

Pairing-free Group-level Knowledge Distillation for Robust Gastrointestinal Lesion Classification in White-Light Endoscopy

Qiang Hu*, Qimei Wang*, Yingjie Guo, Qiang Li, Zhiwei Wang†

Wuhan National Laboratory for Optoelectronics, Huazhong University of Science and Technology
{huqiang77, zwwang}@hust.edu.cn

Abstract

White-Light Imaging (WLI) is the standard for endoscopic cancer screening, but Narrow-Band Imaging (NBI) offers superior diagnostic details. A key challenge is transferring knowledge from NBI to enhance WLI-only models, yet existing methods are critically hampered by their reliance on paired NBI-WLI images of the same lesion, a costly and often impractical requirement that leaves vast amounts of clinical data untapped. In this paper, we break this paradigm by introducing PaGKD, a novel Pairing-free Group-level Knowledge Distillation framework that enables effective cross-modal learning using unpaired WLI and NBI data. Instead of forcing alignment between individual, often semantically mismatched image instances, PaGKD operates at the *group level* to distill more complete and compatible knowledge across modalities. Central to PaGKD are two complementary modules: (1) Group-level Prototype Distillation (GKD-Pro) distills compact group representations by extracting modality-invariant semantic prototypes via shared lesion-aware queries; (2) Group-level Dense Distillation (GKD-Den) performs dense cross-modal alignment by guiding group-aware attention with activation-derived relation maps. Together, these modules enforce global semantic consistency and local structural coherence without requiring image-level correspondence. Extensive experiments on four clinical datasets demonstrate that PaGKD consistently and significantly outperforms state-of-the-art methods, achieving relative AUC improvements of 3.3%, 1.1%, 2.8%, and 3.2%, respectively, establishing a new direction for cross-modal learning from unpaired data.

Code — <https://github.com/Huster-Hq/PaGKD>

Introduction

Endoscopy is a key tool for early detection of gastrointestinal (GI) cancers (Schmitz et al. 2022). Accurate lesion classification from endoscopic images is critical for timely diagnosis and treatment planning; however, manual evaluation is inherently subjective and labor-extensive, driving the development of automated diagnostic systems. Modern endoscopy relies on two imaging modalities: White-Light

Imaging (WLI) and Narrow-Band Imaging (NBI). Among them, NBI enjoys enhanced vascular and mucosal details via spectral filtering, yielding superior lesion visibility and thereby improving classification performance (Tamura et al. 2022; Yue et al. 2023). However, WLI remains the dominant modality in real-world clinical practice due to its universal availability and default configuration in most endoscopic systems (Soffer et al. 2020; Beg et al. 2020), whereas NBI functionality is either absent or underutilized in routine deployments.

This discrepancy raises a pivotal research question: *How can we transfer the rich visual knowledge embedded in NBI to enhance lesion recognition in WLI, particularly when only WLI data is accessible during clinical deployment?* This challenge falls under the broader umbrella of cross-modal knowledge transfer. Emerging studies have initiated investigations into this direction by leveraging NBI as a privileged modality during training to augment WLI-only test-time classification. Representative approaches, such as PolypsAlign (Ma et al. 2022) and CPC-Trans (Wang et al. 2021), align modality-specific features using CNN or Transformer architectures guided by lesion priors. More recently, ADD (Hu et al. 2025) relaxes the requirement for spatial alignment priors and performs holistic image-level alignment on paired yet unregistered images, achieving notable advancements in WLI-only classification performance.

Despite their efficacy, existing cross-modal distillation methods are limited by their reliance on paired NBI-WLI images of the same lesion. This *instance-specific distillation* paradigm assumes that knowledge can be directly transferred between different modality views of the same physical object. While such paired data is informative, collecting it at scale is costly and impractical. More importantly, most real-world clinical datasets are composed of unpaired NBI and WLI images acquired independently from different lesions and patients, leaving a vast amount of data untapped and severely limiting the scalability of current approaches.

In this work, we address a critical yet underexplored challenge: *cross-instance knowledge distillation* for improving lesion classification in WLI endoscopy. Our goal is to leverage abundant *unpaired* clinical data to transfer NBI-derived knowledge into WLI-only classification model. This is both feasible and clinically meaningful, as many modality-specific patterns in NBI, such as vascular and mu-

*These authors contributed equally.

†Corresponding author.

Copyright © 2026, Association for the Advancement of Artificial Intelligence (www.aaai.org). All rights reserved.

cosal structures, exhibit high recurrence across patients and lesions. However, a major bottleneck arises when applying traditional instance-level distillation to unpaired data: individual lesion images often provide only partial and case-specific manifestations of disease. This leads to semantic mismatches and feature incompatibility between the source (NBI) and target (WLI) models, ultimately undermining the effectiveness of cross-modal distillation.

To effectively bridge the modality gap between WLI and NBI images without requiring one-to-one pairing, we propose PaGKD, a **P**airing-free **G**roup-level **K**nowledge **D**istillation framework. Unlike prior methods like ADD (Hu et al. 2025), which align features between individual images, PaGKD operates at the group level by aggregating images with the same lesion class into semantically coherent feature sets. This group-level formulation mitigates sample-level noise and enhances the stability and completeness of cross-modal knowledge transfer. Specifically, PaGKD is built upon two complementary modules: Group-level Prototype Distillation (GKD-Pro) and Group-level Dense Distillation (GKD-Den). GKD-Pro distills high-level, modality-invariant semantics by extracting class-specific prototypes through shared lesion-aware queries. These prototypes act as global semantic anchors and are aligned across modalities via a contrastive objective, enabling the model to abstract consistent concepts from unpaired groups. GKD-Den, in contrast, distills local structural knowledge by generating relation maps from refined activation cues, which capture spatial correspondences between modality-specific groups. These maps supervise a group-aware cross-attention mechanism that restructures features across modalities, enforcing alignment at a dense, region-wise level. Together, GKD-Pro and GKD-Den form a mutually reinforcing mechanism, where GKD-Pro ensures semantic consistency and GKD-Den enforces spatial coherence. This dual-level alignment enables robust knowledge distillation across modalities, even in the absence of image-level pairing.

In summary, our major contributions are as follows.

- We propose PaGKD, a novel pairing-free group-level knowledge distillation framework that enables effective cross-modal learning between WLI and NBI images without requiring one-to-one correspondence.
- We design two complementary modules: GKD-Pro for modality-invariant semantic prototype learning, and GKD-Den for fine-grained spatial alignment. Together, they achieve coherent, class-consistent feature alignment across WLI and NBI modalities.
- We conduct extensive experiments on three colorectal polyp classification datasets and one gastric cancer dataset. PaGKD consistently outperforms state-of-the-art methods, achieving 3.3%, 1.1%, 2.8%, and 3.2% relative improvements in AUC, respectively.

Related Works

Cross-modality Independent Classification (CIC) for GI Diseases

Significant progress has been made in gastrointestinal (GI) disease classification. Most existing methods treat multi-

modal inputs (e.g., WLI and NBI) as modality-agnostic, ignoring potential inter-modal knowledge interaction. Some works focus on architectural improvements, such as integrating multi-level features to enhance classification (Zhang et al. 2016), employing Fourier transforms to mitigate illumination-related noise (Wang et al. 2022). Other studies emphasize feature learning via self-supervised strategies: SSL-CPCD (Xu, Rittscher, and Ali 2024) introduces an additive angular margin within cosine similarity to guide patch-image clustering, and SSL-WCE (Guo and Yuan 2020) adopts adaptive aggregated attention to extract discriminative features. Additionally, some methods focus on learning local lesion-aware features, such as employing attention mechanisms to integrate local and global representations for classification (Wang et al. 2023). However, these approaches often perform suboptimally on WLI images, primarily due to the lack of effective utilization of the richer diagnostic information in NBI, which limits their classification performance in modality-limited scenarios.

Cross-modal Distillation Classification (CDC) for GI Diseases

Cross-modal knowledge distillation is an effective strategy for transferring knowledge from an informative source modality (NBI) to a less informative target modality (WLI) to improve disease classification. Existing CDC methods typically relying on paired images for image-level knowledge transfer. For example, Polypsalign (Wang et al. 2021) adopted discriminator and contrastive loss for global feature alignment; CPC-Trans (Ma et al. 2022) employed a Transformer-based architecture with cross-attention to align patch-level and global features; and ADD (Hu et al. 2025) further performed pixel-wise image-to-image distillation guided by semantic correspondences. However, NBI functionality is either absent or underutilized in routine deployments, leading to large-scale unpaired data. These paired-dependent CDC methods exhibit limited generalizability in such settings and fail to fully exploit the unpaired information. To bridge the modality gap between WLI and NBI without requiring one-to-one pairing, we propose group-level local spatial feature distillation and global prototype distillation, enabling effective and stable unpaired knowledge transfer in challenging GI scenarios.

Method

Overview

As illustrated in Figure 1, our framework PaGKD consists of a trainable WLI classifier and a frozen, pretrained NBI classifier. Both classifiers share the same architecture, following ADD (Hu et al. 2025). To facilitate knowledge transfer between *unpaired* WLI and NBI images, we introduce a group-level feature distillation strategy, instead of regular image-level approaches. In each training iteration, we construct image groups $\{\mathcal{G}_c^{\text{mod}}\}$, where $\mathcal{G}_c^{\text{mod}}$ contains N_c images from class c in modality $\text{mod} \in \{\text{WLI}, \text{NBI}\}$, and H, W denote the input dimensions. Each image within a group is fed independently into its corresponding classifier to extract fea-

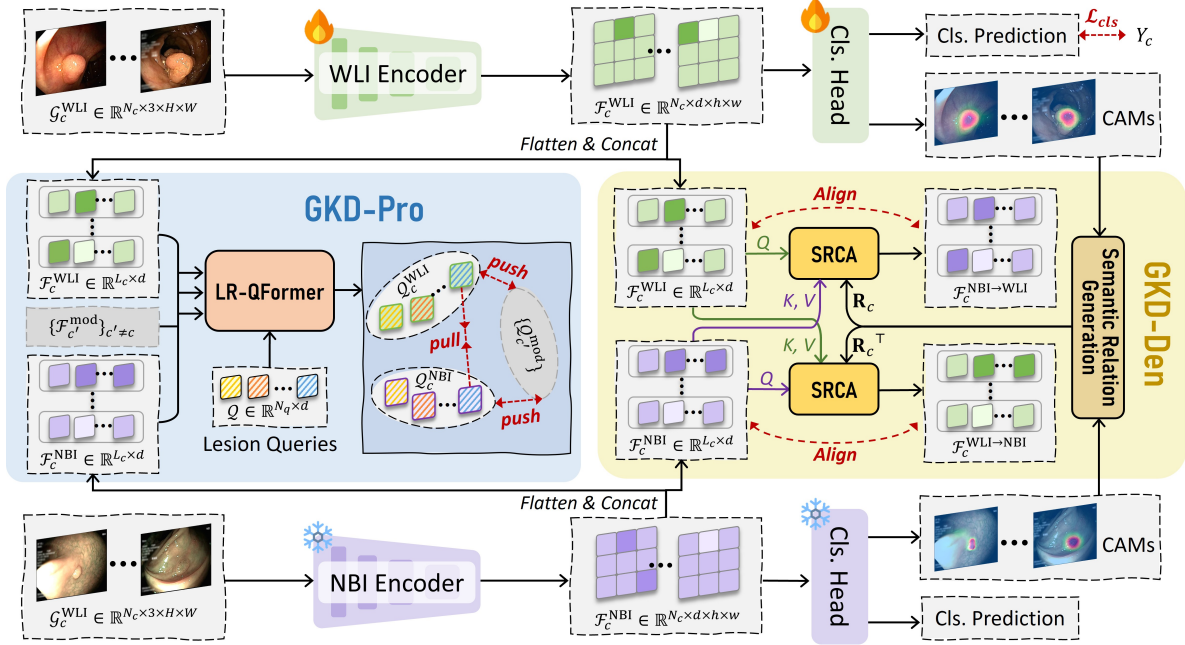


Figure 1: Overview of the proposed PaGKD. It consists of a pre-trained and frozen NBI classifier, a trainable WLI classifier, and two group-level knowledge distillation modules: Group-level Prototype knowledge distillation (GKD-Pro) and Group-level Dense Knowledge Distillation (GKD-Den). During training, given two unpaired image groups from the same class, GKD-Pro collaborates with GKD-Den to achieve robust and multi-granularity cross-modal knowledge distillation between them.

ture maps:

$$\mathcal{G}_c^{\text{mod}} \rightarrow \mathcal{F}_c^{\text{mod}} \in \mathbb{R}^{N_c \times d \times h \times w}, \quad (1)$$

where d , h , and w are the channel, height, and width of the feature maps, respectively.

For group-level operations, we flatten and aggregate all spatial features across the group into a unified representation:

$$\mathcal{F}_c^{\text{mod}} = \{\mathbf{f}_{i,c}^{\text{mod}} \in \mathbb{R}^d\}_{i=1}^{L_c}, \quad \text{with } L_c = N_c \cdot h \cdot w. \quad (2)$$

We then apply two complementary modules for multi-granularity distillation: Group-level Prototype Knowledge Distillation (GKD-Pro) performs alignment of global class-level distributions, while Group-level Dense Knowledge Distillation (GKD-Den) aligns local feature.

Group-level Prototype Knowledge Distillation

To facilitate effective cross-modal knowledge transfer from unpaired data, we first group features that share similar lesion traits. While these group-level features offer richer semantics than individual samples, they may still contain redundant or irrelevant patterns. To distill core lesion representations, we extract feature prototypes that summarize the underlying disease characteristics shared across each group.

Lesion-related Query Transformer for Prototype Representation. We propose a learnable prototype extraction mechanism based on the Lesion-Related Query Transformer (LR-QFormer). It uses a set of shared trainable lesion queries $Q \in \mathbb{R}^{N_q \times d}$ to actively attend to discrimina-

tive group features, where $N_q \ll L_c$. Unlike static clustering, this query-based approach enables adaptive and interpretable abstraction of disease semantics. The queries Q are shared across all groups and modalities, and are designed to be class- and modality-agnostic. They act as latent anchors representing modality-invariant lesion concepts. During inference, their interaction with modality-specific group features reveals how these lesion attributes are expressed in different imaging conditions.

The LR-QFormer is composed of T stacked Transformer blocks. At each layer, the queries are first refined via self-attention SA, then updated by cross-attending CA to the group feature map $\mathcal{F}_c^{\text{mod}} \in \mathbb{R}^{L_c \times d}$ with absolute positional encoding \mathbf{E}_{pos} (Dosovitskiy et al. 2020). Formally:

$$\mathcal{Q}_{t,c}^{\text{mod}} = \text{CA} \left(\text{SA} \left(\mathcal{Q}_{t-1,c}^{\text{mod}} \right), \mathcal{F}_c^{\text{mod}} + \mathbf{E}_{\text{pos}} \right), \quad \mathcal{Q}_{0,c}^{\text{mod}} = \mathcal{Q}. \quad (3)$$

Through this iterative process, each query accumulates modality-specific evidence of its corresponding disease attribute. After T layers, we obtain the refined set $\mathcal{Q}_c^{\text{mod}} = \{\mathbf{q}_{i,c}^{\text{mod}} \in \mathbb{R}^d\}_{i=1}^{N_q}$ which serves as the prototype representation of the group in modality mod . These prototypes are both compact and interpretable, enabling robust cross-modal alignment in subsequent distillation.

Group-level Contrastive Loss for Knowledge Distillation.

To facilitate cross-modal knowledge distillation, we introduce a group-level contrastive objective over the refined lesion query sets $\{\mathcal{Q}_c^{\text{mod}}\}$ generated for each class c and modality $\text{mod} \in \{\text{WLI}, \text{NBI}\}$. The core motivation is that

the same set of lesion queries \mathcal{Q} , shared across groups, captures modality-invariant disease semantics. When interacting with group features from different classes and modalities, these queries produce class-specific prototype representations, which reflect how a disease manifests under a particular imaging modality.

Our goal is to align these group-level representations across modalities for the same class, while separating them across different classes. This encourages the WLI classifier to focus on pathology-relevant patterns distilled from NBI, avoiding modality-specific confounders like color or illumination variations. To this end, we define the similarity between two query sets $\mathcal{Q}_c^{\text{mod}}$ and $\mathcal{Q}_{c'}^{\text{mod}'}$ as the average cosine similarity between queries at the same index:

$$S_{\mathcal{Q}_c^{\text{mod}}, \mathcal{Q}_{c'}^{\text{mod}'}} = \frac{1}{N_q} \sum_{i=1}^{N_q} \frac{(\mathbf{q}_{i,c}^{\text{mod}})^\top \mathbf{q}_{i,c'}^{\text{mod}'}}{\|\mathbf{q}_{i,c}^{\text{mod}}\| \|\mathbf{q}_{i,c'}^{\text{mod}'}\|}. \quad (4)$$

This formulation assumes each query captures a specific semantic aspect of disease (e.g., boundary irregularity, vascular pattern), enabling a fine-grained alignment across groups. We then construct a symmetric contrastive loss that encourages high similarity between WLI-NBI pairs of the same class and low similarity between different classes:

$$\mathcal{L}_{pro} = -\frac{1}{2C} \sum_{c=1}^C \left(\log \frac{\exp(S_{\mathcal{Q}_c^{\text{WLI}}, \mathcal{Q}_c^{\text{NBI}}})}{\sum_{\mathcal{Q}_{c'}^{\text{mod}} \neq \mathcal{Q}_c^{\text{WLI}}} \exp(S_{\mathcal{Q}_c^{\text{WLI}}, \mathcal{Q}_{c'}^{\text{mod}}})} + \log \frac{\exp(S_{\mathcal{Q}_c^{\text{NBI}}, \mathcal{Q}_c^{\text{WLI}}})}{\sum_{\mathcal{Q}_{c'}^{\text{mod}} \neq \mathcal{Q}_c^{\text{NBI}}} \exp(S_{\mathcal{Q}_c^{\text{NBI}}, \mathcal{Q}_{c'}^{\text{mod}}})} \right). \quad (5)$$

This loss term guides the lesion queries to distill disease-centric semantics shared across modalities, thus enabling group-level alignment and effective cross-modal transfer, even under unpaired supervision.

Group-level Dense Knowledge Distillation

While the above-introduced lesion-related queries capture high-level, modality-invariant attributes, they may overlook fine spatial details critical for accurate diagnosis. To complement this, we perform dense alignment between unpaired cross-modal feature groups. In medical imaging, subtle local patterns such as vascular distortion or mucosal texture are often decisive and require pixel-level modeling.

Our approach draws inspiration from ADD (Hu et al. 2025), which aligns images based on Class Activation Maps (CAMs) (Zhou et al. 2016). While ADD operates on individual images, we extend this idea to the group level by aligning all WLI and NBI features within the same class. This enables more stable and comprehensive cross-modal supervision.

Semantic Relation Generation for Inter-class Dense Connection. For each group $\mathcal{G}_c^{\text{mod}}$ in class c , we first compute CAMs from the extracted features, then apply pixel-adaptive refinement (Ru et al. 2022) to enhance spatial coherence. The refined CAMs are flattened into a sequence $\mathcal{M}_c^{\text{mod}} = \{\mathbf{m}_{i,c}^{\text{mod}}\}_{i=1}^{L_c}$, and discretized using a tri-thresholding scheme with thresholds τ_1 and τ_2 . Each element $\mathbf{m}_{i,c}^{\text{mod}} \in \{0, 1, \emptyset\}$ indicates background, lesion, or ambiguity.

To construct a semantic relation matrix $\mathbf{R}_c \in \{0, 1\}^{L_c \times L_c}$ between WLI and NBI groups, we mark a pair of positions (p, q) as semantically aligned (i.e., $r_{(p,q),c} = 1$) if both are confidently labeled (non-ambiguous) and share the same semantic label in their respective binarized CAMs:

$$\mathbf{R}_c(p, q) = \begin{cases} 0 & \text{if } \mathbf{m}_{p,c}^{\text{WLI}} = \mathbf{m}_{q,c}^{\text{NBI}} \text{ and } \mathbf{m}_{p,c}^{\text{WLI}}, \mathbf{m}_{q,c}^{\text{NBI}} \neq \emptyset, \\ -\infty & \text{otherwise.} \end{cases} \quad (6)$$

This relation matrix serves as dense supervision for pixel-level cross-modal distillation. It enables the model to transfer lesion-related knowledge from NBI to WLI at a finer granularity, even in the presence of spatial misalignment or modality-specific artifacts.

Semantic Relation-guided Cross-Attention (SRCA) for Dense Distillation. Given the semantic relation matrix \mathbf{R}_c , we introduce a Semantic Relation-guided Cross-Attention (SRCA) module to align dense local features across unpaired cross-modal groups. This module uses relation-aware attention to reshape spatial structures and enhance semantic consistency during feature interaction. Formally, given group features $\mathcal{F}_c^{\text{WLI}}$ and $\mathcal{F}_c^{\text{NBI}}$, we reconstruct a spatially ‘aligned’ version of NBI in the WLI coordinate space via:

$$\mathcal{F}_c^{\text{NBI} \rightarrow \text{WLI}} = \mathbf{A}(\mathcal{F}_c^{\text{NBI}} W_v) \in \mathbb{R}^{L_c \times d}, \quad (7)$$

$$\mathbf{A} = \text{Softmax} \left(\mathbf{R}_c + \frac{(\mathcal{F}_c^{\text{WLI}} W_q)(\mathcal{F}_c^{\text{NBI}} W_k)^\top}{\sqrt{d/4}} \right),$$

where $W_q, W_k \in \mathbb{R}^{d \times \frac{d}{4}}$ and $W_v \in \mathbb{R}^{d \times d}$ are linear projections. The attention matrix $\mathbf{A} \in \mathbb{R}^{L_c \times L_c}$ guides the spatial reconstruction by aggregating semantically matched positions from NBI features, reshaping them to better align with WLI semantics. A symmetric operation is performed to reconstruct WLI features in the NBI space, i.e., $\mathcal{F}_c^{\text{WLI} \rightarrow \text{NBI}}$, using the transposed relation matrix \mathbf{R}_c^\top . This bidirectional mapping supports dense cross-modal feature fusion under weakly supervised and unpaired conditions.

Compared with ADD (Hu et al. 2025), which aligns single images using CAM-derived correspondences, our method performs attention-guided alignment at the group level, allowing the model to query from richer spatial and semantic contexts. Importantly, our spatial reshaping aggregates evidence across groups, enabling more reliable reconstruction even when individual samples lack complete lesion information. With the two sets of group-level ‘reconstructed’ features $\mathcal{F}_c^{\text{NBI} \rightarrow \text{WLI}} = \{\mathbf{f}_{i,c}^{\text{NBI} \rightarrow \text{WLI}}\}_{i=1}^{L_c}$ and $\mathcal{F}_c^{\text{WLI} \rightarrow \text{NBI}} = \{\mathbf{f}_{i,c}^{\text{WLI} \rightarrow \text{NBI}}\}_{i=1}^{L_c}$ obtained, we define a bidirectional consistency loss to supervise this alignment as formulated in Eq. (8). This loss supervises dense distillation at the group level in both directions and encourages fine-grained cross-modal consistency.

$$\mathcal{L}_{den} = \frac{1}{C(L_c)^2} \sum_{c=1}^C \left(\sum_{p=1}^{L_c} \|\mathbf{f}_{p,c}^{\text{WLI}} - \mathbf{f}_{p,c}^{\text{NBI} \rightarrow \text{WLI}}\|_2 + \sum_{q=1}^{L_c} \|\mathbf{f}_{q,c}^{\text{NBI}} - \mathbf{f}_{q,c}^{\text{WLI} \rightarrow \text{NBI}}\|_2 \right). \quad (8)$$

Type	Dataset	Class	WLI-NBI	Unpaired	
			Pairs	WLI	NBI
Colorectal Polyp	PolypSet	Hyperplasia	63	200	200
		Adenoma	102	250	250
	PICCOLO	Hyperplasia	249	155	35
		Adenoma	620	834	162
		Adenocarcinoma	186	76	17
	IH-Polyp	Hyperplasia	257	1,250	427
Adenoma		299	2,480	494	
GC	IH-GC	Gastric Inflammation	120	215	159
		Gastric Cancer	144	254	144

Table 1: Statistics of experimental datasets.

Network Training and Inference

In the training stage, in addition to GKD-Pro and GKD-Den distillations, i.e., \mathcal{L}_{pro} and \mathcal{L}_{den} , our overall training objective also includes a cross-entropy loss term \mathcal{L}_{cls} for the classification results of the WLI classifier for the ground truth class labels. The total training loss \mathcal{L}_{total} is formulated as:

$$\mathcal{L}_{total} = \mathcal{L}_{pro} + \mathcal{L}_{den} + \mathcal{L}_{cls}, \quad (9)$$

where $\mathcal{L}_{cls} = \frac{1}{C} \sum_{c=1}^C \left(\frac{1}{N_c} \sum_{i=1}^{N_c} \text{CE}(P_{i,c}^{\text{WLI}}, Y_c) \right)$. Here, $\text{CE}(\cdot, \cdot)$ is the cross-entropy loss function, $P_{i,c}^{\text{WLI}}$ is the classification result of the WLI classifier for the i -th image in the WLI image group of class c , and Y_c is the one-hot label corresponding to the class c . In the inference stage, we only keep the WLI classifier, allowing it to receive a single WLI image and output the classification result.

Experiments

Datasets and Evaluation Metrics

Datasets. We conduct experiments on four gastrointestinal (GI) cancer datasets independently, including three colorectal polyp datasets: PICCOLO (Sánchez-Peralta et al. 2020), PolypSet (Li et al. 2021), and an in-house dataset, denoted IH-Polyp, and an in-house gastric cancer (GC) dataset, denoted IH-GC. Detailed statistics are summarized in Table 1.

Evaluation Metrics. To fairly compare the performance of the NBI and WLI classifiers, we perform 5-fold cross-validation on the paired data in all datasets, while the unpaired data are only used as the train set. We ensure that the patient-level separation between the train and test sets. To evaluate performance, we use four metrics: Accuracy (Acc), Precision (Pre), Recall (Rec), F1-score (F1), and Area Under the Curve (AUC). Unless otherwise specified, all evaluations are conducted on the WLI test set.

Implementation Details

PaGKD is implemented using PyTorch (Paszke et al. 2019) and trained on a single NVIDIA 4090 GPU with 24GB memory. Weights of ResNet-50 (He et al. 2016) pre-trained on ImageNet (Deng et al. 2009) are loaded as initialization of the WLI classifier. NBI classifier is pretrained on NBI train set and frozen during distillation. The input size is set to 448×448 , and the batch size is set to 24. We define the

group size for each class within a batch based on the class distribution in the dataset. The number of epochs is set to 100 and the initial learning rate is set to $1e^{-4}$. The optimizer is set to Adam (Kingma and Ba 2014) with $1e^{-8}$ weight decay. We set the number of lesion queries in LR-QFormer N as 12, and set τ_1 and τ_2 as 0.3 and 0.7, respectively, according to the parameter search experiments. Note that PaGKD can also be seamlessly extended to paired data and leverage the inherent prior pairing relationships by simply transforming the input features of GKD-Pro and GKD-Den from group-level to image-level.

Comparison with State-of-the-arts (SOTAs)

We compare our PaGKD with eight SOTA methods on four GI cancer datasets, including four cross-modal independent classification (CIC) methods: SSL-CPCD (Xu, Rittscher, and Ali 2024), SSL-WCE (Guo and Yuan 2020), FFC-Net (Wang et al. 2022), and DLGNet (Wang et al. 2023), and four cross-modal distillation classification (CDC) methods: CPC-Trans (Ma et al. 2022), SAMD (Shen et al. 2023), PolypsAlign (Wang et al. 2021), and ADD (Hu et al. 2025). For fairness, all methods adopt ResNet-50 (He et al. 2016) as backbone, except CPC-Trans using ViT-S (Dosovitskiy et al. 2020), which relies on the Transformer architecture. Since all CIC methods supervise each image independently without relying on paired data, they are trained with paired and unpaired data to maximize the train data. The CDC methods are inherently designed for paired data. For comprehensive comparison, we train them in two settings: only paired data and a combination of paired and unpaired data.

The results are presented in Table 2, from which we derive three key observations: (1) The CDC methods outperform CIC methods in terms of overall trends on four datasets, which proves the efficacy of cross-modal knowledge distillation from the NBI modality to the WLI modality in enhancing WLI GI cancer classification. (2) The introduction of additional unpaired data not only fails to enhance the performance of existing CDC methods but also leads to performance degradation. This confirms that applying these methods, which are originally designed for paired images of the same lesion instance, to unpaired images from different lesions will introduce noise and hinder effective feature learning. (3) When trained on the combination of paired and unpaired data, PaGKD significantly outperforms other SOTAs across all four datasets, achieving relative AUC improvements of at least 3.3%, 1.1%, 2.8%, and 3.2%, respectively, and achieves performance closest to the NBI classifier. Furthermore, even when using only unpaired data, PaGKD achieves performance that is competitive with or even better than the best competing methods. These results demonstrate the effectiveness of PaGKD in utilizing unpaired data.

Ablation Studies

Effectiveness of Key Components. To verify the effectiveness of the two key components of PaGKD: Group-level Prototype Knowledge Distillation (GKD-Pro) and Group-level Dense Knowledge Distillation (GKD-Den), we train three variants of PaGKD by disabling GKD-Pro or/and GKD-Den on the combination of paired and unpaired

Methods	Train Data	PICCOLO (3 Classes)					PolypSet (2 Classes)					IH-Polyp (2 Classes)					IH-GC (2 Classes)				
		Acc	Pre	Rec	F1	AUC	Acc	Pre	Rec	F1	AUC	Acc	Rec	Pre	F1	AUC	Acc	Rec	Pre	F1	AUC
NBI Classifier		80.7	79.1	79.6	79.1	86.9	94.8	96.0	95.1	95.5	97.6	83.8	84.3	89.4	86.8	87.0	81.9	76.7	89.7	83.5	86.3
► <i>Cross-modal Independent Classification (CIC) methods:</i>																					
SSL-CPCD	$\mathcal{D}_p + \mathcal{D}_{unp}$	64.1	65.4	51.4	54.0	72.5	76.7	77.7	85.3	81.3	80.2	64.9	70.1	71.2	70.6	68.9	66.7	74.1	75.5	74.8	70.4
SSL-WCE	$\mathcal{D}_p + \mathcal{D}_{unp}$	66.0	67.8	57.5	59.4	73.9	76.2	76.5	86.3	81.1	81.2	56.8	57.1	90.3	70.0	57.5	64.9	73.9	70.8	72.3	69.6
FFCNet	$\mathcal{D}_p + \mathcal{D}_{unp}$	76.8	80.2	68.3	72.1	80.1	87.2	90.8	87.3	89.0	91.4	64.0	66.7	78.8	72.2	62.9	70.4	<u>89.2</u>	62.3	73.4	70.8
DLGNet [†]	$\mathcal{D}_p + \mathcal{D}_{unp}$	68.2	67.8	57.5	59.4	74.5	80.8	80.5	89.2	84.6	83.4	67.6	74.2	69.7	71.9	63.3	70.4	78.4	75.5	76.9	66.0
► <i>Cross-modal Distillation Classification (CDC) methods:</i>																					
CPC-Trans [†]	\mathcal{D}_p	78.6	76.7	76.2	76.4	86.6	90.1	89.0	95.1	91.9	87.2	69.4	67.4	93.9	78.5	75.6	75.6	73.2	90.6	79.4	75.5
CPC-Trans [†]	$\mathcal{D}_p + \mathcal{D}_{unp}$	78.2	76.7	76.5	76.6	<u>87.2</u>	89.9	88.3	94.7	91.0	86.5	62.2	72.2	59.1	65.0	69.6	67.9	73.3	68.8	71.0	69.8
SAMD	\mathcal{D}_p	75.3	77.0	70.0	72.8	83.1	89.0	89.5	92.2	90.8	90.0	70.3	70.0	74.6	75.8	75.2	73.6	76.7	71.9	74.2	76.6
SAMD	$\mathcal{D}_p + \mathcal{D}_{unp}$	74.3	74.7	67.9	70.1	79.3	87.2	89.2	89.2	89.2	90.1	62.2	76.1	53.0	62.5	70.8	69.6	89.5	53.1	66.7	77.3
PolypsAlign [†]	\mathcal{D}_p	79.4	80.4	74.1	76.3	84.7	91.3	93.1	92.2	92.6	90.6	70.3	71.4	83.3	76.9	74.3	73.2	79.3	71.9	75.4	81.4
PolypsAlign [†]	$\mathcal{D}_p + \mathcal{D}_{unp}$	76.2	88.8	58.8	63.9	81.5	86.0	91.5	84.3	87.8	85.7	60.4	67.2	65.2	66.2	60.6	71.4	75.0	75.0	75.0	73.4
ADD	\mathcal{D}_p	78.1	79.1	74.4	76.4	83.8	93.4	96.1	92.8	94.0	93.7	80.2	76.8	<u>90.5</u>	<u>84.1</u>	<u>82.8</u>	75.0	82.1	71.9	79.2	77.3
ADD	$\mathcal{D}_p + \mathcal{D}_{unp}$	76.8	77.9	71.7	74.2	84.2	91.9	90.7	96.1	93.3	93.6	74.8	78.8	78.8	78.8	75.7	75.0	76.5	81.2	78.8	77.8
PaGKD (Ours)	\mathcal{D}_{unp}	<u>80.8</u>	80.8	<u>77.2</u>	<u>78.8</u>	86.6	<u>94.2</u>	91.8	99.0	<u>95.3</u>	94.0	79.3	77.9	83.9	83.9	82.4	<u>79.4</u>	82.5	<u>88.0</u>	<u>80.5</u>	<u>82.2</u>
PaGKD (Ours)	\mathcal{D}_p	78.6	79.2	73.3	75.3	84.4	93.6	93.3	96.1	94.7	<u>94.2</u>	<u>81.1</u>	<u>83.6</u>	83.0	83.2	81.1	76.8	82.8	75.0	78.7	81.0
PaGKD (Ours)	$\mathcal{D}_p + \mathcal{D}_{unp}$	81.9	<u>81.9</u>	80.6	81.1	90.1	94.8	<u>94.3</u>	<u>97.1</u>	95.7	94.7	82.0	84.8	84.8	84.8	85.1	81.5	86.5	84.9	85.7	84.0

Table 2: Comparisons on three polyp classification datasets: PICCOLO, PolypSet, and IH-Polyp, and a gastric cancer dataset IH-GC. ‘NBI Classifier’ denotes the performance of the pre-trained NBI classifier on the NBI test set. The gray lines indicate that the methods originally designed for paired data receive additional unpaired train data. ‘†’ represents that the method requires additional prior information about the location of polyps during training and inference.

Components		PICCOLO			IH-GC		
GKD-Pro	GKD-Den	Acc	F1	AUC	Acc	F1	AUC
✗	✗	62.8	52.4	71.2	62.5	61.8	66.9
✓	✗	77.7	75.3	83.5	75.0	80.5	75.5
✗	✓	77.3	75.2	85.0	75.0	75.0	77.3
✓	✓	81.9	81.1	90.1	81.5	85.7	84.0

Table 3: Ablation study on the effectiveness of the two key components: GKD-Pro and GKD-Den.

data of both PICCOLO and IH-GC. When both GKD-Pro and GKD-Den are disabled, we actually directly train a ResNet-50 classifier in a cross-modal independent manner. As shown in Table 3, introducing GKD-Pro or GKD-Den individually produces notable performance improvements, demonstrating the effectiveness of each component in isolation. When both components are employed jointly, the model achieves further performance gains. This suggests that their contributions are complementary rather than redundant, enabling effective cross-modal knowledge distillation at both global-level and local-level.

Effectiveness of Key Sub-components. To evaluate the effectiveness of the key sub-components in GKD-Pro and GKD-Den, including LR-QFormer, SRCA and bidirectional consistency loss, we conduct ablation studies by individually disabling each component. Specifically, we ablate LR-QFormer by directly applying the group-level contrastive loss between image feature vectors obtained via average pooling, instead of between prototypes extracted by LR-QFormer. For SRCA, we replace it with the standard cross-attention mechanism (Vaswani et al. 2017) without leverag-

Methods	PICCOLO			IH-GC		
	ACC	F1	AUC	ACC	F1	AUC
PaGKD	81.9	81.1	90.1	81.5	85.7	84.0
► <i>GKD-Pro:</i>						
w/o LR-QFormer	77.0	74.8	85.7	76.8	76.3	78.3
► <i>GKD-Den:</i>						
w/o SRCA	76.6	74.7	83.2	75.6	80.0	76.8
w/o Bidirectional	78.8	78.0	87.8	79.4	83.8	81.9

Table 4: Ablation study on the effectiveness of a key sub-component in GKD-Pro: LR-QFormer, and two key sub-components in GKD-Den: SRCA and bidirectional consistency loss (denoted as ‘Bidirectional’).

ing the semantic relation map. To assess the effectiveness of the bidirectional consistency loss, we retain only the consistency loss term between $\mathbf{f}_{p,c}^{\text{WLI}}$ and $\mathbf{f}_{p,c}^{\text{NBI} \rightarrow \text{WLI}}$ in Eq. (8). As shown in Table 4, for GKD-Pro, removing LR-QFormer leads to a performance degradation. This is attributed to the fact that directly pooled features lack sufficient pathology-specific information, thereby weakening the effectiveness of contrastive learning for the model’s feature learning in lesion classification. For GKD-Den, ablating either SRCA or the bidirectional consistency loss results in varying degrees of performance degradation. The is because that removing SRCA compromises semantic consistency in the reconstruction of the spatially aligned feature, which undermines the reliability of the reconstructed features and weakens the subsequent dense distillation process. Similarly, using only the unidirectional consistency loss limits the completeness of knowledge transfer from NBI domain to WLI domain, thus reducing the effectiveness of dense knowledge distillation.

Level	PICCOLO			IH-GC		
	ACC	F1	AUC	ACC	F1	AUC
► <i>GKD-Pro</i> :						
Group-level	77.7	75.3	83.5	75.0	80.5	75.5
Image-level	64.0	64.9	78.6	67.9	71.9	69.0
► <i>GKD-Den</i> :						
Group-level	77.3	75.2	85.0	75.0	75.0	77.3
Image-level	73.9	69.7	79.9	69.6	75.4	71.9
► <i>Jointly (GKD-Pro+GKD-Den)</i> :						
Group-level	81.9	81.1	90.1	81.5	85.7	84.0
Image-level	75.3	72.1	84.3	73.2	75.4	75.7

Table 5: Comparisons between our proposed group-level distillation components and their image-level variants.

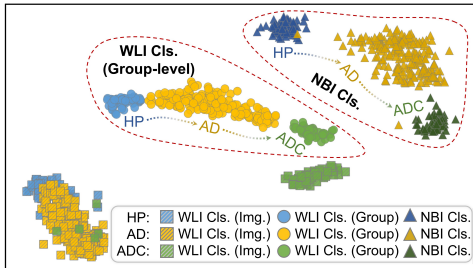


Figure 2: t-SNE visualization of the features on PICCOLO. HP (Hyperplasia) → AD (Adenoma) → ADC (Adenocarcinoma): the severity of the disease increases progressively.

Group-level vs. Image-level Knowledge Distillation. To verify the core idea of PaGKD, *i.e.*, group-level knowledge distillation, we modify GKD-Pro and GKD-Den into image-level variants by replacing the original group inputs with randomly sampled NBI-WLI pairs, and evaluate their performance both individually and jointly. The quantitative results are shown in Table 5. Image-level distillation leads to significant performance drops for both modules when used individually. Specifically, it hinders GKD-Pro from extracting comprehensive pathological-relevant features via LR-QFormer, and reduces key-value diversity in SRCA, impairing the construction of spatially aligned feature in GKD-Den. More notably, their joint application under image-level causes more performance degradation. To clearly show the effect of group-level and image-level distillation on the WLI classifier, we visualize the feature distributions using t-SNE (Maaten and Hinton 2008). As shown in Figure. 2, group-level distillation better preserves inter-class separability and spatial structure aligned with the pre-trained NBI features, whereas image-level distillation yields entangled and inconsistent distributions, demonstrating the superiority of group-level distillation under unpaired settings.

Hyperparameters Configuration. Figure 3 presents the results under different hyperparameter settings: the number of lesion queries in LR-QFormer N and training batch size s , where the batch size s is positively correlated with the group size of each class. PaGKD using $N=12$ and $s=24$ shows the best performance, respectively, and we adopt this configuration as the default setting of PaGKD.

Methods	PICCOLO			IH-GC		
	ACC	F1	AUC	ACC	F1	AUC
► <i>PolypsAlign</i>						
+ CycleGAN	70.2	66.6	78.8	62.5	66.6	59.6
+ SynDiff	73.1	69.0	80.4	69.6	73.3	67.7
► <i>ADD</i>						
+ CycleGAN	74.3	71.3	80.5	66.1	64.1	71.5
+ SynDiff	74.5	72.2	81.7	70.2	66.9	72.4
PaGKD	81.9	81.1	90.1	81.5	85.7	84.0

Table 6: Comparisons with an alternative solution via the unpaired I2I translation method.




Figure 3: Ablation on two hyperparameters: N and s .

Comparison with an Alternative Solution

Existing CDC methods like those in Table 2 are designed for paired images, one possible workaround for unpaired data is to synthesize missing modalities via unpaired image-to-image (I2I) translation, enabling CDC methods to be effectively applied. To evaluate this alternative, we combine two representative I2I methods, CycleGAN (Zhu et al. 2017) and SynDiff (Özbeý et al. 2023), with the two best CDC methods excluding PaGKD in Table 2, *i.e.*, PolypsAlign and ADD. As shown in Table 6, our PaGKD remains superior to this alternative I2I-based solution. These methods have not yet effectively enabled ADD and PolypsAlign to benefit from additional unpaired data, and in some metrics, even lag behind their counterparts trained solely on paired data. The primary reason is that existing I2I methods focus on translating global imaging styles, such as imaging color, while neglecting modality translation of pathology-relevant content, such as vascular structure. As a result, the synthesized images often lack pathological realism, reducing their utility for downstream classification.

Conclusion

In this paper, we propose PaGKD, which addresses the limitations of existing CDC methods that rely on paired data. Motivated by the semantic inconsistencies in unpaired data, we design two group-level distillation modules: GKD-Pro for learning modality-invariant semantic prototypes, and GKD-Den for capturing fine-grained structural features via group-level dense alignment. Our method can be applied to multi-modal learning scenarios without requiring instance-wise correspondence. Extensive experiments on four clinical datasets demonstrate the effectiveness and robustness of PaGKD in real-world GI classification, showcasing its strong clinical applicability. Nonetheless, PaGKD currently adopts a random group formation strategy; future work will explore an adaptive or data-driven grouping strategy.

Acknowledgments

This work was supported in part by National Natural Science Foundation of China (Grant No.62202189), research grants from Wuhan United Imaging Healthcare Surgical Technology Co., Ltd. Thanks to Professor Mei Liu's team for providing the in-house dataset used in this work at the Department of Gastroenterology, Tongji Medical College, Huazhong University of Science and Technology.

References

- Beg, S.; Card, T.; Warburton, S.; Rahman, I.; Wilkes, E.; White, J.; and Ragonath, K. 2020. Diagnosis of Barrett's esophagus and esophageal varices using a magnetically assisted capsule endoscopy system. *Gastrointestinal Endoscopy*, 91(4): 773–781.
- Deng, J.; Dong, W.; Socher, R.; Li, L.-J.; Li, K.; and Fei-Fei, L. 2009. Imagenet: A large-scale hierarchical image database. In *2009 IEEE conference on computer vision and pattern recognition*, 248–255. Ieee.
- Dosovitskiy, A.; Beyer, L.; Kolesnikov, A.; Weissenborn, D.; Zhai, X.; Unterthiner, T.; Dehghani, M.; Minderer, M.; Heigold, G.; Gelly, S.; et al. 2020. An image is worth 16x16 words: Transformers for image recognition at scale. *arXiv preprint arXiv:2010.11929*.
- Guo, X.; and Yuan, Y. 2020. Semi-supervised WCE image classification with adaptive aggregated attention. *Medical Image Analysis*, 64: 101733.
- He, K.; Zhang, X.; Ren, S.; and Sun, J. 2016. Deep residual learning for image recognition. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, 770–778.
- Hu, Q.; Wang, Q.; Chen, J.; Ji, X.; Li, Q.; and Wang, Z. 2025. Holistic White-light Polyp Classification via Alignment-free Dense Distillation of Auxiliary Optical Chromoendoscopy. *arXiv preprint arXiv:2505.19319*.
- Kingma, D. P.; and Ba, J. 2014. Adam: A method for stochastic optimization. *arXiv preprint arXiv:1412.6980*.
- Li, K.; Fathan, M. I.; Patel, K.; Zhang, T.; Zhong, C.; Bansal, A.; Rastogi, A.; Wang, J. S.; and Wang, G. 2021. Colonoscopy polyp detection and classification: Dataset creation and comparative evaluations. *Plos one*, 16(8): e0255809.
- Ma, W.; Zhu, Y.; Zhang, R.; Yang, J.; Hu, Y.; Li, Z.; and Xiang, L. 2022. Toward clinically assisted colorectal polyp recognition via structured cross-modal representation consistency. In *International Conference on Medical Image Computing and Computer-Assisted Intervention*, 141–150. Springer.
- Maaten, L. v. d.; and Hinton, G. 2008. Visualizing data using t-SNE. *Journal of machine learning research*, 9(Nov): 2579–2605.
- Özbey, M.; Dalmaz, O.; Dar, S. U.; Bedel, H. A.; Öztürk, Ş.; Güngör, A.; and Cukur, T. 2023. Unsupervised medical image translation with adversarial diffusion models. *IEEE Transactions on Medical Imaging*, 42(12): 3524–3539.
- Paszke, A.; Gross, S.; Massa, F.; Lerer, A.; Bradbury, J.; Chanan, G.; Killeen, T.; Lin, Z.; Gimelshein, N.; Antiga, L.; et al. 2019. Pytorch: An imperative style, high-performance deep learning library. *Advances in neural information processing systems*, 32.
- Ru, L.; Zhan, Y.; Yu, B.; and Du, B. 2022. Learning affinity from attention: End-to-end weakly-supervised semantic segmentation with transformers. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 16846–16855.
- Sánchez-Peralta, L. F.; Pagador, J. B.; Picón, A.; Calderón, Á. J.; Polo, F.; Andraka, N.; Bilbao, R.; Glover, B.; Saratxaga, C. L.; and Sánchez-Margallo, F. M. 2020. Piccolo white-light and narrow-band imaging colonoscopic dataset: A performance comparative of models and datasets. *Applied Sciences*, 10(23): 8501.
- Schmitz, R.; Werner, R.; Repici, A.; Bisschops, R.; Meinig, A.; Zornow, M.; Messmann, H.; Hassan, C.; Sharma, P.; and Rösch, T. 2022. Artificial intelligence in GI endoscopy: stumbling blocks, gold standards and the role of endoscopy societies. *Gut*, 71(3): 451–454.
- Shen, Y.; Wang, X.; Gao, P.; and Lin, M. 2023. Auxiliary modality learning with generalized curriculum distillation. In *International Conference on Machine Learning*, 31057–31076. PMLR.
- Soffer, S.; Klang, E.; Shimon, O.; Nachmias, N.; Eliakim, R.; Ben-Horin, S.; Kopylov, U.; and Barash, Y. 2020. Deep learning for wireless capsule endoscopy: a systematic review and meta-analysis. *Gastrointestinal endoscopy*, 92(4): 831–839.
- Tamura, N.; Sakaguchi, Y.; Furutani, W.; Matsui, M.; Nagao, S.; Sakuma, N.; Fukagawa, K.; Miura, Y.; Mizutani, H.; Ohki, D.; et al. 2022. Magnifying endoscopy with narrow-band imaging is useful in differentiating gastric cancer from matched adenoma in white light imaging. *Scientific Reports*, 12(1): 8349.
- Vaswani, A.; Shazeer, N.; Parmar, N.; Uszkoreit, J.; Jones, L.; Gomez, A. N.; Kaiser, Ł.; and Polosukhin, I. 2017. Attention is all you need. *Advances in neural information processing systems*, 30.
- Wang, K.-N.; He, Y.; Zhuang, S.; Miao, J.; He, X.; Zhou, P.; Yang, G.; Zhou, G.-Q.; and Li, S. 2022. Ffcnet: Fourier transform-based frequency learning and complex convolutional network for colon disease classification. In *International Conference on Medical Image Computing and Computer-Assisted Intervention*, 78–87. Springer.
- Wang, K.-N.; Zhuang, S.; Ran, Q.-Y.; Zhou, P.; Hua, J.; Zhou, G.-Q.; and He, X. 2023. Dlgnet: A dual-branch lesion-aware network with the supervised gaussian mixture model for colon lesions classification in colonoscopy images. *Medical Image Analysis*, 87: 102832.
- Wang, Q.; Che, H.; Ding, W.; Xiang, L.; Li, G.; Li, Z.; and Cui, S. 2021. Colorectal polyp classification from white-light colonoscopy images via domain alignment. In *Medical Image Computing and Computer Assisted Intervention–MICCAI 2021: 24th International Conference, Strasbourg*,

France, September 27–October 1, 2021, *Proceedings, Part VII 24*, 24–32. Springer.

Xu, Z.; Rittscher, J.; and Ali, S. 2024. SSL-CPCD: Self-supervised learning with composite pretext-class discrimination for improved generalisability in endoscopic image analysis. *IEEE Transactions on Medical Imaging*, 43(12): 4105–4119.

Yue, G.; Zhuo, G.; Li, S.; Zhou, T.; Du, J.; Yan, W.; Hou, J.; Liu, W.; and Wang, T. 2023. Benchmarking polyp segmentation methods in narrow-band imaging colonoscopy images. *IEEE Journal of Biomedical and Health Informatics*, 27(7): 3360–3371.

Zhang, R.; Zheng, Y.; Mak, T. W. C.; Yu, R.; Wong, S. H.; Lau, J. Y.; and Poon, C. C. 2016. Automatic detection and classification of colorectal polyps by transferring low-level CNN features from nonmedical domain. *IEEE journal of biomedical and health informatics*, 21(1): 41–47.

Zhou, B.; Khosla, A.; Lapedriza, A.; Oliva, A.; and Torralba, A. 2016. Learning deep features for discriminative localization. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, 2921–2929.

Zhu, J.-Y.; Park, T.; Isola, P.; and Efros, A. A. 2017. Unpaired image-to-image translation using cycle-consistent adversarial networks. In *Proceedings of the IEEE international conference on computer vision*, 2223–2232.