbf{0}, Z_e)$, this implies:

$$\mathbf{w}_e^\top \phi_e(Z_e) \rightarrow 0, \quad \forall Z_e \in \mathcal{Z}_e. \quad (9)$$

Consequently, the sensitivity to spurious factors vanishes:

$$\mathcal{S}_e = \left\| \frac{\partial f}{\partial Z_e} \right\| = \left\| \mathbf{w}_e^\top \frac{\partial \phi_e}{\partial Z_e} \right\| \rightarrow 0. \quad (10)$$

To minimize the remaining $\mathcal{L}$ on the original samples, the model must re-orient its gradient flow toward $\mathbf{w}_c$, maximizing the reliance on causal features Z_c .

4.4 Counterfactual-Driven GRPO for Causal Alignment

Guided by the theoretical insights in Sec.4.3, we introduce a reinforcement learning framework termed counterfactual-driven GRPO. This stage operationalizes the counterfactual intervention to explicitly reward the model for grounding its diagnosis in causal lesion features.

Counterfactual Normal Sample Synthesis. To eliminate visual bias from background shortcuts, we construct counterfactual samples where lesion features are erased while the environment remains identical. First, the MLLM generates an initial lesion bounding box, which is rigorously refined by experts to define the precise lesion mask $\mathbb{M}$. Second, we apply high-intensity Gaussian smoothing to obliterate diagnostic features within $\mathbb{M}$:

$$\mathbf{x}_{cf} = \mathbf{x} \odot (1 - \mathbb{M}) + \mathcal{G}(\mathbf{x}, \sigma) \odot \mathbb{M}. \quad (11)$$

Finally, we assign a normal label and a corresponding negative reasoning chain to $\mathbf{x}_{cf}$. This paired sample $(\mathbf{x}_{cf}, \mathbf{r}_{cf}, \mathbf{l}_{cf})$ forces the model to ground its diagnosis strictly in lesion features; if it predicts pathology based on the unchanged background in $\mathbf{x}_{cf}$, it incurs a high optimization penalty.

Clinical-Cognition-Centric Rewards. To ensure the reasoning chain $\mathbf{r}$ is both structurally compliant and causally grounded, we design several rewards.

Output Format Reward. To enforce the model adherence to the strict hierarchical structure defined in our clinical cognitive pathway, we design a Format Reward R_{fmt} . The model’s output $\mathbf{y}$ must sequentially cover three critical sections: (1) Location & Imaging Environment, (2) Mucosal Morphology & Focal Lesions, and (3) Surface Texture & Microvascular Architecture. The reward function is defined as an all-or-nothing constraint:

$$R_{fmt}(y) = \mathbb{I} \left(\bigwedge_{s \in \mathcal{S}} (s \in y) \right), \quad (12)$$

where $\mathcal{S}$ represents the set of required section headers and $\mathbb{I}(\cdot)$ is the indicator function. If any section is missing, the reward is 0; otherwise, it is 1. This forces the model to maintain structural integrity during generation.

Clinical Cognition Reward. Merely following the correct format is insufficient; the content must capture specific semiological details. We propose a Clinical Cognition Reward R_{cog} to enforce semantic precision. For each ground truth reasoning chain, we utilize an LLM to pre-extract a set of critical keywords K_{gt} , consisting of exactly three key features for each of the three cognitive sections, totaling $|K_{gt}| = 9$ keywords. During training, we directly verify the presence of these keywords within the generated response $\mathbf{y}$. The reward is calculated as:

$$R_{cog}(\mathbf{y}, K_{gt}) = \frac{1}{9} \sum_{k \in K_{gt}} \mathbb{I}(k \in \mathbf{y}), \quad (13)$$

where $\mathbb{I}(k \in \mathbf{y})$ is an indicator function that returns 1 if the keyword k appears in the generated text $\mathbf{y}$. This mechanism ensures the model explicitly articulates all critical diagnostic criteria across the hierarchy.

Diagnostic Consistency Reward. The Diagnostic Consistency Reward R_{diag} evaluates the final conclusion extracted from the model’s response. Let $\mathbf{l}$ be the diagnosis parsed from $\mathbf{y}$ and $\mathbf{l}_{gt}$ be the ground truth label.

$$R_{diag}(\mathbf{y}, \mathbf{y}_{gt}) = \begin{cases} 1, & \text{if } \mathbf{l} = \mathbf{l}_{gt}, \\ 0, & \text{otherwise.} \end{cases} \quad (14)$$

This reward ensures that the reasoning chain culminates in the correct result.

GRPO Optimization. To align the model with the proposed rewards efficiently, we employ Group Relative Policy Optimization (GRPO), which estimates the baseline directly from the group average of sampled outputs. For each input query $\mathbf{q}$, we sample G outputs $\{y_1, \dots, y_G\}$ from the current policy $\pi_{\theta_{old}}$. We first compute the total reward $r_i = R_{fmt}(y_i) + \lambda_1 R_{cog}(y_i) + \lambda_2 R_{diag}(y_i)$ for each output y_i . To reduce gradient variance, we calculate the normalized group advantage $\hat{A}_i = (r_i - \mu_r) / (\sigma_r + \epsilon)$, where μ_r and σ_r are the mean and standard deviation of the rewards within the sampled group. Finally, we optimize the policy π_θ by maximizing the following surrogate objective alongside a KL divergence penalty to prevent deviation from the reference model π_{ref} :

$$\mathcal{J}_{GRPO}(\theta) = \mathbb{E}_{q \sim D} \left[\frac{1}{G} \sum_{i=1}^G \left(\min \left(\rho_i \hat{A}_i, \text{clip}(\rho_i, 1 - \epsilon, 1 + \epsilon) \hat{A}_i \right) - \beta \mathbb{D}_{KL}(\pi_\theta || \pi_{ref}) \right) \right] \quad (15)$$

where $\rho_i = \pi_\theta(y_i|q) / \pi_{\theta_{old}}(y_i|q)$ is the probability ratio.

5 Experiments

5.1 Experiment Setup

Implementation Details. We implement the two-stage training using the SWIFT framework with eight NVIDIA L20 GPUs. Stage 1 performs SFT for

Table 1: Quantitative comparison on five gastrointestinal benchmarks. We evaluate CogAlign against diverse models. Abbreviations: CI. (CrohnIPI), GV. (GastroVision), HK. (HyperKvasir), KC. (Kvasir-Capsule), SA. (The SEE-AI Project).

Model	CI.	GV.	HK.	KC.	SA.	Average
<i>Large Foundation Models</i>						
Gemini 3 Flash	20.87%	38.46%	43.24%	18.32%	15.01%	20.69%
Gemini 3 Pro	30.58%	44.73%	44.40%	21.83%	19.20%	24.82%
GPT-5 Nano	1.94%	3.99%	10.81%	2.77%	5.14%	5.06%
GPT-5 Mini	10.19%	11.97%	20.46%	6.50%	9.04%	10.04%
GPT-5.2	6.80%	18.80%	33.20%	5.32%	8.32%	11.13%
Qwen3-VL-Flash	43.20%	56.98%	61.00%	30.35%	31.65%	36.93%
Qwen3-VL-Plus	52.91%	64.10%	72.78%	34.72%	33.63%	41.16%
<i>Medical Foundation Models</i>						
Hulu-Med-4B	18.45%	13.68%	7.92%	6.50%	6.55%	7.72%
Hulu-Med-7B	19.42%	13.39%	9.46%	10.86%	6.22%	8.58%
<i>Small Foundation Models</i>						
Qwen3-VL-2B	18.93%	32.48%	33.20%	11.71%	12.01%	16.05%
Qwen3-VL-4B	36.89%	52.99%	50.39%	22.04%	25.50%	30.03%
Qwen3-VL-8B	39.32%	47.01%	67.57%	30.14%	29.22%	35.30%
<i>SFT on The Proposed Dataset</i>						
Qwen3-VL-2B (SFT)	41.26%	73.50%	87.26%	50.16%	48.39%	54.49%
Qwen3-VL-4B (SFT)	55.34%	76.07%	86.10%	64.75%	55.23%	61.98%
Qwen3-VL-8B (SFT)	62.14%	76.92%	89.38%	72.74%	58.77%	66.31%
<i>Our Proposed Models</i>						
CogAlign-2B	50.00%	73.79%	89.77%	53.99%	50.96%	57.40%
CogAlign-4B	59.22%	76.35%	89.19%	66.77%	57.22%	64.05%
CogAlign-8B	63.11%	77.21%	91.51%	74.01%	60.18%	67.67%

400 steps using the AdamW optimizer (learning rate 1×10^{-4} , cosine scheduler) and a global batch size of 128. The vision encoder and aligner are frozen, while we apply LoRA [16] (rank 16, $\alpha = 32$) to all linear modules, capping sequence length at 2048 tokens. Stage 2 applies GRPO [12] for 200 steps. This phase continues LoRA optimization with a global batch size of 256, a reduced learning rate of 1×10^{-6} , and a KL-divergence penalty $\beta = 0.04$. For each query, we sample $G = 8$ generations and compute an additive reward weighting format, clinical cognition, and diagnostic consistency at 1.0, 1.0, and 2.0, respectively.

Baselines. We evaluate the performance of CogAlign against a comprehensive suite of SoTA models. For the large foundation models, we include proprietary systems such as Gemini 3 Flash, Gemini 3 Pro, GPT-5.2, GPT-5 Mini, and GPT-5 Nano. We also benchmark against the Qwen3-VL series [3], specifically Qwen3-VL-Flash and Qwen3-VL-Plus. To assess the effectiveness of domain specific adaptation, we compare our framework with medical foundation models including HuluMed-4B and HuluMed-7B [18]. Furthermore, we evaluate small scale foundation models such as Qwen3-VL-2B, Qwen3-VL-4B, and Qwen3-VL-8B [3]. To isolate the contributions of our strategy, we include three internal variants: Qwen3-VL-2B (SFT), Qwen3-VL-4B (SFT) and Qwen3-VL-8B (SFT).

Evaluation Details. We evaluate CogAlign on a comprehensive test suite comprising a total of 4,779 endoscopic samples across five distinct datasets. These benchmarks include CrohnIPI [40], GastroVision [17], HyperKvasir [4], Kvasir-Capsule [36], and The SEE-AI Project [48]. Notably, The SEE-AI Project presents a significantly higher diagnostic challenge as it contains 235 multi-label

Table 2: Breakdown of Single-Label vs. Multi-Label diagnostic accuracy. We evaluate the ability to identify concurrent pathologies.

Model	Single-Label	Multi-Label	Average
<i>Large Foundation Models</i>			
Gemini 3 Flash	21.68%	1.70%	20.69%
Gemini 3 Pro	26.06%	0.85%	24.82%
GPT-5 Nano	5.30%	0.43%	5.06%
GPT-5 Mini	10.43%	2.55%	10.04%
GPT-5.2	11.69%	0.43%	11.13%
Qwen3-VL-Flash	38.34%	9.79%	36.93%
Qwen3-VL-Plus	42.76%	10.21%	41.16%
<i>Medical Foundation Models</i>			
Hulu-Med-4B	8.12%	0.00%	7.72%
Hulu-Med-7B	9.02%	0.00%	8.58%
<i>Small Foundation Models</i>			
Qwen3-VL-2B	16.81%	1.28%	16.05%
Qwen3-VL-4B	31.27%	5.96%	30.03%
Qwen3-VL-8B	36.77%	6.81%	35.30%
<i>SFT on The Proposed Dataset</i>			
Qwen3-VL-2B (SFT)	56.91%	7.66%	54.49%
Qwen3-VL-4B (SFT)	64.66%	10.21%	61.98%
Qwen3-VL-8B (SFT)	69.19%	10.64%	66.31%
<i>Our Proposed Models</i>			
CogAlign-2B	59.93%	8.09%	57.38%
CogAlign-4B	66.81%	10.64%	64.05%
CogAlign-8B	70.47%	13.62%	67.67%

samples, requiring the model to identify co-occurring pathologies simultaneously. We report accuracy as the primary evaluation metric. For the multi-label cases, we employ a strict accuracy standard where a prediction is considered correct only if it exactly matches the complete set of ground truth pathologies.

5.2 Main Results

Comparison with Large Foundation Models. Despite their massive parameter scales, general-purpose MLLMs often struggle in specialized medical contexts. As illustrated in Tab. 1, proprietary models like Gemini 3 Pro and GPT-5 series achieve moderate performance but lack consistency. Qwen3-VL-Plus perform better, yet they still fall short in challenging scenarios like Kvasir-Capsule and The SEE-AI Project. In contrast, our CogAlign achieves a remarkable accuracy, surpassing Qwen3-VL-Plus by a significant margin.

Comparison with Medical Foundation Models. Specialized medical models do not exhibit a competitive edge. This underperformance can be attributed to their training paradigms, which often focus on general medical visual-question answering rather than the rigorous, fine-grained visual recognition required for gastrointestinal endoscopy. The clinical cognition alignment strategy effectively bridges this gap, ensuring the model attends to lesion features.

5.3 Multi-Label Disease Diagnosis

In real-world clinical environments, patients frequently present with concurrent gastrointestinal pathologies, requiring models to identify multiple co-occurring



Fig. 4: Case study between CogAlign and baseline models. The top row demonstrates CogAlign’s ability to detect a subtle polyp via hierarchical clinical cognition. The bottom row highlights CogAlign’s robustness to visual noise in identifying erosion.

conditions rather than a single dominant lesion. As shown in Tab. 2, general foundation models struggle significantly in this setting, often exhibiting tunnel vision where secondary pathologies are ignored; for instance, specialized medical models like Hulu-Med-7B completely fail to detect multi-label cases. In contrast, our CogAlign framework demonstrates superior performance. This improvement confirms that our hierarchical reasoning chain and counterfactual-driven reinforcement learning effectively force the model to conduct a comprehensive scan of the mucosal surface rather than fixating on spurious or singular features.

5.4 Case Study

To provide a qualitative evaluation of our proposed approach, we present a comparative case study in Fig. 4. The top row illustrates the superiority of our framework over the general foundation model Qwen3 VL Plus. In this scenario, the endoscopic image contains a subtle polyp. The general model fails to identify the lesion and incorrectly predicts a normal mucosa. Conversely, our model leverages the internalized clinical cognitive pathway to systematically analyze the image. By sequentially evaluating the anatomical location, mucosal morphology, and microscopic details, our model accurately detects the lobulated protruding lesion and correctly concludes the diagnosis as polyps.

The bottom row highlights the effectiveness of our counterfactual driven reinforcement learning stage by comparing the full pipeline against the Base-SFT-8B variant. The input image is heavily obscured by environmental noise, specifically frothy bile stained mucus and bubbles. The Base-SFT-8B model, lacking causal diagnostic grounding, is misled by these environmental artifacts and hallucinates a normal diagnosis. In contrast, our fully trained model successfully ignores the




Fig. 5: Detailed analysis of model robustness and masking strategies. (a) Performance degradation under spot interference. (b) Comparison of masking techniques.

spurious visual noise. Guided by the causal alignment phase, it focuses precisely on the superficial mucosal disruption and accurately identifies the erosion.

5.5 Robustness Analysis

To evaluate the resilience of our proposed framework against environmental interference, we conduct a robustness analysis by applying simulated spot interference to the test images. This technique explicitly simulates the mucosal bubbles and specular reflections that frequently corrupt clinical endoscopic observations. As illustrated in Fig. 5(a), the baseline models fine tuned only with SFT suffer a severe degradation in diagnostic accuracy when exposed to visual perturbations. This vulnerability indicates that standard training paradigms overfit to spurious background correlations. In contrast, the complete CogAlign framework exhibits remarkable stability, maintaining high performance across all model scales.

5.6 Selection of Masking Strategy

The generation of counterfactual normal samples requires obliterating pathological evidence while preserving the surrounding contextual environment. We investigate the impact of different erasure techniques by comparing solid white masking against high intensity Gaussian blurring. As depicted in Fig. 5(b), employing a Gaussian blur to synthesize counterfactuals yields consistently higher diagnostic accuracy compared to utilizing solid white patches. We attribute this performance discrepancy to the naturalness of the modified images. Solid white masks introduce sharp artificial boundaries and out of distribution visual signals that can destabilize the reinforcement learning optimization process. Conversely, Gaussian blurring effectively neutralizes the diagnostic features while maintaining a smooth visual texture, thereby providing a reliable reference for causal rectification and enabling the model to accurately isolate lesion representations.

5.7 Ablation Study

Effect of Clinical Cognition Alignment. To validate the necessity of bridging the gap between general reasoning and standardized clinical protocols, we compare the performance of the vanilla foundation models against those fine-tuned




Fig. 6: Ablation study of individual modules in the CogAlign. We examine the contribution of SFT and the proposed reinforcement learning rewards.

on our hierarchical clinical cognition dataset. As observed in Fig. 6, applying our clinical cognition alignment via SFT dramatically significantly boosts this performance. This substantial improvement confirms that explicitly internalizing the expert cognitive flow is essential for unlocking the potential of MLLMs.

Effect of Clinical Cognition Reward. To assess the impact of rewards, we conduct an ablation study by removing the Clinical Cognition Reward R_{cog} from the full reward schema. As shown in Fig. 6, removing R_{cog} leads to a noticeable degradation in performance. Specifically, in the absence of constraints on semantic clinical features, the model’s intermediate reasoning degrades into vague descriptions that lack genuine visual-pathological grounding.

Effect of Diagnostic Consistency Reward. We further evaluate the contribution of the Diagnostic Consistency Reward R_{diag} . By excluding R_{diag} and relying solely on the format and cognition rewards, the model focuses heavily on generating descriptive text but occasionally fails to draw the correct conclusion from its own analysis. Experimental results in Fig. 6 indicate that removing this reward causes a significant decline in diagnostic accuracy.

6 Conclusion

In this paper, we proposed CogAlign, a novel framework designed to bridge the cognitive gap between general MLLMs and the rigorous standards of gastrointestinal diagnosis. Addressing the critical challenges of clinical cognitive misalignment and causal disconnect, we introduced a systematic clinical cognition alignment strategy. First, we constructed a hierarchical clinical cognition dataset and employed SFT to internalize expert-level diagnostic logic, compelling the model to strictly follow a trajectory from anatomical localization and morphological evaluation to micro-detail analysis. Second, guided by our theoretical analysis on shortcut convergence, we implemented a counterfactual-driven GRPO strategy. By utilizing counterfactual normal samples and clinical-cognition-centric rewards, we enforced causal rectification, ensuring diagnoses are grounded in pathological lesion features. Extensive experiments across five diverse benchmarks demonstrate that CogAlign establishes a new SoTA, significantly enhancing diagnostic performance in complex clinical scenarios.

References

1. Alkhaldi, A., Alnajim, R., Alabdullatef, L., Alyahya, R., Chen, J., Zhu, D., Alsinan, A., Elhoseiny, M.: Minigpt-med: Large language model as a general interface for radiology diagnosis. arXiv preprint arXiv:2407.04106 (2024)
2. Azad, R., Kazerouni, A., Heidari, M., Aghdam, E.K., Molaei, A., Jia, Y., Jose, A., Roy, R., Merhof, D.: Advances in medical image analysis with vision transformers: a comprehensive review. *Medical image analysis* **91**, 103000 (2024)
3. Bai, S., Cai, Y., Chen, R., Chen, K., Chen, X., Cheng, Z., Deng, L., Ding, W., Gao, C., Ge, C., et al.: Qwen3-vl technical report. arXiv preprint arXiv:2511.21631 (2025)
4. Borgli, H., Thambawita, V., Smedsrud, P.H., Hicks, S., Jha, D., Eskeland, S.L., Randel, K.R., Pogorelov, K., Lux, M., Nguyen, D.T.D., et al.: Hyperkvasir, a comprehensive multi-class image and video dataset for gastrointestinal endoscopy. *Scientific data* **7**(1), 283 (2020)
5. Chen, J., Cai, Z., Liu, Z., Yang, Y., Wang, R., Xiao, Q., Feng, X., Su, Z., Guo, J., Wan, X., et al.: Shizhengpt: Towards multimodal llms for traditional chinese medicine. arXiv preprint arXiv:2508.14706 (2025)
6. Chen, J., Gui, C., Ouyang, R., Gao, A., Chen, S., Chen, G.H., Wang, X., Cai, Z., Ji, K., Wan, X., et al.: Towards injecting medical visual knowledge into multimodal llms at scale. In: *Proceedings of the 2024 conference on empirical methods in natural language processing*. pp. 7346–7370 (2024)
7. Deria, A., Kumar, K., Dukre, A.M., Segal, E., Khan, S., Razzak, I.: Medmo: Grounding and understanding multimodal large language model for medical images. arXiv preprint arXiv:2602.06965 (2026)
8. Doi, K.: Computer-aided diagnosis in medical imaging: historical review, current status and future potential. *Computerized medical imaging and graphics* **31**(4-5), 198–211 (2007)
9. Dosovitskiy, A., Beyer, L., Kolesnikov, A., Weissenborn, D., Zhai, X., Unterthiner, T., Dehghani, M., Minderer, M., Heigold, G., Gelly, S., et al.: An image is worth 16x16 words: Transformers for image recognition at scale. arXiv preprint arXiv:2010.11929 (2020)
10. Fan, D.P., Ji, G.P., Zhou, T., Chen, G., Fu, H., Shen, J., Shao, L.: Pranet: Parallel reverse attention network for polyp segmentation. In: *International conference on medical image computing and computer-assisted intervention*. pp. 263–273. Springer (2020)
11. Google: Gemini 3 pro: the frontier of vision ai (2025), <https://blog.google/innovation-and-ai/technology/developers-tools/gemini-3-pro-vision/>
12. Guo, D., Yang, D., Zhang, H., Song, J., Wang, P., Zhu, Q., Xu, R., Zhang, R., Ma, S., Bi, X., et al.: Deepseek-r1 incentivizes reasoning in llms through reinforcement learning. *Nature* **645**(8081), 633–638 (2025)
13. He, K., Zhang, X., Ren, S., Sun, J.: Deep residual learning for image recognition. In: *Proceedings of the IEEE conference on computer vision and pattern recognition*. pp. 770–778 (2016)
14. He, Q., Bano, S., Stoyanov, D., Zuo, S.: Divgi: delve into digestive endoscopy image classification. *International Journal of Computer Assisted Radiology and Surgery* **20**(7), 1513–1520 (2025)
15. Hu, B.C., Ji, G.P., Shao, D., Fan, D.P.: Pranet-v2: Dual-supervised reverse attention for medical image segmentation. *Computational Visual Media* (2026)

16. Hu, E.J., Shen, Y., Wallis, P., Allen-Zhu, Z., Li, Y., Wang, S., Wang, L., Chen, W., et al.: Lora: Low-rank adaptation of large language models. *Iclr* **1**(2), 3 (2022)
17. Jha, D., Sharma, V., Dasu, N., Tomar, N.K., Hicks, S., Bhuyan, M.K., Das, P.K., Riegler, M.A., Halvorsen, P., Bagci, U., et al.: Gastrovision: A multi-class endoscopy image dataset for computer aided gastrointestinal disease detection. In: *Workshop on machine learning for multimodal healthcare data*. pp. 125–140. Springer (2023)
18. Jiang, S., Wang, Y., Song, S., Hu, T., Zhou, C., Pu, B., Zhang, Y., Yang, Z., Feng, Y., Zhou, J.T., et al.: Hulu-med: A transparent generalist model towards holistic medical vision-language understanding. *arXiv preprint arXiv:2510.08668* (2025)
19. Kröner, P.T., Engels, M.M., Glicksberg, B.S., Johnson, K.W., Mzaik, O., van Hooft, J.E., Wallace, M.B., El-Serag, H.B., Krittanawong, C.: Artificial intelligence in gastroenterology: A state-of-the-art review. *World journal of gastroenterology* **27**(40), 6794 (2021)
20. Lai, Y., Zhong, J., Li, M., Zhao, S., Li, Y., Psounis, K., Yang, X.: Med-r1: Reinforcement learning for generalizable medical reasoning in vision-language models. *IEEE Transactions on Medical Imaging* (2026)
21. Li, C., Wong, C., Zhang, S., Usuyama, N., Liu, H., Yang, J., Naumann, T., Poon, H., Gao, J.: Llava-med: Training a large language-and-vision assistant for biomedicine in one day. *Advances in Neural Information Processing Systems* **36**, 28541–28564 (2023)
22. Lin, T., Zhang, W., Li, S., Yuan, Y., Yu, B., Li, H., He, W., Jiang, H., Li, M., Song, X., et al.: Healthgpt: A medical large vision-language model for unifying comprehension and generation via heterogeneous knowledge adaptation. *arXiv preprint arXiv:2502.09838* (2025)
23. Liu, S., Zheng, B., Chen, W., Peng, Z., Yin, Z., Shao, J., Hu, J., Yuan, Y.: Endobench: A comprehensive evaluation of multi-modal large language models for endoscopy analysis. *arXiv preprint arXiv:2505.23601* (2025)
24. Moor, M., Huang, Q., Wu, S., Yasunaga, M., Dalmia, Y., Leskovec, J., Zakka, C., Reis, E.P., Rajpurkar, P.: Med-flamingo: a multimodal medical few-shot learner. In: *Machine learning for health (ML4H)*. pp. 353–367. PMLR (2023)
25. Motta, J.P., Wallace, J.L., Buret, A.G., Deraison, C., Vergnolle, N.: Gastrointestinal biofilms in health and disease. *Nature reviews Gastroenterology & hepatology* **18**(5), 314–334 (2021)
26. Mullappilly, S.S., Kurpath, M.I., Mohamed, O., Zidan, M., Khan, F., Khan, S., Anwer, R., Cholakkal, H.: Medix-r1: Open ended medical reinforcement learning. *arXiv preprint arXiv:2602.23363* (2026)
27. Mullappilly, S.S., Kurpath, M.I., Pieri, S., Alseieri, S.Y., Cholakkal, S., Aldahmani, K., Khan, F., Anwer, R., Khan, S., Baldwin, T., et al.: Bimedix2: Bio-medical expert lmm for diverse medical modalities. *arXiv preprint arXiv:2412.07769* (2024)
28. Nath, V., Li, W., Yang, D., Myronenko, A., Zheng, M., Lu, Y., Liu, Z., Yin, H., Law, Y.M., Tang, Y., et al.: Vila-m3: Enhancing vision-language models with medical expert knowledge. In: *Proceedings of the Computer Vision and Pattern Recognition Conference*. pp. 14788–14798 (2025)
29. Ning, J., Li, W., Tang, C., Lin, J., Ma, C., Zhang, C., Liu, J., Chen, Y., Gao, S., Liu, L., et al.: Unimedvl: Unifying medical multimodal understanding and generation through observation-knowledge-analysis. *arXiv preprint arXiv:2510.15710* (2025)
30. Pan, J., Liu, C., Wu, J., Liu, F., Zhu, J., Li, H.B., Chen, C., Ouyang, C., Rueckert, D.: Medvlm-r1: Incentivizing medical reasoning capability of vision-language

- models (vlms) via reinforcement learning. In: International Conference on Medical Image Computing and Computer-Assisted Intervention. pp. 337–347. Springer (2025)
31. Ramoni, D., Scuricini, A., Carbone, F., Liberale, L., Montecucco, F.: Artificial intelligence in gastroenterology: Ethical and diagnostic challenges in clinical practice. *World Journal of Gastroenterology* **31**(10), 102725 (2025)
 32. Roth, M., Nowak, M.V., Krenzer, A., Puppe, F.: Domain-adaptive pre-training of self-supervised foundation models for medical image classification in gastrointestinal endoscopy. *arXiv preprint arXiv:2410.21302* (2024)
 33. Shah, H., Tamuly, K., Raghunathan, A., Jain, P., Netrapalli, P.: The pitfalls of simplicity bias in neural networks. *Advances in Neural Information Processing Systems* **33**, 9573–9585 (2020)
 34. Shool, S., Adimi, S., Saboori Amleshi, R., Bitaraf, E., Golpira, R., Tara, M.: A systematic review of large language model (llm) evaluations in clinical medicine. *BMC Medical Informatics and Decision Making* **25**(1), 117 (2025)
 35. Shu, Y., Liu, C., Chen, R., Li, D., Dai, B.: Fleming-vl: Towards universal medical visual reasoning with multimodal llms. *arXiv preprint arXiv:2511.00916* (2025)
 36. Smedsrud, P.H., Thambawita, V., Hicks, S.A., Gjestang, H., Nedrejord, O.O., Næss, E., Borgli, H., Jha, D., Berstad, T.J.D., Eskeland, S.L., et al.: Kvasir-capsule, a video capsule endoscopy dataset. *Scientific Data* **8**(1), 142 (2021)
 37. Soleymanjahi, S., Huebner, J., Elmansy, L., Rajashekar, N., Lüdtke, N., Paracha, R., Thompson, R., Grimshaw, A.A., Foroutan, F., Sultan, S., et al.: Artificial intelligence-assisted colonoscopy for polyp detection: a systematic review and meta-analysis. *Annals of internal medicine* **177**(12), 1652–1663 (2024)
 38. Sun, H., Jiang, Y., Lou, W., Zhang, Y., Li, W., Wang, L., Liu, M., Liu, L., Wang, X.: Chiron-o1: Igniting multimodal large language models towards generalizable medical reasoning via mentor-intern collaborative search. *arXiv preprint arXiv:2506.16962* (2025)
 39. Thieme, A.H., Zheng, Y., Machiraju, G., Sadee, C., Mittermaier, M., Gertler, M., Salinas, J.L., Srinivasan, K., Gyawali, P., Carrillo-Perez, F., et al.: A deep-learning algorithm to classify skin lesions from mpox virus infection. *Nature medicine* **29**(3), 738–747 (2023)
 40. Vallée, R., De Maissin, A., Coutrot, A., Mouchère, H., Bourreille, A., Normand, N.: Crohnipi: An endoscopic image database for the evaluation of automatic crohn’s disease lesions recognition algorithms. In: *Medical Imaging 2020: Biomedical Applications in Molecular, Structural, and Functional Imaging*. vol. 11317, pp. 440–446. SPIE (2020)
 41. Vaswani, A., Shazeer, N., Parmar, N., Uszkoreit, J., Jones, L., Gomez, A.N., Kaiser, Ł., Polosukhin, I.: Attention is all you need. *Advances in neural information processing systems* **30** (2017)
 42. Wang, W., Ma, Z., Wang, Z., Wu, C., Ji, J., Chen, W., Li, X., Yuan, Y.: A survey of llm-based agents in medicine: How far are we from baymax? Findings of the Association for Computational Linguistics: ACL 2025 pp. 10345–10359 (2025)
 43. Wang, Z., Luo, X., Jiang, X., Li, D., Qiu, L.: Llm-radjudge: Achieving radiologist-level evaluation for x-ray report generation. *arXiv preprint arXiv:2404.00998* (2024)
 44. Xiao, C., Xu, T., Ma, S., Jiang, Y., Gao, H., Wu, Y.: Reversible primitive-composition alignment for continual vision-language learning. In: *The Fourteenth International Conference on Learning Representations* (2026)
 45. Xu, M., Liu, H., Peng, N., Chen, Q., Xiao, C.: Affordance-first decomposition for continual learning in video-language understanding. In: *Proceedings of the*

- IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR). pp. 3908–3919 (June 2026)
46. Xu, W., Chan, H.P., Li, L., Aljunied, M., Yuan, R., Wang, J., Xiao, C., Chen, G., Liu, C., Li, Z., et al.: Lingshu: A generalist foundation model for unified multimodal medical understanding and reasoning. arXiv preprint arXiv:2506.07044 (2025)
 47. Yanase, J., Triantaphyllou, E.: A systematic survey of computer-aided diagnosis in medicine: Past and present developments. *Expert Systems with Applications* **138**, 112821 (2019)
 48. Yokote, A., Umeno, J., Kawasaki, K., Fujioka, S., Fuyuno, Y., Matsuno, Y., Yoshida, Y., Imazu, N., Miyazono, S., Moriyama, T., et al.: Small bowel capsule endoscopy examination and open access database with artificial intelligence: the see-artificial intelligence project. *DEN open* **4**(1), e258 (2024)
 49. Zhang, J., Zhao, C., Xiao, C., Duan, R., Mo, W., Gao, H., Wang, W.: Pi-cca: Prompt-invariant cca certificates for replay-free continual multimodal learning. In: The Fourteenth International Conference on Learning Representations (2026)
 50. Zhang, K., Zhou, R., Adhikarla, E., Yan, Z., Liu, Y., Yu, J., Liu, Z., Chen, X., Davison, B.D., Ren, H., et al.: A generalist vision-language foundation model for diverse biomedical tasks. *Nature medicine* **30**(11), 3129–3141 (2024)
 51. Zhao, W., Wu, C., Fan, Y., Qiu, P., Zhang, X., Sun, Y., Zhou, X., Zhang, S., Peng, Y., Wang, Y., et al.: An agentic system for rare disease diagnosis with traceable reasoning. *Nature* pp. 1–10 (2026)
 52. Zhou, H., Tang, J., Zhang, J., Li, Y., Xiao, C., Hou, L., Ke, Z., Yao, J.: Comem: Compositional concept-graph memory for vision-language adaptation. In: The Fourteenth International Conference on Learning Representations (2026)
 53. Zhou, S., Wang, S., Yuan, Z., Shi, M., Shang, Y., Yang, D.: Gsq-tuning: Group-shared exponents integer in fully quantized training for llms on-device fine-tuning. In: Findings of the Association for Computational Linguistics: ACL 2025. pp. 22971–22988 (2025)
 54. Zhou, Y., Li, X., Wang, Q., Shen, J.: Visual in-context learning for large vision-language models. In: Findings of the Association for Computational Linguistics: ACL 2024. pp. 15890–15902 (2024)
 55. Zhou, Y., Song, L., Shen, J.: Improving medical large vision-language models with abnormal-aware feedback. In: Proceedings of the 63rd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers). pp. 12994–13011 (2025)
 56. Zhou, Y., Song, L., Shen, J.: Mam: Modular multi-agent framework for multi-modal medical diagnosis via role-specialized collaboration. In: Findings of the Association for Computational Linguistics: ACL 2025. pp. 25319–25333 (2025)
 57. Zhou, Y., Zheng, H., Chen, D., Yang, H., Han, W., Shen, J.: From medical llms to versatile medical agents: A comprehensive survey (2026)