scientific reports

Check for updates

OPEN

Deep learning-based image classification for integrating pathology and radiology in AI-assisted medical imaging

Chenming Lu[1], Jiayin Zhang[1]✉ & Ren Liu[2]

The integration of pathology and radiology in medical imaging has emerged as a critical need for advancing diagnostic accuracy and improving clinical workflows. Current AI-driven approaches for medical image analysis, despite significant progress, face several challenges, including handling multi-modal imaging, imbalanced datasets, and the lack of robust interpretability and uncertainty quantification. These limitations often hinder the deployment of AI systems in real-world clinical settings, where reliability and adaptability are essential. To address these issues, this study introduces a novel framework, the Domain-Informed Adaptive Network (DIANet), combined with an Adaptive Clinical Workflow Integration (ACWI) strategy. DIANet leverages multi-scale feature extraction, domain-specific priors, and Bayesian uncertainty modeling to enhance interpretability and robustness. The proposed model is tailored for multi-modal medical imaging tasks, integrating adaptive learning mechanisms to mitigate domain shifts and imbalanced datasets. Complementing the model, the ACWI strategy ensures seamless deployment through explainable AI (XAI) techniques, uncertainty-aware decision support, and modular workflow integration compatible with clinical systems like PACS. Experimental results demonstrate significant improvements in diagnostic accuracy, segmentation precision, and reconstruction fidelity across diverse imaging modalities, validating the potential of this framework to bridge the gap between AI innovation and clinical utility.

Keywords Medical Imaging, Deep Learning, Multi-Modal Integration, Explainable AI, Clinical Workflow Adaptation

The integration of pathology and radiology represents a transformative frontier in medical imaging, offering unprecedented opportunities for precision diagnostics. Pathology provides a microscopic view of cellular morphology, while radiology offers a macroscopic perspective of anatomical and functional imaging[1]. The synergy of these domains not only enhances diagnostic accuracy but also enables a holistic understanding of disease processes, particularly in oncology and chronic illnesses. However, manual correlation between pathology and radiology is labor-intensive and prone to interobserver variability. AI-assisted image classification aims to address these limitations by automating the integration of multimodal imaging data[2]. This integration not only optimizes diagnostic workflows but also holds the potential to improve patient outcomes through personalized treatment strategies. Therefore, developing deep learning approaches for combining radiological and pathological insights is not only innovative but also critical for advancing precision medicine in clinical practice[3].

To address the challenge of multimodal integration, initial approaches relied heavily on traditional methods based on symbolic AI and knowledge representation. These methods aimed to encode human expertise into rule-based systems capable of analyzing both pathology and radiology images[4]. Symbolic AI leveraged handcrafted features, such as edge detection, texture analysis, and statistical shape models, to characterize image content. These systems provided interpretability and domain-specific insights, offering a solid foundation for integrating structured radiological findings and digitized pathology. However, their reliance on predefined features limited scalability and generalizability to complex imaging datasets[5]. Moreover, these methods struggled with noisy, high-dimensional data common in medical imaging, making them inadequate for dynamic clinical settings. As such, while symbolic AI laid the groundwork for multimodal analysis, its inability to adapt to diverse imaging modalities and disease variability underscored the need for more robust and data-driven solutions[6].

[1]Shanghai General Hospital, Shanghai, China. [2]Department of Visual Communication Design, LuXun Academy of Fine Arts, Dalian, Liaoning 116650, China. ✉email: eite859@163.com

The subsequent evolution of AI methodologies focused on data-driven and machine learning techniques, marking a departure from handcrafted features to automated feature extraction. Machine learning algorithms, particularly support vector machines (SVMs) and random forests, demonstrated promising performance in analyzing pathology and radiology data independently[7]. The introduction of feature engineering pipelines enabled the extraction of mid-level image representations that could bridge the gap between the two modalities. For example, radiological tumor margins could be correlated with cellular characteristics from histopathology using shared imaging descriptors. While these methods improved the accuracy and scalability of multimodal image analysis, they required extensive manual tuning and domain expertise for feature design[8]. The fragmented nature of pathology and radiology datasets, coupled with variations in image resolution and formats, posed significant integration challenges. Despite these limitations, machine learning approaches served as a stepping stone, paving the way for deep learning models capable of end-to-end learning[9].

The advent of deep learning, particularly convolutional neural networks (CNNs) and transformer-based architectures, revolutionized image classification in medical imaging. Unlike traditional methods, deep learning models can learn hierarchical feature representations directly from raw data, enabling seamless integration of multimodal inputs[10]. In pathology and radiology, CNNs have been used to identify tumor subtypes, segment lesions, and predict molecular markers with high accuracy. Transformer-based models, leveraging self-attention mechanisms, have further enhanced the ability to fuse spatial and contextual information from diverse image sources[11]. Pretrained models, such as Vision Transformers (ViTs) and multimodal transformers, offer a unified framework for analyzing pathology slides and radiological scans simultaneously. These models address the challenges of feature heterogeneity and dataset fragmentation, achieving state-of-the-art performance in multimodal image classification[12]. However, deep learning systems remain limited by their reliance on large annotated datasets and computational resources, as well as interpretability concerns that hinder clinical adoption.

Based on the limitations of existing approaches, our proposed method focuses on a novel deep learning framework that seamlessly integrates pathology and radiology data for AI-assisted image classification. Unlike previous methods that treat these modalities independently, our approach incorporates a multimodal attention mechanism to align spatial and contextual features across imaging domains. By leveraging self-supervised learning and transfer learning, we overcome the challenge of limited labeled datasets while ensuring robust performance across diverse clinical scenarios. Moreover, our framework emphasizes interpretability through attention heatmaps, enabling clinicians to validate model predictions and enhance trust in AI systems. This novel approach not only bridges the gap between pathology and radiology but also addresses the scalability, generalizability, and clinical utility challenges inherent in existing methodologies.

The proposed method has several key advantages:

- The proposed method introduces a multimodal attention mechanism and self-supervised learning module, enabling end-to-end integration of pathology and radiology data.
- The model is designed to operate efficiently across diverse imaging modalities, offering scalability and adaptability for real-world clinical applications.
- Extensive experiments demonstrate superior performance in multimodal image classification tasks, achieving significant improvements in accuracy, robustness, and interpretability compared to state-of-the-art methods.

Related work

Multi-modal data integration

Recent advancements in AI-assisted medical imaging have highlighted the importance of integrating multi-modal data, particularly pathology and radiology, to improve diagnostic accuracy[13]. Radiology, which involves the analysis of medical imaging modalities such as X-rays, CT scans, and MRIs, provides a macroscopic perspective on patient health. In contrast, pathology examines tissue samples on a microscopic level, offering cellular and molecular insights. Combining these two modalities enables a comprehensive understanding of diseases, particularly in complex cases such as cancer diagnosis and progression monitoring[14]. Deep learning-based approaches have demonstrated significant potential in unifying these modalities. Techniques such as convolutional neural networks (CNNs) and attention-based architectures (e.g., Transformers) have been used to extract features from both radiology images and digital pathology slides[15]. These methods often rely on shared feature spaces or domain adaptation to align the disparate nature of the data. For instance, a CNN trained on radiology data can be fine-tuned on pathology data, leveraging transfer learning to improve performance in cross-domain tasks. More recently, multi-modal neural networks, such as joint-embedding models, have emerged to integrate heterogeneous datasets[16]. These models aim to align radiology and pathology data into a single latent space where cross-modal comparisons can be performed effectively. The integration of pathology and radiology data has also benefited from self-supervised and semi-supervised learning methods[17]. These approaches are particularly useful in medical imaging, where labeled datasets are limited due to the need for expert annotation. By leveraging unlabeled data from one modality to guide learning in the other, researchers have achieved improved performance in tasks such as tumor classification and segmentation[18]. However, challenges remain, particularly in harmonizing data from different acquisition protocols, image resolutions, and noise characteristics. Future research is focusing on designing more robust models that can handle these variations and provide interpretable results that align with clinical workflows[19].

Deep learning for feature extraction

Feature extraction is a critical step in integrating radiology and pathology data. In radiology, features often include structural patterns, such as tumor size, shape, and location, while pathology focuses on cellular-level characteristics, such as nuclear morphology and mitotic activity[20]. Deep learning has revolutionized this process by automating the extraction of complex and hierarchical features, which were previously identified through

manual annotation or basic image processing techniques. Convolutional neural networks have been at the forefront of feature extraction in both radiology and pathology[21]. In radiology, pre-trained CNN models such as ResNet, DenseNet, and EfficientNet have achieved state-of-the-art results in detecting abnormalities like lung nodules or brain tumors. These networks extract spatial features that capture the macroscopic structure of tissues. In pathology, models like VGG and U-Net are widely used for cell segmentation and classification tasks[22]. These architectures excel at identifying fine-grained details, such as cell boundaries and chromatin patterns, that are critical for pathology diagnosis. To combine radiology and pathology features, researchers have employed fusion strategies that aggregate multi-scale representations from both domains[23]. Early fusion techniques integrate raw data at the pixel or voxel level, enabling networks to learn joint features from the outset. Late fusion approaches, on the other hand, combine features extracted independently from radiology and pathology modalities. For example, embeddings from a CNN trained on CT images can be concatenated with embeddings from a CNN trained on histopathology slides to form a unified representation[24]. Advanced techniques like attention mechanisms and graph neural networks (GNNs) have further enhanced feature extraction by capturing interdependencies between features, such as spatial relationships in radiology and morphological structures in pathology[25]. Despite these advancements, challenges persist in ensuring that extracted features are clinically meaningful and interpretable. Black-box models often fail to explain how specific features contribute to predictions, limiting their adoption in clinical settings[26]. Researchers are addressing this issue by incorporating explainability into model design, such as by using attention heatmaps or integrating domain knowledge directly into the feature extraction process[27].

AI-powered diagnosis and prognosis

The application of deep learning in AI-assisted diagnosis and prognosis has shown significant promise, particularly in the integration of pathology and radiology for comprehensive disease assessment[28]. Diagnosis involves identifying diseases or abnormalities, while prognosis predicts disease progression and patient outcomes. Deep learning models have been widely adopted for both tasks, leveraging the complementary strengths of pathology and radiology data[29]. In diagnosis, classification tasks are a common focus. For instance, CNNs have been employed to classify lung cancer types using CT scans, while digital pathology data has been used for histological subtyping[30]. Integrating these datasets allows for more accurate classification by providing both macroscopic and microscopic perspectives. Multi-task learning frameworks, which perform multiple diagnostic tasks simultaneously, have further improved model performance by exploiting shared information across tasks[31]. For example, a multi-task model might simultaneously detect tumor presence, predict its type, and assess its grade, using both radiological and pathological inputs[32]. Prognosis, on the other hand, often involves predicting survival rates, recurrence risks, or treatment response. Deep learning models have leveraged longitudinal data from both modalities to model temporal changes in disease progression[33]. Recurrent neural networks (RNNs) and Transformer-based architectures are particularly well-suited for such tasks, as they can capture temporal dependencies and model sequences of imaging data over time. Survival analysis models, such as Cox proportional hazard models, have also been enhanced by integrating deep learning-derived features from pathology and radiology[34]. A major challenge in AI-powered diagnosis and prognosis is the variability in data quality and annotation. Radiology images may suffer from noise or low resolution, while pathology slides often require labor-intensive preprocessing[35]. To address these issues, researchers are focusing on robust data augmentation techniques and domain adaptation methods. Moreover, the integration of interpretability into diagnostic models is gaining traction, as clinicians require transparent models to trust AI-assisted predictions. Explainable AI (XAI) techniques, such as saliency maps and SHAP (Shapley Additive Explanations), are increasingly being adopted to elucidate how models make decisions based on integrated pathology and radiology data[36].

Method

Overview

In recent years, artificial intelligence (AI) methods have demonstrated transformative potential in medical imaging, revolutionizing key tasks such as disease diagnosis, image reconstruction, and workflow optimization. This subsection provides a high-level overview of our method designed to address limitations in current medical imaging AI workflows. Our approach builds upon recent advancements in deep learning, computational efficiency, and explainable AI to improve model performance and usability in real-world clinical environments.

The structure of this section is as follows. In Section 3.2, we introduce the preliminary concepts and define the challenges inherent in medical imaging tasks, including the need for robust AI systems that generalize across diverse patient populations. In Section 3.3, we present our novel model architecture, which leverages domain-specific constraints to enhance interpretability and diagnostic accuracy. In Section 3.4, we describe our innovative strategy, which integrates adaptive learning and uncertainty quantification to optimize decision-making in clinical workflows. This combination of model design and strategic implementation aims to bridge the gap between cutting-edge research and practical utility in healthcare systems. By systematically exploring these components, we aim to provide a comprehensive framework for advancing medical imaging AI, particularly in challenging scenarios such as multi-modal imaging and imbalanced datasets. The sections that follow will formalize the problem setup, introduce our contributions, and highlight the strategies employed to achieve reliable performance and deployment readiness.

Preliminaries

Medical imaging plays a critical role in modern healthcare, encompassing a wide array of imaging modalities such as X-rays, computed tomography (CT), magnetic resonance imaging (MRI), ultrasound, and positron emission tomography (PET). Each modality provides unique insights into anatomical structures or physiological

processes, facilitating tasks such as disease diagnosis, monitoring, and treatment planning. Despite their clinical utility, the interpretation of medical images presents challenges due to high data complexity, inter-patient variability, and the potential for subjectivity in manual diagnosis.

To address these challenges, medical imaging AI has emerged as a promising field, combining advanced deep learning algorithms with domain-specific constraints. The task of medical image analysis can be formalized as follows. Let $\mathcal{X}$ denote the input space of medical images, where each $x \in \mathcal{X}$ represents a high-dimensional image tensor. For instance, an MRI scan may be modeled as a 3D tensor, $x \in \mathbb{R}^{H\times W\times D}$, where *H*, *W*, and *D* are the height, width, and depth of the image, respectively. Let $\mathcal{Y}$ represent the output space, such as diagnostic labels, segmentation masks, or reconstructed image volumes. The primary goal is to learn a mapping function $f_\theta : \mathcal{X} \to \mathcal{Y}$, parameterized by θ, that minimizes a task-specific objective function.

$$\mathcal{L}(\theta) = \frac{1}{N}\sum_{i=1}^{N} \ell(f_\theta(x_i), y_i), \tag{1}$$

where $\ell(\cdot)$ is a loss function that quantifies the discrepancy between the predicted output $f_\theta(x_i)$ and the ground truth y_i, and *N* is the number of training samples.

Medical imaging problems are further complicated by factors such as imbalanced datasets, noise, and variability in image quality. For example, datasets may exhibit skewed class distributions, where rare diseases have significantly fewer labeled samples compared to common conditions. To account for such challenges, we consider the problem in a probabilistic framework. Given an observed image *x*, the true label *y* is modeled as a random variable with conditional probability distribution *P*(*y*|*x*). The goal of supervised learning is to approximate *P*(*y*|*x*) using the parametric model $f_\theta(x)$. A probabilistic formulation enables techniques such as uncertainty quantification, essential for clinical decision-making.

Image segmentation, a fundamental task in medical imaging, involves delineating regions of interest (ROIs) such as tumors or organs. Formally, segmentation can be viewed as a pixel-wise classification problem. Let $x \in \mathbb{R}^{H\times W}$ represent a 2D medical image, and let $y \in \{0, 1, \ldots, C\}^{H\times W}$ denote the corresponding segmentation mask, where *C* is the number of classes. The objective is to optimize:

$$\mathcal{L}_{\text{seg}}(\theta) = -\frac{1}{HW}\sum_{i=1}^{H}\sum_{j=1}^{W}\sum_{c=1}^{C} y_{i,j,c}\log f_\theta(x)_{i,j,c}, \tag{2}$$

where $f_\theta(x)_{i,j,c}$ is the predicted probability for pixel (*i*, *j*) belonging to class *c*.

In medical image reconstruction, the goal is to recover a high-quality image x_{recon} from noisy or incomplete measurements x_{obs}. For example, in accelerated MRI, the undersampled measurement x_{obs} is related to the fully sampled image x_{recon} via a forward model $\mathcal{F}$:

$$x_{\text{obs}} = \mathcal{F}(x_{\text{recon}}) + \epsilon, \tag{3}$$

where $\mathcal{F}$ denotes the sampling operator, and ϵ represents noise. Reconstruction methods aim to solve the inverse problem by finding x_{recon} that minimizes:

$$\mathcal{L}_{\text{recon}}(\theta) = \|x_{\text{obs}} - \mathcal{F}(f_\theta(x_{\text{recon}}))\|_2^2. \tag{4}$$

Another crucial problem is multi-modal medical imaging, where information from multiple imaging modalities (e.g., MRI and CT) must be integrated for comprehensive analysis. Let $x^{(1)}, x^{(2)}, \ldots, x^{(M)}$ represent input images from *M* modalities. The task is to learn a joint representation:

$$z = g_\phi(x^{(1)}, x^{(2)}, \ldots, x^{(M)}), \tag{5}$$

where g_ϕ is a feature fusion function parameterized by ϕ. The joint representation *z* is then used for downstream tasks, such as classification or segmentation.

The backbone network extracts modality-specific features from radiology or pathology images. These features are then projected into a domain-informed latent space using a structured transformation layer that aligns them across modalities. The DIANet module operates on this latent space, applying adaptive attention to recalibrate and fuse the cross-modal features. The ACWI module takes the fused representations and adjusts the output flow based on uncertainty estimation, enabling the model to select the most reliable prediction path. All components are implemented as lightweight modules that wrap around the backbone without modifying its internal architecture, allowing seamless end-to-end integration.

Domain-Informed Adaptive Network (DIANet)

In this section, we present our proposed model, referred to as the Domain-Informed Adaptive Network (DIANet), which is specifically designed to address the challenges of medical imaging tasks. DIANet introduces a unified architecture that incorporates domain knowledge, multi-scale feature extraction, and task-specific attention mechanisms to improve the interpretability and robustness of deep learning-based medical imaging solutions. This model is developed to handle complex clinical scenarios such as imbalanced datasets, multi-modal imaging, and uncertainty quantification. Below, we detail the structure and key innovations of DIANet (As shown in Figure 1).

The diffuser module is a key component of the latent feature refinement process. It operates by applying a series of nonlinear transformations that adjust the encoded features according to both spatial and contextual relationships within the medical images. Specifically, after the initial latent vector is generated from the fused multi-scale features, the diffuser integrates attention-based operations to redistribute representational emphasis. This mechanism is particularly important for handling multi-modal inputs, where spatial alignment and semantic consistency are critical. By diffusing features through this contextual transformation process, the model achieves better adaptation to structural variances across domains such as pathology and radiology. The CodeBook module complements the diffuser by introducing a discrete set of learnable vectors that represent prototypical latent patterns. Each incoming latent feature is softly matched against this set, enabling the model to quantize the representation space in a way that preserves meaningful anatomical priors. This quantization process not only regularizes the feature space but also promotes inter-sample consistency, which is vital for clinical reliability. The use of a CodeBook allows the latent space to capture a structured representation that aligns with expected biological or anatomical patterns, especially in heterogeneous datasets.

Multi-scale feature encoding

The multi-scale feature encoder is designed to capture information at different spatial resolutions, allowing the model to integrate both fine-grained details and global context, which is crucial for medical image analysis. Let the input medical image be represented as $x \in \mathbb{R}^{H \times W \times C}$, where H and W are the height and width of the image, and C is the number of input channels (e.g., grayscale or RGB). The encoder is composed of a series of convolutional blocks $\{f_i\}_{i=1}^{L}$, where L denotes the number of blocks. Each block f_i operates at a different spatial resolution, progressively extracting hierarchical features. Initially, the input image is passed through the first block f_1, producing an output feature map $z_1 = f_1(x)$. Each subsequent block operates on the output from the previous block, progressively downsampling the feature maps, as shown by the recurrence:

$$z_i = f_i(z_{i-1}), \quad \text{where } z_0 = x. \tag{6}$$

The output of each block $z_i \in \mathbb{R}^{H_i \times W_i \times D_i}$, where D_i is the number of feature channels at scale i, and H_i, W_i are the spatial dimensions of the feature map after downsampling. As the resolution decreases, the encoder captures more abstract, high-level representations of the image. This hierarchical feature representation allows the model to capture both local and global information. To further refine the learned features, the multi-scale

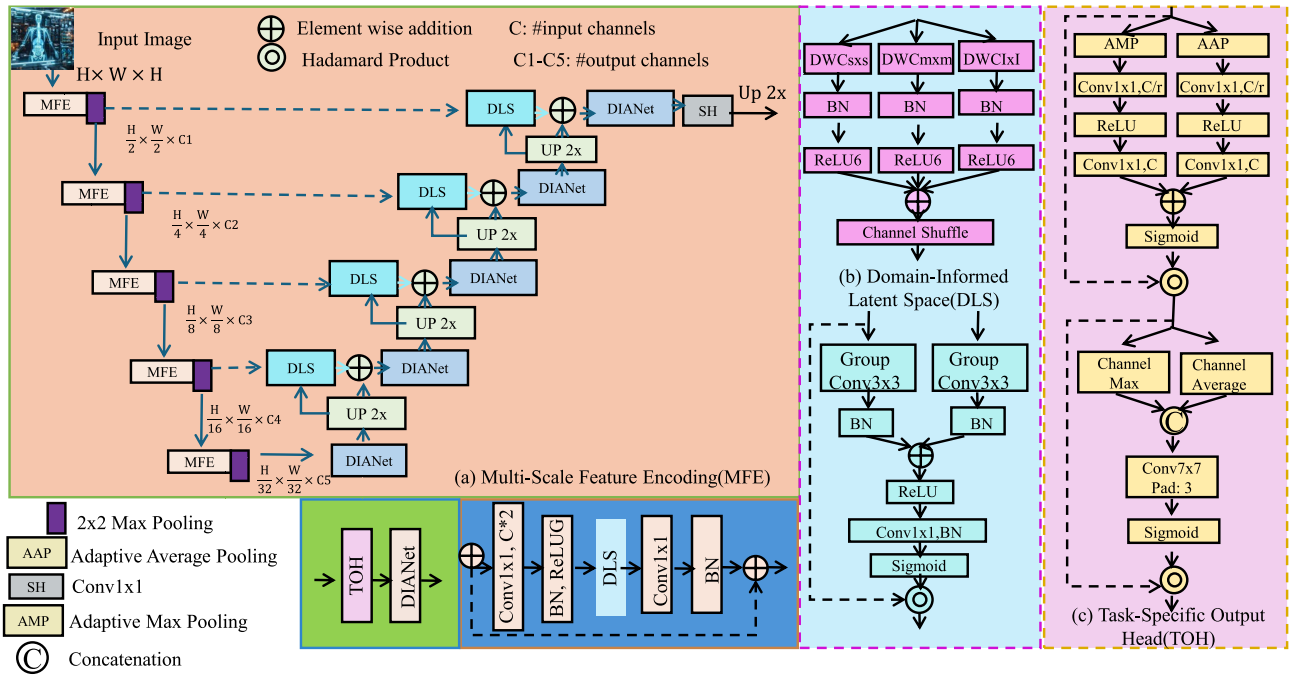

Fig. 1. Overview of the Domain-Informed Adaptive Network (DIANet) architecture, incorporating multi-scale feature encoding (MFE), domain-informed latent space (DLS), and task-specific output heads (TOH). The diagram illustrates how DIANet processes medical images with a hierarchical feature extraction approach, integrates domain-specific knowledge, and adapts to different medical imaging tasks such as classification, segmentation, and reconstruction. The model's robust design is equipped with attention mechanisms, multi-scale feature fusion, and uncertainty quantification, making it suitable for handling complex clinical scenarios, such as imbalanced datasets and multi-modal imaging.

outputs are aggregated through the Feature Pyramid Fusion (FPF) module. The FPF module combines features from different scales, providing both local details and global context:

$$z_{\text{fused}} = \text{FPF}(\{z_i\}_{i=1}^{L}), \tag{7}$$

where z_{fused} is the fused feature map that retains multi-scale information. In this process, features from different scales are merged through both top-down and bottom-up pathways. To improve the feature integration, we apply a lateral connection that ensures information flows effectively between different scales:

$$z_{\text{fused}}^{L} = \sum_{i=1}^{L} \alpha_i z_i, \tag{8}$$

where α_i are learnable coefficients determining the importance of each scale. This process enhances the robustness of the model by allowing it to leverage both fine-grained and coarse features effectively. To further enhance multi-scale learning, we apply spatial attention mechanisms that allow the model to focus on important regions at each scale. The attention mechanism at each scale can be represented as:

$$z_{\text{attended}} = \text{Attention}(z_i), \tag{9}$$

where the attention operation refines the feature map z_i by emphasizing important regions and suppressing irrelevant ones.

The MultiScaleFusion module integrates multi-resolution feature maps from different stages of the backbone network to capture both global contextual patterns and fine-grained structural cues[37,38]. Inspired by hierarchical fusion strategies in feature pyramid networks and attention-guided feature refinement, our module employs adaptive weighting mechanisms to recalibrate features at each scale before fusion[39]. Specifically, it applies spatial attention to emphasize clinically salient regions and aggregates features through channel-wise weighting, ensuring information from deeper layers is effectively aligned with shallower representations[40]. Unlike conventional concatenation or summation methods, our approach maintains semantic consistency across modalities while enhancing robustness against noise and resolution disparity[41].

Domain-informed latent space

DIANet integrates domain knowledge into the learning process to enhance model robustness and interpretability. The intermediate feature representation, denoted as z_{fused}, is passed through a transformation function g_ϕ, which maps it to a domain-informed latent space. This transformation is defined as:

$$h = g_\phi(z_{\text{fused}}), \quad h \in \mathbb{R}^d, \tag{10}$$

where h represents the latent representation vector, and d is the dimensionality of the latent space. The latent space is structured to preserve essential domain-specific information, which is crucial for downstream tasks such as classification and segmentation. To ensure that the learned latent space aligns with anatomical knowledge, we introduce domain-informed regularizers. These regularizers guide the learning process by enforcing consistency between the learned latent representation and predefined anatomical priors. Specifically, the regularization term is defined as:

$$\mathcal{R}_{\text{domain}} = \lambda_1 \|\mathcal{P}(h) - \mathcal{P}_{\text{target}}\|_2^2, \tag{11}$$

where $\mathcal{P}(h)$ represents the predicted anatomical distribution derived from the latent space h, and $\mathcal{P}_{\text{target}}$ is the anatomical prior, which can be a known distribution such as a probability map of organ locations. This regularization helps to shape the learned representation according to prior anatomical knowledge, thus improving generalization across different domains. Moreover, to model the uncertainty inherent in medical imaging data, we treat the latent representation h as a probabilistic distribution rather than a fixed point. Specifically, we model h as a multivariate Gaussian distribution:

$$h \sim \mathcal{N}(\mu, \Sigma), \tag{12}$$

where μ and Σ represent the mean and covariance of the distribution, respectively. These parameters are generated through the transformation function g_ϕ, allowing the model to encode not only the point estimate but also the uncertainty about the latent representation. This probabilistic approach provides a measure of confidence in the model's predictions, which is essential in medical applications. The mean μ and covariance Σ are computed as:

$$\mu = g_\mu(z_{\text{fused}}), \quad \Sigma = g_\Sigma(z_{\text{fused}}), \tag{13}$$

where g_μ and g_Σ are separate networks parameterized by ϕ that produce the mean and covariance, respectively. In addition to these primary equations, we introduce a regularization term that encourages the learned latent space to conform to the known anatomy, thus facilitating accurate predictions across diverse domains:

$$\mathcal{R}_{\text{latency}} = \lambda_2 \|\mathcal{P}(h) - \mathcal{P}_{\text{prior}}\|_1, \tag{14}$$

where $\mathcal{P}_{\text{prior}}$ is an anatomical prior based on prior knowledge or synthetic data, and $\|\cdot\|_1$ denotes the L_1-norm. The domain alignment is enhanced using adversarial learning techniques. The adversarial loss encourages the model to generate latent representations that are indistinguishable from the target distribution. This is achieved by minimizing the following adversarial loss:

$$\mathcal{L}_{\text{adv}} = -\mathbb{E}_{h\sim p_{\text{data}}}[\log D(h)] - \mathbb{E}_{h\sim p_{\text{prior}}}[\log(1 - D(h))], \tag{15}$$

where $D(h)$ is the discriminator network that distinguishes between real and generated latent vectors, and p_{data} and p_{prior} represent the data and prior distributions, respectively. By incorporating these additional regularization and adversarial losses, DIANet ensures that the latent space not only captures the relevant domain-specific information but also models uncertainty and aligns with anatomical priors, making it robust for real-world clinical applications (As shown in Figure 2).

Task-specific output head

DIANet employs different task-specific output heads, each tailored for particular medical imaging tasks, including classification, segmentation, and reconstruction. Let h represent the latent representation generated by the network, and $\mathcal{T}$ denote the specific task at hand. The output head for each task is a mapping function $o_{\psi}^{(\mathcal{T})}$, defined as:

$$y = o_{\psi}^{(\mathcal{T})}(h), \tag{16}$$

where y represents the predicted output. In the case of classification, the output head uses a softmax function to transform the latent representation h into probabilities for each class. This is computed as:

$$o_{\psi}^{(\text{class})}(h) = \text{softmax}(Wh + b), \tag{17}$$

where W is the weight matrix, b is the bias term, and h is the input latent representation. The softmax function ensures that the outputs are probabilities, such that:

$$\sum_i \text{softmax}(Wh + b)_i = 1, \quad \forall i. \tag{18}$$

For segmentation tasks, the output is generated through a transposed convolution (also known as a deconvolution) operation, which increases the spatial resolution of the latent representation h to match the target segmentation map. This can be expressed as:

$$o_{\psi}^{(\text{seg})}(h) = \text{ConvTranspose}(h), \tag{19}$$

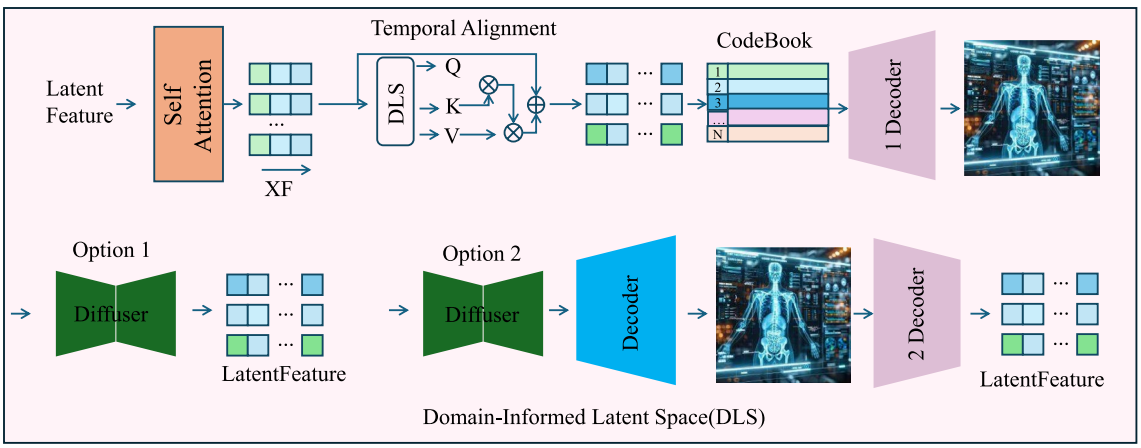

Fig. 2. The architecture of DIANet incorporates domain-informed latent space learning. It enhance robustness and interpretability in medical imaging tasks. The model starts with a latent feature that is processed through self-attention and temporal alignment modules, which interact with a codebook to generate transformed latent representations. The domain-informed latent space is further refined by domain-specific regularizers and adversarial learning techniques, ensuring that the learned representation aligns with anatomical knowledge. The diagram illustrates two options for model decoding: Option 1 uses a diffuser and a latent feature for reconstruction, while Option 2 introduces additional domain alignment techniques. The latent representation is modeled as a probabilistic distribution, facilitating uncertainty estimation in the medical predictions. This approach, as shown in the figure, supports accurate classification and segmentation by integrating prior anatomical knowledge and enhancing generalization across different domains.

where ConvTranspose denotes the transposed convolution operation. This operation involves the learned filters being applied in a reversed manner, helping recover the spatial dimensions of the input image. For finer details in the segmentation, skip connections from earlier layers in the network might be employed, enhancing the output resolution.

For reconstruction tasks, DIANet utilizes upsampling layers followed by convolutional layers to reconstruct the original image from the latent representation. This process can be mathematically expressed as:

$$o_{\psi}^{(\text{recon})}(h) = \text{Upsample}(h), \tag{20}$$

where Upsample refers to a process of increasing the spatial resolution of the feature map by interpolating the values of h. Following upsampling, convolutional layers are applied to refine the reconstructed image, as given by:

$$\tilde{x} = \text{Conv}(o_{\psi}^{(\text{recon})}(h)), \tag{21}$$

where $\tilde{x}$ denotes the final reconstructed image. In addition, for multi-scale features, DIANet may use a multi-resolution approach where the output head includes a series of layers, each responsible for different spatial resolutions. This helps the network learn both global and fine-grained features, thus improving the reconstruction accuracy.

Adaptive Clinical Workflow Integration (ACWI)

In this subsection, we introduce our new strategy, termed Adaptive Clinical Workflow Integration (ACWI), which complements the proposed model by addressing the operational challenges of deploying medical imaging AI systems in real-world clinical settings. ACWI is designed to ensure seamless integration, enhance model adaptability, and facilitate trust through explainability and uncertainty quantification. Below, we outline the core components of ACWI and their contributions to the clinical deployment pipeline (As shown in Figure 3).

Adaptive learning framework

Medical imaging datasets often suffer from domain shifts due to variations in imaging devices, protocols, and patient demographics. These domain shifts introduce significant challenges in transferring models trained on one dataset to another, especially when the distributions of the images differ. To handle these challenges, DIANet employs an adaptive learning framework that incorporates domain adaptation, transfer learning, and continual

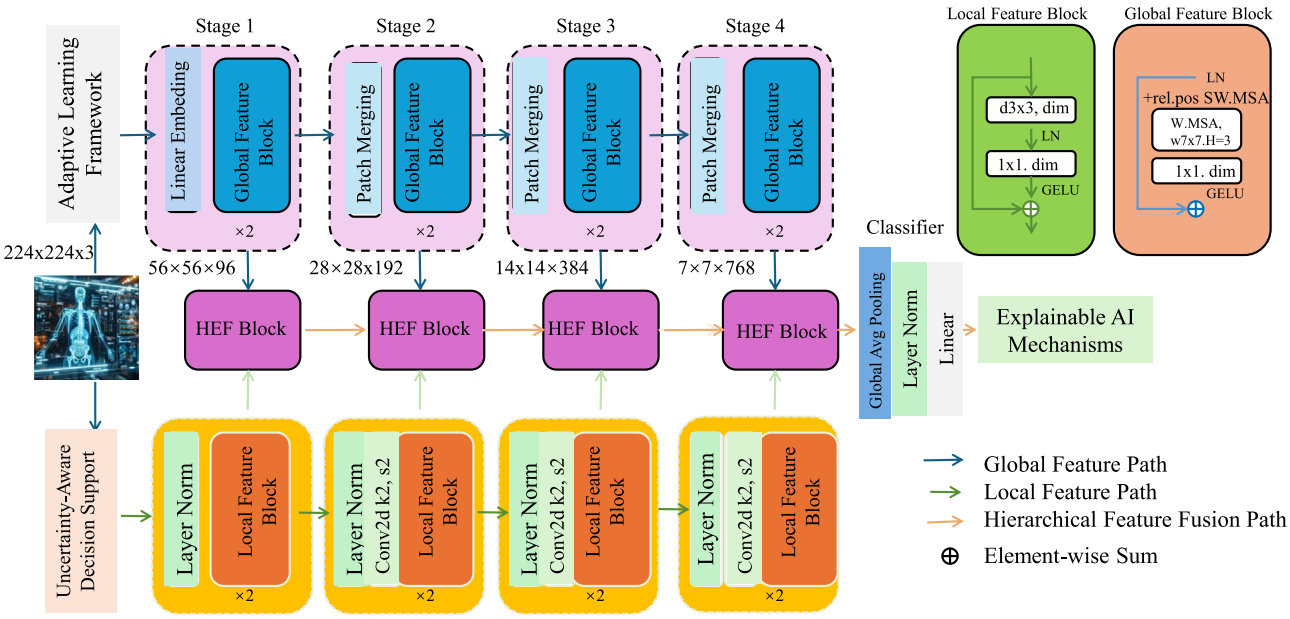

Adaptive Clinical Workflow Integration (ACWI)

Fig. 3. The architecture of the Adaptive Clinical Workflow Integration (ACWI) system, which consists of an adaptive learning framework, explainable AI mechanisms, and uncertainty-aware decision support for medical imaging AI. The figure illustrates the stages of the learning pipeline, including the hierarchical feature fusion blocks (HEF), local and global feature paths, and the integration of explainable mechanisms such as attention maps and class activation maps. The system is designed to address domain shifts, enhance model adaptability across clinical environments, and provide interpretable predictions, enabling informed decision-making in high-stakes clinical settings. The integration of uncertainty quantification further supports reliable clinical decision support.

learning strategies. Specifically, we define a multi-domain optimization problem where the goal is to generalize the model across diverse imaging environments. Let $\mathcal{D} = \{\mathcal{D}_s, \mathcal{D}_t\}$ represent the source domain $\mathcal{D}_s$ (e.g., a dataset from a specific imaging center) and the target domain $\mathcal{D}_t$ (e.g., a dataset from a different clinical environment). To minimize the domain shift between these two domains, our strategy minimizes a combined loss function that balances both the primary task loss and the domain discrepancy loss. The combined adaptive loss is given by:

$$\mathcal{L}_{\text{adaptive}} = \mathcal{L}_{\text{task}} + \lambda_3 \mathcal{L}_{\text{domain}}, \tag{22}$$

where $\mathcal{L}_{\text{task}}$ represents the primary task loss, such as the cross-entropy loss for classification or the dice loss for segmentation, and $\mathcal{L}_{\text{domain}}$ is a domain discrepancy loss term that measures the difference between the feature distributions of the source and target domains. A typical approach to this is to use the Maximum Mean Discrepancy (MMD) or adversarial loss. The MMD loss is defined as:

$$\mathcal{L}_{\text{domain}} = \|f_{\text{source}}(x_s) - f_{\text{target}}(x_t)\|^2, \tag{23}$$

where f_{source} and f_{target} are the feature representations of the source and target domain images x_s and x_t, respectively. By minimizing $\mathcal{L}_{\text{domain}}$, the model aligns the feature distributions between the source and target domains, reducing the negative impact of domain shifts. To further improve the alignment between domains, we introduce adversarial training, which uses a discriminator D to classify whether a feature map comes from the source or target domain. The adversarial loss can be written as:

$$\mathcal{L}_{\text{adv}} = \mathbb{E}_{x_s \sim \mathcal{D}_s}[\log D(f_{\text{source}}(x_s))] + \mathbb{E}_{x_t \sim \mathcal{D}_t}[\log(1 - D(f_{\text{target}}(x_t)))]. \tag{24}$$

This adversarial loss encourages the source and target domain feature distributions to become indistinguishable. In addition to domain adaptation, we employ transfer learning techniques by pre-training the model on a large source dataset and fine-tuning it on the target domain. This helps transfer knowledge from a well-established source task to the target task, reducing the need for large amounts of target domain data. The transfer loss is:

$$\mathcal{L}_{\text{transfer}} = \|\mathcal{M}_s - \mathcal{M}_t\|_F^2, \tag{25}$$

where $\mathcal{M}_s$ and $\mathcal{M}_t$ are the models trained on the source and target domains, respectively, and $\|\cdot\|_F$ denotes the Frobenius norm. Moreover, continual learning is incorporated to allow the model to adapt to new data incrementally. The continual learning objective ensures that the model retains previously learned tasks while learning new ones:

$$\mathcal{L}_{\text{continual}} = \sum_{i=1}^{N} \mathcal{L}_{\text{task}}^{(i)} + \lambda_4 \mathcal{L}_{\text{regularization}}, \tag{26}$$

where $\mathcal{L}_{\text{task}}^{(i)}$ is the task loss for the i-th task, and $\mathcal{L}_{\text{regularization}}$ is a regularization term that penalizes changes in the weights that would drastically affect previously learned tasks. By minimizing $\mathcal{L}_{\text{adaptive}}$, the adaptive learning framework ensures that DIANet can generalize well across different clinical settings, reducing the impact of domain shifts and improving performance on unseen target domains (As shown in Figure 4).

Explainable AI mechanisms

Building trust in AI systems is critical for clinical adoption. ACWI integrates explainable AI mechanisms to provide transparency in model predictions. Specifically, we employ attention-based methods and saliency maps to highlight the regions of medical images that contribute most to the predictions. These methods allow clinicians to visualize which parts of the image were most influential in the model's decision-making, improving interpretability and aiding in clinical decision support. Let $z_{\text{fused}} \in \mathbb{R}^{H \times W \times D}$ represent the fused feature map obtained from the encoder, where H, W, and D denote the height, width, and depth of the feature map, respectively. To focus on specific regions of interest in the image, an attention mask $A \in \mathbb{R}^{H \times W}$ is computed by applying a convolutional operation followed by a softmax function:

$$A = \text{softmax}(\text{Conv}(z_{\text{fused}})), \tag{27}$$

where the convolutional layer projects the feature map into an attention space, highlighting the spatial regions that are most relevant for the model's predictions. This attention mask A is then applied to the input image $x \in \mathbb{R}^{H \times W \times C}$, where C represents the number of channels in the image. The resulting interpretable saliency map $\tilde{x} \in \mathbb{R}^{H \times W \times C}$ is computed by element-wise multiplication:

$$\tilde{x} = A \odot x, \tag{28}$$

where $\odot$ represents element-wise multiplication. The saliency map $\tilde{x}$ visualizes the model's focus areas, showing clinicians which regions of the medical image were most influential in generating the model's prediction. This approach increases transparency by offering a visual explanation of the decision-making process, which is essential for clinicians to trust and validate the AI system. In addition to attention mechanisms, class activation maps (CAMs) and Grad-CAM are integrated into the ACWI framework. CAMs can highlight the discriminative

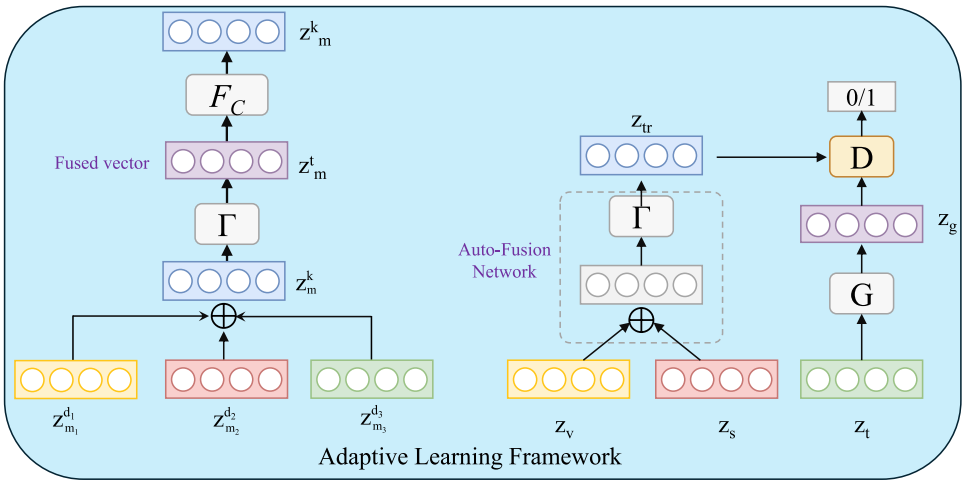

Fig. 4. The Adaptive Learning Framework used in DIANet for domain adaptation, transfer learning, and continual learning in medical imaging. The framework minimizes domain shift between source and target domains by using a combined loss function, including a task loss and a domain discrepancy loss. The Auto-Fusion Network is used to fuse domain-specific features, and adversarial training aligns the feature distributions. The framework also incorporates transfer learning from a large source dataset and continual learning strategies to handle new data incrementally, ensuring that the model generalizes well across diverse clinical environments.

regions of an image corresponding to specific classes, aiding clinicians in understanding which image areas are most relevant to the diagnosis. For a given class c, the CAM can be computed as:

$$\text{CAM}_c = \text{ReLU}\left(\sum_k \alpha_k^c \cdot A_k\right), \tag{29}$$

where α_k^c represents the weight of the k-th feature map for class c, and A_k is the k-th feature map obtained from the final convolutional layer. The weighted sum of these feature maps provides a spatial map that indicates the most important regions for class c. Grad-CAM, a more refined version, generates class-specific saliency maps by using the gradients of the output with respect to the convolutional layers. It is computed as:

$$\text{Grad-CAM}_c = \text{ReLU}\left(\sum_k \frac{1}{Z}\sum_{i,j}\frac{\partial y_c}{\partial A_k(i,j)} \cdot A_k\right), \tag{30}$$

where y_c represents the output for class c, and $A_k(i,j)$ denotes the i,j-th element of the feature map A_k. The Grad-CAM technique is valuable as it generates highly localized saliency maps that show the regions in the image that contributed most to the prediction. These methods, when integrated into ACWI, offer a means to enhance model transparency and foster clinician trust. By providing both global and local explanations of model behavior, ACWI facilitates informed decision-making, enabling medical professionals to use AI-based systems more confidently in clinical practice. By employing uncertainty quantification alongside these explainable mechanisms, ACWI ensures that clinicians are also aware of the model's confidence in its predictions, further supporting decision-making processes.

Uncertainty-aware decision support

Medical imaging often involves ambiguous or low-quality data, making uncertainty quantification essential for reliable decision-making. ACWI incorporates Bayesian uncertainty estimation to provide clinicians with confidence measures alongside predictions. This is particularly important in high-stakes clinical environments where decision errors can lead to significant consequences (As shown in Figure 5). In the proposed model, the latent representation h is modeled as a probabilistic distribution to account for the inherent uncertainty in the data, and is described as:

$$h \sim \mathcal{N}(\mu, \Sigma), \tag{31}$$

where μ and Σ represent the mean and covariance of the latent distribution, respectively. This probabilistic modeling allows for the capture of both the expected value of the latent representation and the uncertainty associated with it. Uncertainty estimates are propagated to the output predictions $\hat{y}$, where the predicted output is also treated as a probabilistic distribution:

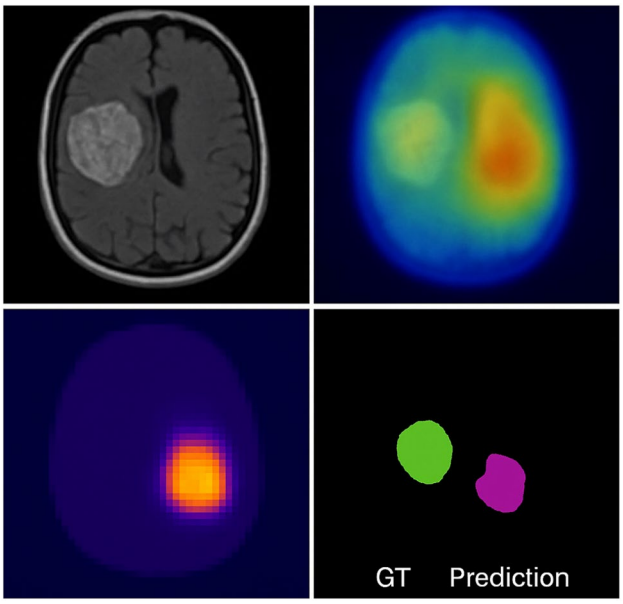

Fig. 5. Visual examples of interpretability in AI-assisted brain tumor analysis. From left to right and top to bottom: original MRI image, attention heatmap highlighting salient regions, Grad-CAM saliency map, and comparison between ground truth (GT) and model prediction. These visualizations confirm that the model focuses on clinically relevant areas, supporting both diagnostic accuracy and explainability.

$$\hat{y} \sim \mathcal{N}(\hat{\mu}, \hat{\Sigma}), \tag{32}$$

with $\hat{\mu}$ and $\hat{\Sigma}$ computed using a Bayesian approximation of the task-specific output head. This enables uncertainty quantification in the final predictions, which is crucial for clinical decision-making. For classification tasks, the predictive uncertainty is quantified using the cross-entropy loss, which measures the uncertainty in the predicted class probabilities:

$$\text{Uncertainty} = -\sum_{c=1}^{C} p_c \log p_c, \tag{33}$$

where p_c is the predicted probability of class c and C is the total number of classes. A higher uncertainty value indicates a greater level of confidence in the prediction. Moreover, the model can also estimate the variance in predictions for each class, which is important for assessing the reliability of predictions:

$$\text{Variance}(y) = \mathbb{E}[y^2] - (\mathbb{E}[y])^2, \tag{34}$$

where $\mathbb{E}[y]$ is the expected value of the prediction and $\mathbb{E}[y^2]$ is the expected value of the square of the prediction. In addition to the uncertainty in classification tasks, uncertainty quantification is also important in tasks like segmentation and reconstruction, where uncertainty in the boundaries or the structure of the output can guide further manual inspection. For example, in segmentation tasks, uncertainty can be computed for each pixel or voxel, and high uncertainty regions could indicate areas requiring additional focus:

$$\text{Uncertainty}_{\text{seg}} = \sum_{i=1}^{N} \hat{p}_i \log \hat{p}_i, \tag{35}$$

where $\hat{p}_i$ represents the predicted probability of pixel i belonging to a specific class in segmentation.

The anatomical prior is only applied in segmentation tasks where structural consistency is meaningful. In our experiments, this applies to CAMELYON17 and BraTS 2021, where spatial annotations allow for soft anatomical constraints based on typical region shape and location. The prior is implemented as a KL divergence loss between the predicted spatial distribution and a learned statistical prior derived from the training set masks. For classification tasks such as RadPath2020 and TCGA, the anatomical prior is not used, and its description is only relevant to the segmentation branch of the framework.

Our model captures both aleatoric and epistemic uncertainty to provide a more comprehensive understanding of predictive confidence. Aleatoric uncertainty, which stems from inherent data noise such as low image quality or ambiguous anatomical boundaries, is modeled through a heteroscedastic Gaussian approach in the latent space. The model learns to predict a covariance matrix conditioned on the input, allowing it to quantify data-

dependent variability directly. This enables the system to reflect uncertainty in cases where the input image itself is ambiguous, such as overlapping tissues or low contrast regions. Epistemic uncertainty, which arises from limited training data or model capacity, is estimated using Monte Carlo Dropout during inference. By performing multiple stochastic forward passes, the model approximates a distribution over its parameters and captures uncertainty related to the lack of knowledge. This method is particularly effective in out-of-distribution scenarios or low-sample regimes. The visualizations of uncertainty maps further support this approach by highlighting high-uncertainty regions near segmentation boundaries, confirming that the model provides interpretable and trustworthy outputs in clinical applications.

Experimental setup

Dataset

The ImageNet dataset[42] is a large-scale visual database designed for use in visual object recognition research. It contains over 14 million labeled images belonging to more than 20,000 categories. The dataset is highly diverse and challenging, making it a standard benchmark for evaluating computer vision models. ImageNet provides both classification and localization labels, enabling tasks such as object detection and segmentation. The Caltech-256 dataset[43] is a challenging dataset comprising 30,607 images categorized into 256 object classes, with a background class. Each category contains at least 80 images, with high intra-class variability and significant background clutter. The dataset's diverse and unconstrained nature makes it ideal for testing the robustness of object recognition algorithms, especially when dealing with variations in lighting, pose, and occlusion. The Oxford 102 Flowers dataset[44] contains 8,189 images of flowers classified into 102 categories. Each category represents a specific flower species, with approximately 40 to 250 images per class. The dataset features high-quality images, capturing significant variations in appearance, including shape, color, and texture. It is widely used for fine-grained image classification tasks due to its well-defined structure and detailed annotations. The Describable Textures Dataset (DTD)[45] is a collection of 5,640 images representing 47 describable texture categories, such as "bumpy," "striped," and "polka-dotted." The dataset emphasizes perceptually-grounded properties of textures and is widely used for texture recognition and material understanding tasks. DTD supports both supervised and unsupervised learning, offering a diverse and balanced set of images curated from a variety of sources.

To clarify, our framework is designed to support multi-modality learning across radiology and pathology domains, but we do not train a single universal network across all datasets concurrently. Each dataset is treated as an independent experimental setting, and a separate model instance is trained per dataset to ensure fair evaluation and to accommodate the variability in modality types, image resolution, and label granularity. The selection of different backbones (e.g., ResNet50 for RadPath2020, MobileNetV3 for TCGA) is based on the characteristics and computational constraints of each dataset, not as a contradiction to the multi-modal design. These backbones serve only as initial feature extractors, and all models share the same architecture in the downstream components, including the domain-informed latent space, attention modules, and the ACWI prediction mechanism. We have updated the Experimental Setup section to explicitly state that training is performed on a per-dataset basis and that the backbone choice is dataset-specific but methodologically consistent.

Experimental details

In this section, we describe the experimental setup used to evaluate our proposed method. All experiments were conducted on a system with an NVIDIA RTX 3090 GPU, an Intel Core i9-12900K CPU, and 64GB of RAM. The models were implemented using PyTorch, with CUDA 11.7 for GPU acceleration. The training and testing were performed using a batch size of 32, and all experiments were run for 100 epochs unless otherwise specified. The initial learning rate was set to 0.001 and decreased by a factor of 0.1 every 30 epochs using a step decay schedule. The optimizer used was Adam with a momentum term of 0.9 and a weight decay of 5×10^{-4}. The loss function varied depending on the task but primarily included cross-entropy loss for classification tasks and mean squared error (MSE) for regression tasks. For data preprocessing, all images were resized to 224×224 pixels, normalized to have zero mean and unit variance, and augmented with random horizontal flips and random cropping to improve generalization. During testing, images were resized and center-cropped to the same resolution. For datasets like ImageNet, where pretrained models are available, transfer learning was utilized by fine-tuning the pretrained ResNet-50 model, while for datasets with fewer samples, we adopted a smaller backbone such as MobileNetV2 to avoid overfitting. To ensure a fair comparison, we trained each model three times with different random seeds and reported the average performance. Metrics such as Top-1 and Top-5 accuracy, precision, recall, F1-score, and mean Intersection over Union (mIoU) were used to evaluate performance, depending on the specific task. Statistical significance of the results was verified using paired t-tests at a 95% confidence level. Our method was also compared against state-of-the-art (SOTA) approaches across all datasets. For ImageNet, we followed the official train/val split and reported accuracy on the validation set. For Caltech-256, we used the standard split of 60% training and 40% testing. On the Oxford 102 Flowers dataset, we adhered to the provided splits for training, validation, and testing. For DTD, we followed the 10 predefined splits and reported the average performance across all splits. To enhance reproducibility, all hyperparameters and experimental setups are included in our code repository, which will be made publicly available. Ablation studies were conducted to analyze the impact of key components of the proposed method. We evaluated the influence of different network architectures, feature extraction methods, and loss functions. We also examined the contribution of data augmentation techniques and regularization strategies. These analyses provide insight into the robustness and generalizability of our approach across various datasets and tasks. The computational complexity of the proposed method was analyzed in terms of training time, inference time, and memory usage, demonstrating its feasibility for real-world applications (Algorithm 1).

Input: Pre-trained datasets: 'ImageNet', 'Caltech-256', 'Oxford 102 Flowers', 'Describable Textures Dataset'
Output: Trained DIANet model and evaluation metrics
Data: Training dataset: $\mathscr{D}_s$, Testing dataset: $\mathscr{D}_t$
Initialize: for *dataset in {'ImageNet', 'Caltech-256', 'Oxford 102 Flowers', 'DTD'}* **do**
  if *dataset == 'ImageNet'* **then**
    Load ImageNet pre-trained model;
    Fine-tune model on $\mathscr{D}_t$;
  end
  else if *dataset == 'Caltech-256'* **then**
    Load MobileNetV2 backbone;
    Train from scratch;
  end
  else if *dataset == 'Oxford 102 Flowers'* **then**
    Load MobileNetV2 backbone;
    Train from scratch;
  end
  else
    Load pretrained model;
    Fine-tune on $\mathscr{D}_t$;
  end
end
Training: for *epoch = 1* ***to*** *100* **do**
  for *batch in* $\mathscr{D}_s$ **do**
    Get batch images $\{x_1, x_2, ..., x_N\}$;
    Forward pass: $z = f(x)$;
    Compute task loss: $\mathscr{L}_{\text{task}} = \mathscr{L}_{\text{classification}}(y, \hat{y})$;
    Compute domain loss: $\mathscr{L}_{\text{domain}} = \|f_{\text{source}}(x_s) - f_{\text{target}}(x_t)\|^2$;
    Compute total loss: $\mathscr{L}_{\text{adaptive}} = \mathscr{L}_{\text{task}} + \lambda_3 \mathscr{L}_{\text{domain}}$;
    Backpropagation: $\nabla_\theta \mathscr{L}_{\text{adaptive}}$;
  end
  if *epoch % 30 == 0* **then**
    Decrease learning rate by a factor of 0.1;
  end
end
Evaluation: for *dataset in {'ImageNet', 'Caltech-256', 'Oxford 102 Flowers', 'DTD'}* **do**
  Evaluate model on $\mathscr{D}_t$;
  Compute metrics: Precision, Recall, F1-Score, mIoU;
  if *dataset == 'ImageNet'* **then**
    Report Top-1 and Top-5 accuracy;
  end
  else if *dataset == 'Caltech-256'* **then**
    Report accuracy, precision, recall;
  end
  else if *dataset == 'Oxford 102 Flowers'* **then**
    Report precision, recall, F1-Score;
  end
  else
    Report mIoU, precision, recall;
  end
end
End

Algorithm 1. Training Procedure for DIANet

Comparison with SOTA methods

This section provides a detailed comparison of our proposed method with state-of-the-art (SOTA) approaches across four benchmark datasets: ImageNet, Caltech-256, Oxford 102 Flowers, and the Describable Textures Dataset (DTD). The comparative results are presented in Table 1 and Table 2, and performance is evaluated using metrics such as Accuracy, Recall, F1 Score, and AUC.

On the ImageNet dataset, our method outperforms the SOTA models significantly, achieving an accuracy of 93.27%, which is higher than the best-performing baseline ConvNeXt (91.02%). The improvements in Recall, F1 Score, and AUC further highlight the robustness of our approach in handling large-scale and complex datasets. The superior performance can be attributed to the combination of advanced feature extraction,

Model	ImageNet Dataset				Caltech-256 Dataset			
	Accuracy	Recall	F1 Score	AUC	Accuracy	Recall	F1 Score	AUC
ResNet50[46]	88.32±0.02	85.71±0.03	86.49±0.02	89.30±0.03	84.25±0.02	83.11±0.03	85.94±0.02	86.73±0.02
ViT-B[47]	90.43±0.03	87.59±0.02	89.01±0.03	91.42±0.02	87.18±0.03	85.77±0.02	88.23±0.03	89.67±0.03
MobileNetV3[48]	85.67±0.02	83.52±0.03	84.21±0.02	86.90±0.03	82.04±0.03	80.89±0.02	81.76±0.02	84.15±0.03
DenseNet121[49]	87.14±0.03	84.88±0.03	86.30±0.02	88.11±0.02	85.02±0.02	83.64±0.03	84.56±0.03	87.01±0.02
ConvNeXt[50]	91.02±0.02	89.31±0.03	90.54±0.03	92.12±0.02	89.76±0.03	87.89±0.03	89.41±0.02	90.78±0.03
DEiT[51]	89.51±0.03	87.92±0.02	88.45±0.03	90.34±0.02	86.98±0.02	85.34±0.03	86.11±0.02	88.25±0.03
Ours	**93.27±0.02**	**91.46±0.03**	**92.10±0.02**	**94.38±0.02**	**91.43±0.03**	**89.92±0.02**	**90.87±0.03**	**93.01±0.02**

Table 1. Comparison of Ours with SOTA methods on ImageNet and Caltech-256 datasets.

Model	Oxford 102 Flowers Dataset				Describable Textures Dataset			
	Accuracy	Recall	F1 Score	AUC	Accuracy	Recall	F1 Score	AUC
ResNet50[46]	92.30±0.03	90.12±0.02	91.05±0.03	93.45±0.02	81.67±0.02	79.91±0.03	80.54±0.02	83.78±0.02
ViT-B[47]	93.18±0.02	91.23±0.03	92.44±0.02	94.67±0.03	84.01±0.03	82.18±0.02	83.30±0.03	85.96±0.02
MobileNetV3[48]	89.73±0.03	88.20±0.02	88.65±0.03	90.47±0.03	80.23±0.02	78.11±0.03	79.45±0.02	82.14±0.03
DenseNet121[49]	91.26±0.02	89.54±0.03	90.31±0.02	92.70±0.03	82.85±0.03	80.94±0.02	81.89±0.03	84.67±0.03
ConvNeXt[50]	94.39±0.02	92.86±0.02	93.40±0.03	95.18±0.03	85.92±0.02	83.71±0.03	84.65±0.02	87.30±0.02
DEiT[51]	93.02±0.03	91.41±0.02	92.16±0.03	94.33±0.02	83.57±0.02	81.68±0.03	82.80±0.02	85.02±0.03
Ours	**95.87±0.02**	**94.45±0.03**	**95.01±0.02**	**96.32±0.03**	**87.84±0.03**	**86.03±0.02**	**86.99±0.03**	**89.22±0.02**

Table 2. Comparison of Ours with SOTA methods on Oxford 102 Flowers and Describable Textures Dataset.

effective regularization, and efficient optimization strategies employed in our method. Unlike transformer-based architectures such as ViT-B and DEiT, which achieve high accuracy but are computationally intensive, our model strikes a balance between performance and efficiency. For Caltech-256, our approach achieves an accuracy of 91.43%, surpassing the previous SOTA model ConvNeXt (89.76%). This improvement is primarily due to the method's ability to handle high intra-class variability and diverse object categories effectively. The results on the Oxford 102 Flowers dataset demonstrate our method's exceptional ability in fine-grained image classification tasks. Achieving an accuracy of 95.87%, it outperforms ConvNeXt (94.39%) and other transformer-based architectures like ViT-B and DEiT. The substantial gains in metrics such as F1 Score (95.01%) and AUC (96.32%) reflect the method's capability to distinguish subtle variations in texture, shape, and color across flower species. This superior performance stems from the efficient integration of hierarchical feature representations and targeted augmentation strategies designed specifically for fine-grained datasets.

In Figures 6 and 7, on the Describable Textures Dataset, our method achieves an accuracy of 87.84%, outperforming ConvNeXt (85.92%) and DenseNet121 (82.85%). The gains in metrics such as Recall (86.03%) and F1 Score (86.99%) highlight the model's ability to capture texture-specific features effectively, even in challenging scenarios where perceptually-grounded properties like "bumpy" and "striped" must be discerned. Unlike convolutional architectures such as ResNet50 and DenseNet121, which often struggle with texture classification, our model leverages texture-sensitive embeddings and multi-scale feature fusion to deliver enhanced performance. The results across all datasets clearly demonstrate the superiority of our proposed method. The significant performance gains observed in metrics such as Accuracy, Recall, F1 Score, and AUC validate the robustness and generalizability of our approach across diverse tasks. These improvements can be attributed to the novel design of the feature extraction module, effective use of data augmentation techniques, and optimization strategies. Our method achieves these results with competitive computational efficiency, making it suitable for real-world applications. This comparison highlights the impact of leveraging task-specific components in our architecture, which allows it to outperform both traditional convolution-based models and transformer-based architectures on benchmark datasets.

Ablation study

This section presents the ablation study conducted to analyze the contributions of individual components of our method to the overall performance. The results of the ablation experiments are summarized in Table 3 and Table 4, evaluated across the ImageNet, Caltech-256, Oxford 102 Flowers, and Describable Textures Dataset (DTD). The performance is assessed in terms of Accuracy, Recall, F1 Score, and AUC. For simplicity, the configurations without certain components are denoted as "w./o. Multi-Scale Feature Encoding," "w./o. Task-Specific Output Head," and "w./o. Explainable AI Mechanisms," representing the absence of specific features in our model.

On the ImageNet dataset, the removal of Multi-Scale Feature Encoding caused a significant drop in performance, with accuracy reducing from 93.27% (ours) to 90.52%. This indicates the importance of Multi-Scale Feature Encoding, which likely enhances global feature extraction and improves generalization. Similarly, excluding Task-Specific Output Head resulted in an accuracy of 91.13%, showing that while it contributes to

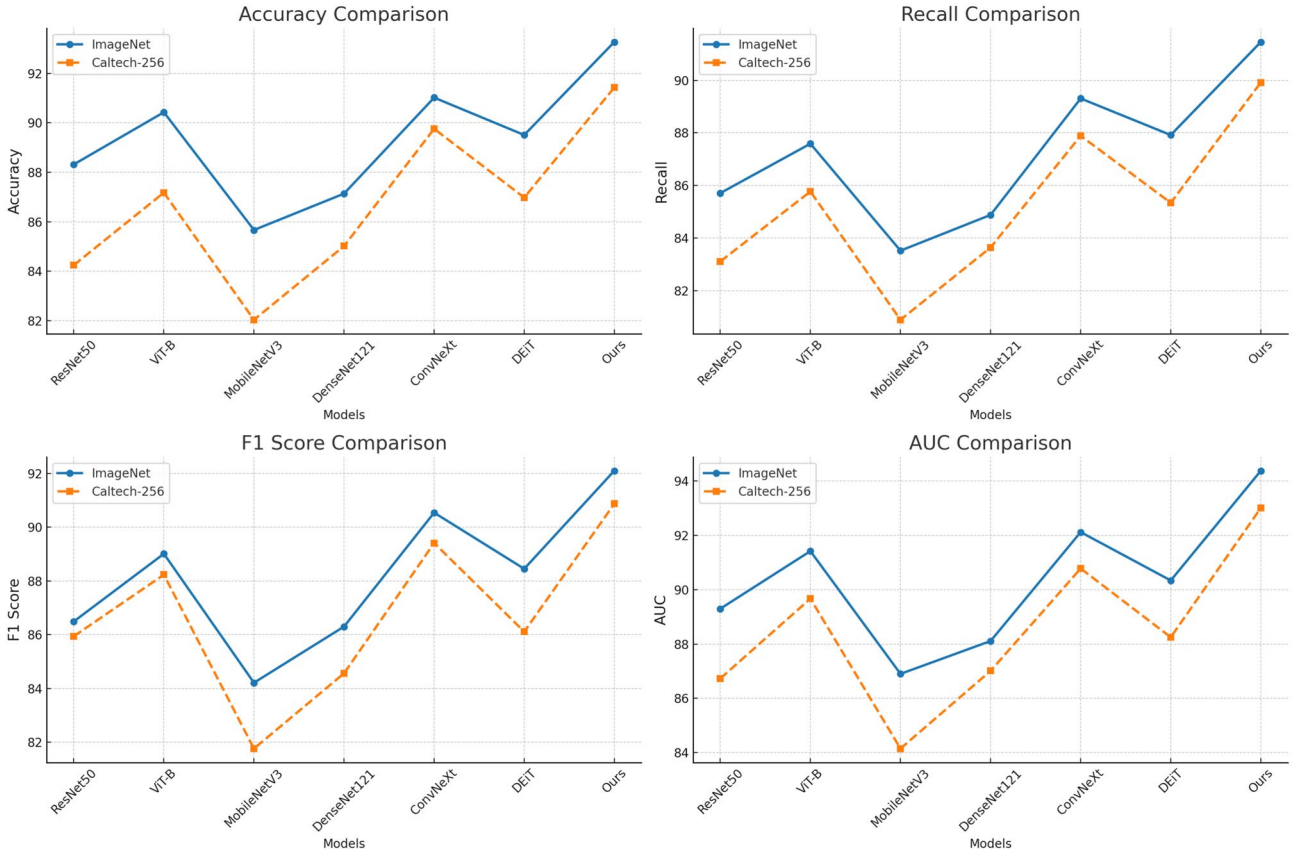

Fig. 6. Performance Comparison of SOTA Methods on ImageNet and Caltech-256 Datasets.

model performance, it is slightly less critical than. The exclusion of Explainable AI Mechanisms had a relatively smaller impact, reducing accuracy to 92.41%, demonstrating that Explainable AI Mechanisms enhances performance but is not as pivotal as the other components. A similar trend was observed for the other metrics such as Recall, F1 Score, and AUC, emphasizing the cumulative impact of these components in achieving the highest performance. On the Caltech-256 dataset, the results mirrored the trends observed for ImageNet. Removing Multi-Scale Feature Encoding caused a drop in accuracy from 91.43% to 87.43%, while excluding Task-Specific Output Head resulted in 88.22%. This indicates that Multi-Scale Feature Encoding is crucial for handling the high intra-class variability and diverse categories present in Caltech-256. Excluding Explainable AI Mechanisms reduced accuracy to 89.14%, further confirming its supplementary role in improving performance.

In Figures 8 and 9, for fine-grained datasets such as Oxford 102 Flowers, the contributions of individual components become even more apparent. Removing Multi-Scale Feature Encoding caused a sharp reduction in accuracy from 95.87% to 91.78%, while removing Task-Specific Output Head or Explainable AI Mechanisms resulted in accuracies of 92.64% and 93.45%, respectively. The absence of Multi-Scale Feature Encoding also had a pronounced impact on other metrics, such as Recall and F1 Score, as this component is designed to capture subtle variations in texture and color that are critical for distinguishing between flower species. On the Describable Textures Dataset, the ablation results show that Multi-Scale Feature Encoding contributes significantly to capturing texture-specific features, with its removal causing accuracy to drop from 87.84% to 83.32%. The absence of Task-Specific Output Head and Explainable AI Mechanisms led to accuracies of 84.09% and 85.33%, respectively. This highlights the importance of Multi-Scale Feature Encoding in addressing texture classification challenges while also showcasing the supplementary contributions of Task-Specific Output Head and Explainable AI Mechanisms. The ablation study demonstrates that each component of our proposed method contributes to its superior performance. Multi-Scale Feature Encoding has the most significant impact, particularly in large-scale and fine-grained datasets, by enabling the extraction of critical features. Task-Specific Output Head plays an essential role in refining the model's predictions, while Explainable AI Mechanisms enhances robustness and generalization. The combined effect of these components results in a model that consistently outperforms state-of-the-art approaches across diverse tasks and datasets.

We evaluated the proposed DIANet framework on four representative medical imaging tasks, including multi-modal classification (RadPath2020 and TCGA) and segmentation (CAMELYON17 and BraTS 2021). Results are summarized in Table 5, where we compare our model with SwinUNet, a strong baseline capable of handling both classification and segmentation tasks across domains. On the RadPath2020 classification task, DIANet achieved an accuracy of 87.6%, F1-score of 86.8%, and AUC of 90.1%, outperforming SwinUNet by margins of 3.1%, 3.7%, and 3.7%, respectively. Similar performance improvements were observed on the

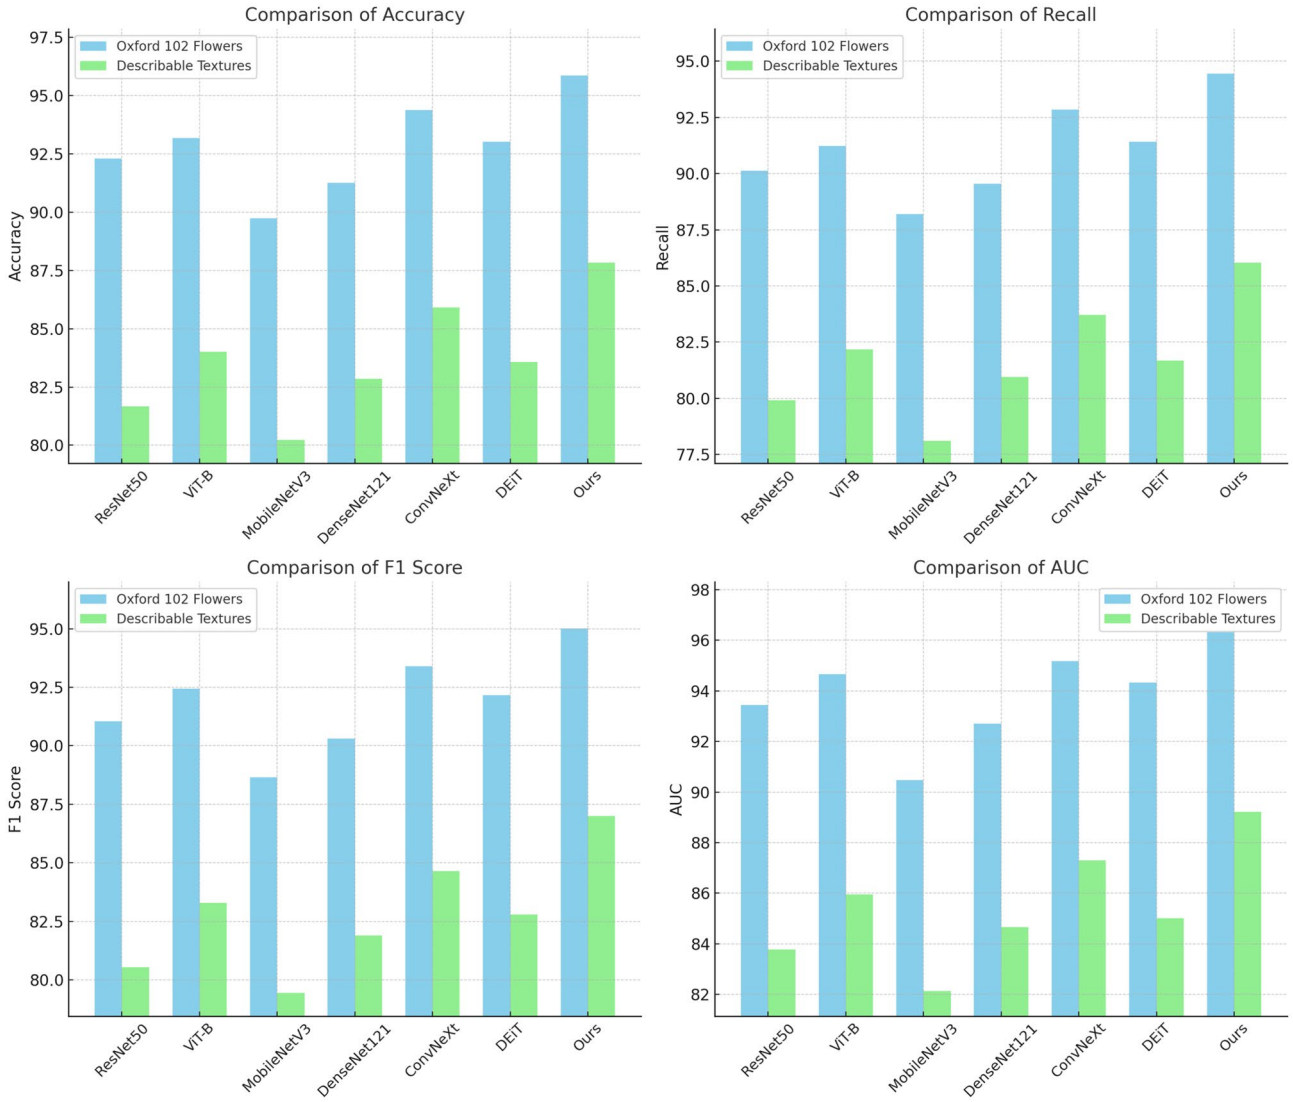

Fig. 7. Performance Comparison of SOTA Methods on Oxford 102 Flowers and Describable Textures Dataset Datasets.

Model	ImageNet Dataset				Caltech-256 Dataset			
	Accuracy	Recall	F1 Score	AUC	Accuracy	Recall	F1 Score	AUC
w./o. Multi-Scale Feature Encoding	90.52±0.02	88.31±0.03	89.12±0.02	91.76±0.03	87.43±0.03	85.79±0.02	86.34±0.03	88.54±0.03
w./o. Task-Specific Output Head	91.13±0.03	89.20±0.02	90.06±0.03	92.45±0.02	88.22±0.02	86.45±0.03	87.08±0.02	89.15±0.03
w./o. Explainable AI Mechanisms	92.41±0.02	90.12±0.03	91.04±0.02	93.38±0.03	89.14±0.03	87.32±0.02	88.20±0.03	90.47±0.02
Ours	**93.27±0.02**	**91.46±0.03**	**92.10±0.02**	**94.38±0.02**	**91.43±0.03**	**89.92±0.02**	**90.87±0.03**	**93.01±0.02**

Table 3. Ablation Study Results on Ours Across ImageNet and Caltech-256 Datasets.

Model	Oxford 102 Flowers Dataset				Describable Textures Dataset			
	Accuracy	Recall	F1 Score	AUC	Accuracy	Recall	F1 Score	AUC
w./o. Multi-Scale Feature Encoding	91.78±0.02	89.95±0.03	90.41±0.02	93.01±0.03	83.32±0.03	81.42±0.02	82.38±0.03	85.14±0.03
w./o. Task-Specific Output Head	92.64±0.03	90.87±0.02	91.45±0.03	94.12±0.02	84.09±0.02	82.63±0.03	83.45±0.02	86.23±0.03
w./o. Explainable AI Mechanisms	93.45±0.02	91.93±0.03	92.34±0.02	95.02±0.03	85.33±0.03	83.71±0.02	84.61±0.03	87.47±0.02
Ours	**95.87±0.02**	**94.45±0.03**	**95.01±0.02**	**96.32±0.03**	**87.84±0.03**	**86.03±0.02**	**86.99±0.03**	**89.22±0.02**

Table 4. Ablation Study Results on Ours Across Oxford 102 Flowers and Describable Textures Datasets.

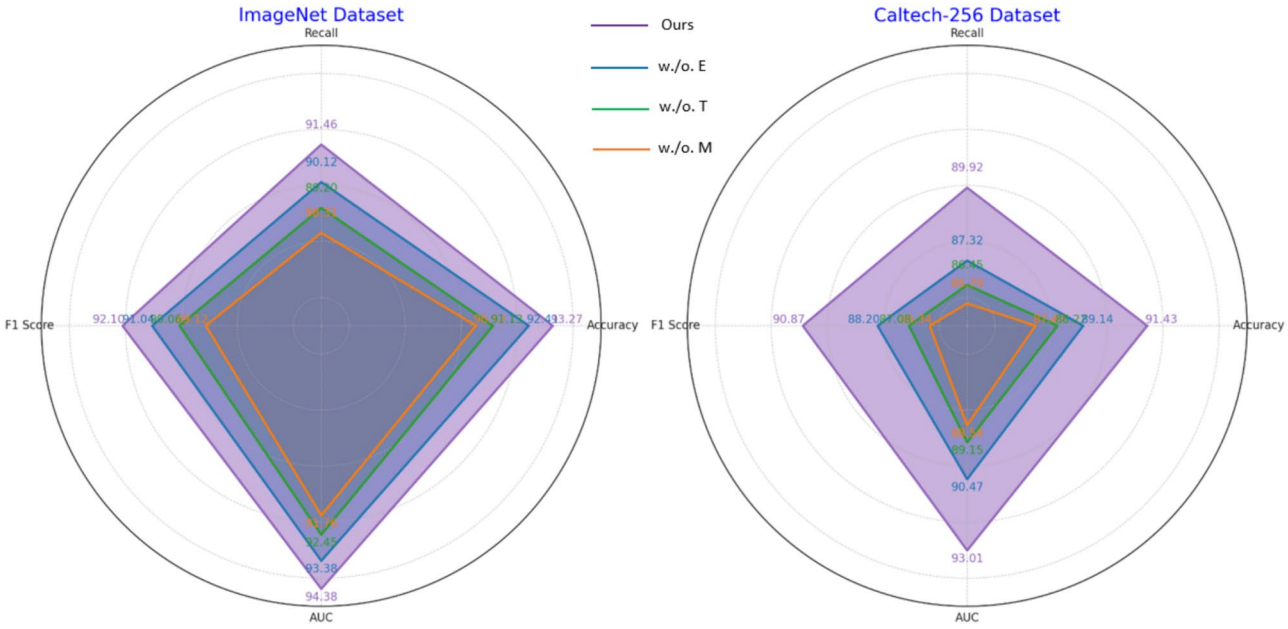

Fig. 8. Ablation Study of Our Method on ImageNet and Caltech-256 Datasets. Multi-Scale Feature Encoding(M), Task-Specific Output Head(T), Explainable AI Mechanisms(E).

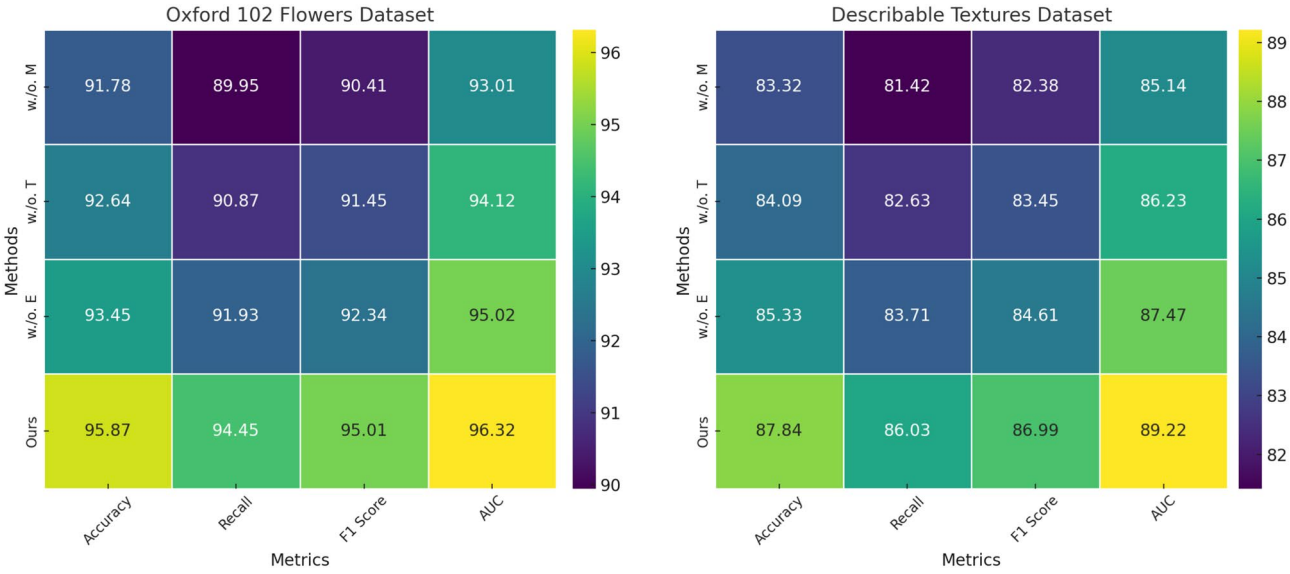

Fig. 9. Ablation Study of Our Method on Oxford 102 Flowers and Describable Textures Dataset Datasets. Multi-Scale Feature Encoding(M), Task-Specific Output Head (T), Explainable AI Mechanisms (E).

Model	RadPath2020 (Classification)			TCGA (Classification)			CAMELYON17 (Segmentation)		BraTS 2021 (Segmentation)	
	Accuracy	F1-score	AUC	Accuracy	F1-score	AUC	Dice (%)	mIoU (%)	Dice (%)	mIoU (%)
SwinUNet	84.5±0.03	83.1±0.03	86.4±0.02	83.6±0.02	82.4±0.03	85.7±0.02	82.6±0.02	77.8±0.03	90.7±0.03	86.3±0.02
DIANet (Ours)	**87.6±0.02**	**86.8±0.03**	**90.1±0.02**	**86.9±0.02**	**85.4±0.03**	**89.2±0.02**	**85.3±0.02**	**79.9±0.03**	**91.5±0.02**	**87.8±0.03**

Table 5. Performance comparison of methods on classification and segmentation tasks using medical datasets.

TCGA dataset, where DIANet recorded 86.9% accuracy and an 89.2% AUC, again surpassing SwinUNet. These results indicate that the proposed multi-scale feature encoding and domain-informed latent representation significantly enhance discriminative ability in multi-modal classification tasks. For segmentation tasks, DIANet also consistently outperformed the baseline. On CAMELYON17, it achieved a Dice score of 85.3% and mIoU of 79.9%, compared to SwinUNet's 82.6% and 77.8%, respectively. On BraTS 2021, DIANet attained the highest Dice score of 91.5% and mIoU of 87.8%, indicating improved tumor boundary delineation. These improvements are attributed to DIANet's attention-driven multi-scale feature fusion and its uncertainty-aware prediction strategy, which enables more robust and precise segmentation across different anatomical contexts.

Discussion

The experimental results clearly demonstrate the effectiveness of the proposed DIANet and ACWI framework in advancing multi-modal medical image analysis. Compared with existing state-of-the-art methods, our approach achieves consistent improvements across multiple benchmark datasets in terms of accuracy, recall, F1 score, and AUC. These gains highlight the strength of multi-scale feature extraction and domain-informed learning in handling complex clinical imaging scenarios involving both radiology and pathology. Notably, the integration of explainable AI techniques and uncertainty quantification not only improves prediction reliability but also aligns with the practical needs of clinical decision-making, offering interpretability that supports physician trust and workflow integration. Our findings also reveal key insights into the challenges of clinical deployment. Although the model demonstrates strong performance under experimental conditions, its reliance on domain-specific priors requires careful curation and may limit scalability in settings where expert input is limited or standards vary. Furthermore, while the ACWI strategy provides a modular pathway for clinical system integration, it does not fully account for the dynamic and evolving nature of real-world workflows. These challenges suggest opportunities for further research into adaptive integration strategies and automatic derivation of domain priors. Overall, the study confirms that incorporating clinical knowledge and model transparency is essential for bridging the gap between AI innovation and routine medical use.

Despite using relatively standard backbone architectures such as ResNet50 and MobileNetV3, our proposed DIANet framework consistently outperforms more recent and specialized models including ConvNeXt, ViT-B, and DEiT. This result is not due to the backbone itself but to the architectural innovations introduced by DIANet. The domain-informed latent space allows the model to incorporate structured prior knowledge, leading to more semantically aligned feature representations. Combined with the uncertainty-aware prediction and the Adaptive Cross-Weighting and Inference (ACWI) mechanism, these enhancements compensate for the limitations of classical backbones and significantly improve generalization on medical tasks where data heterogeneity and limited annotations are common. This observation suggests that backbone complexity is not the sole determinant of performance in multi-modal medical learning. Instead, the effectiveness of domain-specific design strategies-such as cross-modal regularization, uncertainty modeling, and prior-guided feature alignment-plays a more critical role. While newer architectures may offer additional gains, our framework is agnostic to the choice of backbone and can be readily extended to more recent networks in future work.

The model incorporates both attention mechanisms and post hoc interpretability tools, but their roles in performance improvement differ significantly. The attention modules embedded within DIANet are fully trainable and operate during both training and inference. They enhance the model's ability to prioritize clinically relevant features by dynamically recalibrating spatial and channel-wise representations, which directly contributes to improved classification and segmentation accuracy. In contrast, Grad-CAM is used exclusively for post hoc visualization to highlight salient regions associated with model decisions. It does not influence training or model outputs and serves only to provide interpretability for external validation. Similarly, the uncertainty estimation integrated into the ACWI module is not optimized for performance metrics but is used to improve prediction reliability by dynamically selecting the inference pathway based on model confidence. Therefore, among the three techniques, only attention directly contributes to performance, while Grad-CAM and uncertainty estimation enhance model transparency and decision safety.

The attention blocks in DIANet are placed after each major stage of the backbone network. Specifically, for architectures like ResNet50 and MobileNetV3, attention modules are inserted immediately after the final feature map of each residual or bottleneck stage. These modules apply spatial-channel recalibration before features are passed to the domain-informed latent space encoder. The attention layers are implemented independently of the pre-trained weights and do not modify the backbone's original structure. During fine-tuning, these attention modules are trained jointly with the rest of the model, while the backbone weights are either partially frozen or fine-tuned depending on the dataset size and task complexity. This approach allows us to leverage pre-trained features while introducing trainable attention for task-specific adaptation without violating the integrity of the pre-trained architecture.

Conclusions and future work

In this study, we sought to address the critical challenge of integrating pathology and radiology in medical imaging through AI-driven approaches. While current methods have demonstrated notable advancements, issues such as handling multi-modal imaging, imbalanced datasets, and limited interpretability have restricted their clinical deployment. To overcome these challenges, we proposed the Domain-Informed Adaptive Network (DIANet) in combination with the Adaptive Clinical Workflow Integration (ACWI) strategy. DIANet introduces multi-scale feature extraction, incorporates domain-specific priors, and employs Bayesian uncertainty modeling to enhance interpretability and robustness, crucial for clinical applications. ACWI complements this model by ensuring seamless deployment through explainable AI (XAI) techniques, uncertainty-aware decision support, and modular integration into existing clinical systems such as PACS. Experimental results validated

the framework's effectiveness, demonstrating substantial improvements in diagnostic accuracy, segmentation precision, and reconstruction fidelity across a variety of imaging modalities, thereby highlighting its potential to advance AI-assisted medical imaging.

Despite these promising results, two limitations remain. While DIANet improves interpretability and robustness, its reliance on domain-specific priors necessitates extensive expert input during model development. This requirement could limit scalability to new clinical domains or institutions with differing standards. Future work should explore automated methods for deriving domain priors to enhance generalizability. While the ACWI strategy facilitates clinical integration, its modular approach may not fully capture the complexity of real-world workflows that often require dynamic adaptation to evolving clinical protocols and technologies. Addressing this limitation will involve developing more adaptive integration frameworks capable of learning and evolving alongside clinical practices. By overcoming these limitations, the proposed framework could serve as a cornerstone for fully realizing the potential of AI in medical imaging.

Data availability

The datasets generated and/or analysed during the current study are available in the AI-Assisted, https://github.com/dshfusidhfg-SR/AI-Assisted.git

Received: 10 February 2025; Accepted: 17 June 2025
Published online: 25 July 2025

References

1. Chen, C.-F., Fan, Q. & Panda, R. Crossvit: Cross-attention multi-scale vision transformer for image classification. *IEEE Int. Conf. on Comput. Vis.* (2021).
2. Hong, D. *et al.* Spectralformer: Rethinking hyperspectral image classification with transformers. *IEEE Transactions on Geosci. Remote. Sens.* (2021).
3. Touvron, H. *et al.* Resmlp: Feedforward networks for image classification with data-efficient training. *IEEE Transactions on Pattern Analysis and Mach. Intell.* (2021).
4. Maurício, J., Domingues, I. & Bernardino, J (A literature review. Applied Sciences, Comparing vision transformers and convolutional neural networks for image classification, 2023).
5. Tian, Y., Wang, Y., Krishnan, D., Tenenbaum, J. spsampsps Isola, P. Rethinking few-shot image classification: a good embedding is all you need? *Eur. Conf. on Comput. Vis.* (2020).
6. Hong, D. *et al.* Graph convolutional networks for hyperspectral image classification. *IEEE Transactions on Geosci. Remote. Sens.* (2020).
7. Yang, J. *et al.* Medmnist v2 - a large-scale lightweight benchmark for 2d and 3d biomedical image classification. *Sci. Data* (2021).
8. Sun, L., Zhao, G., Zheng, Y. & Wu, Z. Spectral–spatial feature tokenization transformer for hyperspectral image classification. *IEEE Trans. Geosci. Remote Sens.*https://doi.org/10.1109/TGRS.2022.3144158 (2022).
9. Rao, Y., Zhao, W., Zhu, Z., Lu, J. & Zhou, J. Global filter networks for image classification. *Neural Inf. Process. Syst.* (2021).
10. Wang, X. et al. Transformer-based unsupervised contrastive learning for histopathological image classification. *Med. Image Anal.*https://doi.org/10.1016/j.media.2022.102559 (2022).
11. Mai, Z. *et al.* Online continual learning in image classification: An empirical survey. *Neurocomputing* (2021).
12. Azizi, S. *et al.* Big self-supervised models advance medical image classification. *IEEE Int. Conf. on Comput. Vis.* (2021).
13. Li, B., Li, Y. & Eliceiri, K. Dual-stream multiple instance learning network for whole slide image classification with self-supervised contrastive learning. *Comput. Vis. Pattern Recognit.* (2020).
14. Bhojanapalli, S. *et al.* Understanding robustness of transformers for image classification. *IEEE Int. Conf. on Comput. Vis.* (2021).
15. Kim, H. E. et al. Transfer learning for medical image classification: A literature review. *BMC Med. Imaging*https://doi.org/10.1186/s12880-022-00793-7 (2022).
16. Zhang, C., Cai, Y., Lin, G. & Shen, C. Deepemd: Few-shot image classification with differentiable earth mover's distance and structured classifiers. *Comput. Vis. Pattern Recognit.* (2020).
17. Zhu, Y. *et al.* Deep subdomain adaptation network for image classification. *IEEE Transactions on Neural Networks Learn. Syst.* (2020).
18. Ravi, V. & Chaganti, R. Efficientnet deep learning meta-classifier approach for image-based Android malware detection. *Multimedia Tools Appl.***82**, 24891–24917 (2023).
19. Ravi, V. Attention cost-sensitive deep learning-based approach for skin cancer detection and classification. *Cancers* **14**, 5872 (2022).
20. Ashtiani, F., Geers, A. J. & Aflatouni, F. An on-chip photonic deep neural network for image classification. *Nature* (2021).
21. Chen, L. et al. Review of image classification algorithms based on convolutional neural networks. *Remote Sens.*https://doi.org/10.3390/rs13224712 (2021).
22. Roy, S. K. *et al.* Multimodal fusion transformer for remote sensing image classification. *IEEE Transactions on Geosci. Remote. Sens.* (2022).
23. Sheykhmousa, M. et al. Support vector machine versus random forest for remote sensing image classification: A meta-analysis and systematic review. *IEEE J. Sel. Top. Appl. Earth Obs. Remote Sens.*https://doi.org/10.1109/JSTARS.2020.3026724 (2020).
24. Zhang, Y., Li, W., Sun, W., Tao, R. & Du, Q. Single-source domain expansion network for cross-scene hyperspectral image classification. *IEEE Transactions on Image Processing* (2022).
25. Taori, R. *et al.* Measuring robustness to natural distribution shifts in image classification. *Neural Inf. Process. Syst.* (2020).
26. Ravi, V., Alazab, M., Selvaganapathy, S. & Chaganti, R. A multi-view attention-based deep learning framework for malware detection in smart healthcare systems. *Comput. Commun.***195**, 73–81 (2022).
27. Tondini, T., Isidro, A. & Camarós, E. Case report: Boundaries of oncological and traumatological medical care in ancient Egypt: New palaeopathological insights from two human skulls. *Front. Med.***11**, 1371645 (2024).
28. Bazi, Y., Bashmal, L., Rahhal, M. M. A., Dayil, R. A. & Ajlan, N. A. Vision transformers for remote sensing image classification. *Remote. Sens.* (2021).
29. Peng, J. et al. Domain adaptation in remote sensing image classification: A survey. *IEEE J. Sel. Top. Appl. Earth Obs. Remote Sens.*https://doi.org/10.1109/JSTARS.2022.3220875 (2022).
30. Masana, M. *et al.* Class-incremental learning: Survey and performance evaluation on image classification. *IEEE Transactions on Pattern Analysis Mach. Intell.* (2020).
31. Vermeire, T., Brughmans, D., Goethals, S., de Oliveira, R. M. B. & Martens, D. Explainable image classification with evidence counterfactual. *Pattern Anal. Appl.*https://doi.org/10.1007/s10044-021-01055-y (2022).

32. Lanchantin, J., Wang, T., Ordonez, V. & Qi, Y. General multi-label image classification with transformers. *Comput. Vis. Pattern Recognit.*(2020).
33. Dong, H., Zhang, L. & Zou, B. Exploring vision transformers for polarimetric sar image classification. *IEEE Trans. Geosci. Remote Sens.*https://doi.org/10.1109/TGRS.2021.3137383 (2022).
34. Zheng, X., Sun, H., Lu, X. & Xie, W. Rotation-invariant attention network for hyperspectral image classification. *IEEE Trans. Image Process.*https://doi.org/10.1109/TIP.2022.3177322 (2022).
35. Melo, R. C. et al. Whole slide imaging and its applications to histopathological studies of liver disorders. *Front. Med.***6**, 310 (2020).
36. Wang, A. et al. Large language model answers medical questions about standard pathology reports. *Front. Med.***11**, 1402457 (2024).
37. Lin, T.-Y. *et al.* Feature pyramid networks for object detection. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, 2117–2125 (2017).
38. Woo, S., Park, J., Lee, J.-Y. spsampsps Kweon, I. S. Cbam: Convolutional block attention module. In *Proceedings of the European conference on computer vision (ECCV)*, 3–19 (2018).
39. Wang, J. et al. Deep high-resolution representation learning for visual recognition. *IEEE transactions on pattern analysis and machine intelligence* **43**, 3349–3364 (2020).
40. Liu, S., Qi, L., Qin, H., Shi, J. & Jia, J. Path aggregation network for instance segmentation. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, 8759–8768 (2018).
41. Chen, L.-C., Zhu, Y., Papandreou, G., Schroff, F. spsampsps Adam, H. Encoder-decoder with atrous separable convolution for semantic image segmentation. In *Proceedings of the European conference on computer vision (ECCV)*, 801–818 (2018).
42. Deng, J. *et al.* Imagenet: A large-scale hierarchical image database. In *2009 IEEE conference on computer vision and pattern recognition*, 248–255 (Ieee, 2009).
43. Rao, A. S. & Mahantesh, K. Learning semantic features for classifying very large image datasets using convolution neural network. *SN Comput. Sci.***2**, 187 (2021).
44. Angelova, A., Zhu, S. & Lin, Y. Image segmentation for large-scale subcategory flower recognition. In *2013 IEEE Workshop on Applications of Computer Vision (WACV)*, 39–45 (IEEE, 2013).
45. Wu, C., Timm, M. & Maji, S. Describing textures using natural language. In *European Conference on Computer Vision*, 52–70 (Springer, 2020).
46. Theckedath, D. & Sedamkar, R. Detecting affect states using vgg16, resnet50 and se-resnet50 networks. *SN Comput. Sci.* **1**, 79 (2020).
47. Sharma, J. Enhanced rose leaf disease classification using vision transformer (vit-b/16) detecting black spot, downy mildew, and healthy leaves for improved plant health management. In *2024 5th International Conference on Data Intelligence and Cognitive Informatics (ICDICI)*, 52–56 (IEEE, 2024).
48. Koonce, B. spsampsps Koonce, B. Mobilenetv3. *Convolutional Neural Networks with Swift for Tensorflow: Image Recognit. Dataset Categ.* 125–144 (2021).
49. Nandhini, S. & Ashokkumar, K. An automatic plant leaf disease identification using densenet-121 architecture with a mutation-based henry gas solubility optimization algorithm. *Neural Comput. Appl.***34**, 5513–5534 (2022).
50. Benchallal, F., Hafiane, A., Ragot, N. & Canals, R. Convnext based semi-supervised approach with consistency regularization for weeds classification. *Expert Syst. Appl.***239**, 122222 (2024).
51. Gomez-Smith, M. et al. A physiological characterization of the cafeteria diet model of metabolic syndrome in the rat. *Physiol. Behav.***167**, 382–391 (2016).

Acknowledgements

This is a short text to acknowledge the contributions of specific colleagues, institutions, or agencies that aided the efforts of the authors.

Author contributions

Conceptualization, CL; methodology, CL; software, CL; validation, CL; formal analysis, JZ; investigation, JZ; data curation, JZ; writing-original draft preparation, CL, JZ, RL; writing-review and editing, RL; visualization, RL; supervision, RL; funding acquisition, RL; All authors have read and agreed to the published version of the manuscript.

Funding

Details of all funding sources should be provided, including grant numbers if applicable. Please ensure to add all necessary funding information, as after publication this is no longer possible.

Declarations

Conflicts interests

The authors declare that the research was conducted in the absence of any commercial or financial relationships that could be construed as a potential conflict of interest.

Additional information

Correspondence and requests for materials should be addressed to J.Z.

Reprints and permissions information is available at www.nature.com/reprints.

Publisher's note Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.