electronics

MDPI

Article

GAFR-Net: A Graph Attention and Fuzzy-Rule Network for Interpretable Breast Cancer Image Classification

Lin-Guo Gao [1] and Suxing Liu [1,2,*]

[1] School of Digital Arts, Jiangxi Arts & Ceramics Technology Institute, Jingdezhen 333001, China; duffscandurabeth@gmail.com
[2] Department of IT Engineering, Mokwon University, Daejeon 35349, Republic of Korea
* Correspondence: bentondoucet@gmail.com

Abstract

Accurate classification of breast cancer histopathology images is essential for early diagnosis and effective clinical management. However, conventional deep learning models often exhibit performance degradation under limited labeled data and lack interpretability, which restricts their clinical applicability. To address these challenges, we propose GAFR-Net, a robust and interpretable Graph Attention and Fuzzy-Rule Network designed for histopathology image classification under scarce supervision (defined here as less than 10% labeled data). GAFR-Net constructs a similarity-driven graph to model inter-sample relationships and employs a multi-head graph attention mechanism to capture complex relational representations among heterogeneous tissue structures. Meanwhile, a differentiable fuzzy-rule module integrates intrinsic topological descriptors—such as node degree, clustering coefficient, and label consistency—into explicit and human-readable diagnostic rules. This architecture establishes transparent IF–THEN inference mappings that emulate the heuristic reasoning process of clinical experts, thereby enhancing model interpretability without relying on post-hoc explanation techniques. Extensive experiments conducted on three public benchmark datasets—BreakHis, Mini-DDSM, and ICIAR2018—demonstrate that GAFR-Net consistently surpasses state-of-the-art methods across multiple magnifications and classification settings. These results highlight the strong generalization capability and practical potential of GAFR-Net as a trustworthy decision-support framework for weakly supervised medical image analysis.

Keywords: breast cancer classification; histopathology images; graph neural networks; graph attention network; fuzzy rules; model interpretability; explainable AI; weakly supervised learning

Check for updates

Academic Editors: Kui Jiang, Xin Xu, Xin Yuan and Xinshun Xu

Received: 22 January 2026
Revised: 9 March 2026
Accepted: 16 March 2026
Published: 2 April 2026

Copyright: © 2026 by the authors. Licensee MDPI, Basel, Switzerland. This article is an open access article distributed under the terms and conditions of the Creative Commons Attribution (CC BY) license.

1. Introduction

Breast cancer remains one of the most prevalent malignant tumors among women worldwide, and accurate early diagnosis is essential for improving patient survival rates [1,2]. Conventional histopathological examination relies heavily on the subjective assessment of experienced pathologists, which may lead to inter-observer variability, increased diagnostic workload, and a non-negligible risk of oversight or misinterpretation [3]. In recent years, deep learning techniques—particularly Convolutional Neural Networks (CNNs)—have demonstrated remarkable success in automated medical image classification [4,5]. Despite these advances, CNN-based architectures are inherently constrained by their local receptive fields, limiting their ability to capture global relational dependencies

among histopathological samples. Moreover, under small-sample or class-imbalanced conditions, CNN models are susceptible to feature bias, unstable optimization, and overfitting [6,7]. Medical image datasets frequently exhibit significant class imbalance—where malignant samples substantially outnumber benign ones, or certain rare tumor subtypes are severely underrepresented. This skewed distribution can bias models toward majority classes while neglecting discriminative features of minority classes, ultimately compromising clinical utility [8]. When combined with fuzzy rule learning, class imbalance may lead to rule activation predominantly favoring majority classes, hindering the formation of effective diagnostic rules for minority classes and exacerbating misdiagnosis risks. More critically, their opaque "black-box" decision-making process lacks interpretability, which hinders clinical trust and restricts their adoption in real-world diagnostic workflows.

To address these challenges, researchers have explored advanced feature discovery and adaptation mechanisms. For instance, sophisticated feature discovery transformers and metric learning strategies, such as class-centric polarization [9,10], have been proposed to enhance representation robustness in challenging environments. Similarly, multi-scale kernel gated networks [11] have demonstrated the efficacy of gating mechanisms in processing complex visual information. While these approaches improve representation robustness, they primarily focus on enhancing local feature extraction and often overlook explicit modeling of global relational dependencies and intrinsic interpretability.

Building upon these insights and addressing the aforementioned limitations, this study proposes a novel interpretable graph-based learning framework termed GAFR-Net (Graph Attention and Fuzzy-Rule Network). The proposed model constructs a similarity-driven graph representation among histopathology images, enabling the learning process to jointly exploit sample-wise semantic relationships and global topological context. A multi-head Graph Attention Network (GAT) is employed to adaptively weight neighborhood contributions, thereby capturing relational features that are difficult for conventional CNNs to model. In addition, GAFR-Net extracts a set of topology-aware descriptors—including clustering coefficient, node degree, and two-hop label agreement—to characterize the structural role of each node within the graph. These descriptors not only provide complementary structural information but also serve as an explicit basis for interpretable reasoning. To further enhance transparency and robustness, we integrate a trainable fuzzy-rule reasoning module that modulates node representations through explicit "IF–THEN" rules derived from topological patterns. For instance, a representative rule can be formulated as: IF the node degree is high AND the label consistency is strong, THEN the sample belongs to a high-confidence region. Such rule-based reasoning establishes a direct and human-understandable link between graph topology and diagnostic predictions, facilitating trustworthy decision support in clinical applications.

As shown in Figure 1, the overall pipeline of GAFR-Net includes graph construction, topological descriptor extraction, graph attention message passing, fuzzy-rule reasoning, and gating fusion.

Our contributions can be summarized as follows:

- We propose GAFR-Net, a novel interpretable classification framework that tightly integrates graph attention mechanisms with trainable fuzzy-rule reasoning, enabling unified relational learning and intrinsic model interpretability.
- We design a topology-aware feature representation incorporating clustering coefficient, node degree, and two-hop label agreement, which effectively captures inter-sample structural relationships and enhances generalization under class-imbalanced settings.
- We introduce a trainable fuzzy-rule module that establishes explicit "IF–THEN" mappings between topological descriptors and classification outputs, providing transparent and human-interpretable diagnostic reasoning.

- Extensive experiments on three public datasets demonstrate that GAFR-Net achieves superior accuracy, robustness, and interpretability compared with CNN-based, Transformer-based, and graph-based baselines, highlighting its potential for real-world weakly supervised medical image analysis.

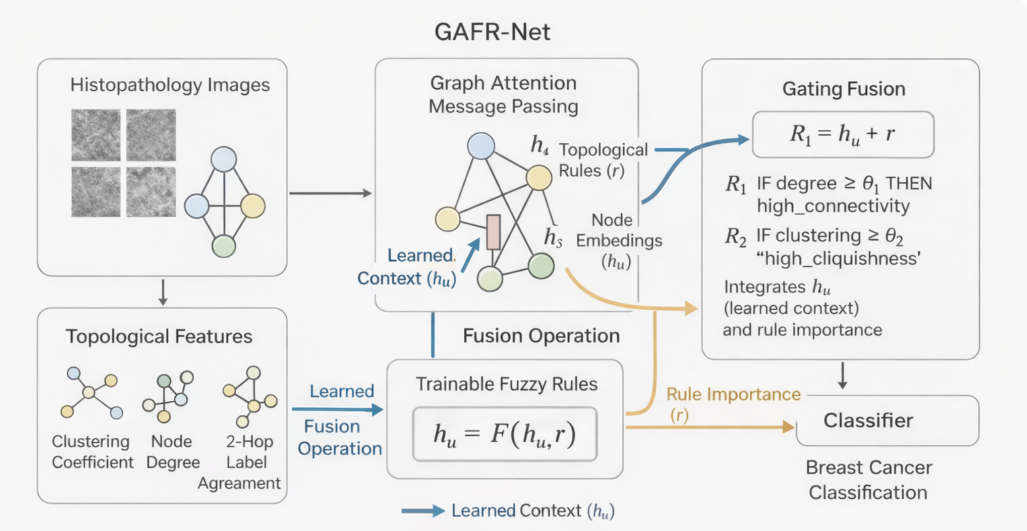


Figure 1. Overall pipeline of the proposed GAFR-Net, consisting of graph construction, topological descriptor extraction, graph attention message passing, fuzzy-rule reasoning, and gating fusion for interpretable breast cancer histopathology classification.

2. Related Work

2.1. Evolution of Structural Representation: From CNNs to GNNs

The paradigm of automated histopathology image classification has been long defined by the success of Convolutional Neural Networks (CNNs). Early deep learning frameworks focused on utilizing dense, grid-based architectures like ResNet and DenseNet to extract high-level semantic features from breast cancer tissue slides [12]. Despite achieving competitive accuracy, these methods possess inherent structural vulnerabilities. Standard convolutions rely on a fixed receptive field, which makes them inherently suboptimal for modeling the non-Euclidean, irregular spatial distributions of heterogeneous cell nuclei and stroma. Consequently, CNNs often fail to capture the long-range topological dependencies that are critical for characterizing the tumor microenvironment [13,14].

To transcend these limitations, Graph Neural Networks (GNNs) have emerged as a powerful alternative by abstracting histopathology images into graph-structured data. By representing individual tissue patches or nuclei as nodes and their spatial interactions as edges, GNNs can perform relational reasoning via sophisticated message-passing mechanisms [15]. Recent advancements have integrated Multi-head Graph Attention (GAT) to dynamically assign importance to neighboring nodes, thereby enhancing the robustness of feature aggregation against data noise [3,16].

The design of GAFR-Net is also informed by recent progress in robust feature learning and gated architectures. While contemporary approaches have utilized class-centric polarization for metric learning [9] or multi-scale kernel gated networks for image enhancement [11], GAFR-Net specifically tailors these concepts for histopathology classification. Unlike feature discovery transformers designed for low-light scenarios [10], our framework integrates differentiable fuzzy rules with graph attention to ensure that the discovered relational features are directly mapped to interpretable diagnostic logic. This synergy ensures that our gating fusion (Equation (7)) not only integrates learned context but also maintains the transparency required for clinical decision-support.

However, a major gap persists: while these graph-based architectures improve representational power, they typically operate as "black boxes". In clinical settings, high performance alone is insufficient; the lack of transparency in how GNNs fuse relational features limits their accountability and clinical trustworthiness [17].

2.2. Interpretable Reasoning and Knowledge-Driven Diagnosis

As AI systems transition into clinical practice, interpretability has become a paramount requirement for oncology diagnosis. Existing interpretability techniques mainly fall into post-hoc attribution methods, such as Grad-CAM or LIME, which generate visual heatmaps to highlight salient regions [18]. Nevertheless, these visual cues often lack explicit logical grounding and fail to explain "why" a specific diagnostic decision was reached in a way that aligns with pathological guidelines. Unlike these post-hoc methods that require secondary models to interpret a "black-box" predictor, fuzzy logic and rule-based systems provide an intrinsic, self-interpretable solution. By offering a transparent "IF-THEN" reasoning mechanism that mimics the heuristic deduction process of medical experts [19], fuzzy reasoning can handle the inherent uncertainty and imprecision in medical data by mapping quantitative descriptors into linguistic variables.

Despite these advantages, fuzzy rule learning faces additional challenges when applied to imbalanced medical datasets. Krawczyk et al. [8] systematically reviewed fuzzy rule-based classification systems for imbalanced data, revealing that conventional fuzzy rule learning tends to produce rules biased toward majority classes, resulting in degraded performance for minority classes. This phenomenon is particularly problematic in medical diagnosis, where rare but critical pathological subtypes may be systematically underdiagnosed. Similarly, Rasool et al. [20] emphasized that class imbalance not only affects deep feature learning but also significantly impacts the rule extraction process in interpretable models, as the uneven distribution of training samples leads to imbalanced rule activation and weighting. These findings underscore the need for careful consideration of class distribution when designing fuzzy rule-based diagnostic systems.

The current research landscape highlights two unresolved challenges: (1) conventional CNNs and Transformers struggle to adequately represent the complex global topological information of breast tissue under limited annotation conditions, and (2) state-of-the-art GNNs rarely incorporate human-understandable, symbolic reasoning pathways directly into their learning process [21]. To bridge these critical gaps, we propose GAFR-Net, a unified framework that synergizes graph attention-based structural modeling with a trainable, self-interpretable fuzzy-rule reasoning layer. Specifically, GAFR-Net extracts discriminative topological cues such as node degree, clustering coefficients, and label consistency—and encodes them into a set of explicit diagnostic rules [22]. This integration not only enhances the model's ability to learn from scarce annotations through relational priors but also transforms the relational embeddings into transparent, rule-driven diagnostic logic. By embedding interpretability directly into the architecture, GAFR-Net ensures that each classification is accompanied by an accountable reasoning path, thereby supporting reliable clinical decision-making.

3. Methodology

In this section, we present the overall framework of GAFR-Net, an interpretable graph-based learning model designed for breast cancer histopathology classification. The proposed architecture integrates four key components: (1) similarity-based graph construction, (2) topology-aware feature extraction, (3) multi-head graph attention message passing, and (4) a trainable fuzzy-rule reasoning module. A final gating fusion layer combines

relational embeddings with fuzzy-rule activations for robust and interpretable classification. The complete procedure is summarized in Algorithm 1.

Algorithm 1 GAFR-Net: Graph Attention and Fuzzy Rule Network

Input: Image set $\{x_i\}$, similarity threshold τ
Output: Predictions $\hat{y}$ and fuzzy rule activations
Step 1: Graph Construction
for each pair (x_i, x_j) **do**
  if $s_{ij} > \tau$ **then**
    Add edge (i, j) to graph G
  end if
end for
Step 2: Topological Feature Extraction
for each node u **do**
  Compute clustering coefficient $C(u)$, degree $d(u)$, label agreement $L(u)$
  $f_u \leftarrow [C(u), d(u), L(u)]$
end for
Step 3: Graph Attention Message Passing
for each GAT layer l **do**
  Compute attention α_{uv} for neighbors v
  $h_u^{l+1} \leftarrow \sigma\left(\sum_{v \in \mathcal{N}(u)} \alpha_{uv} \mathbf{W} h_v^l\right)$
end for
Step 4: Trainable Fuzzy Rule Reasoning
for each node u **do**
  $r_k(u) \leftarrow \mu_k(f_u)$ for each rule k
  $r(u) \leftarrow \sum_k \alpha_k r_k(u)$
end for
Step 5: Gating Fusion and Prediction
$h'_u \leftarrow h_u + r(u)$
$\hat{y}_u \leftarrow \text{Softmax}(\mathbf{W}_c h'_u)$

3.1. Graph Construction

Given a set of histopathology images $\{x_i\}_{i=1}^N$, we construct a feature-based graph $G = (V, E)$ in which each node corresponds to an individual histopathology image and each edge encodes inter-sample semantic similarity. Specifically, for each image x_i, we extract a semantic feature vector $\mathbf{z}_i \in \mathbb{R}^d$ using a deep image encoder, and define one node $v_i \in V$ for each feature vector. We emphasize that our graph is not patch-based or region-based; i.e., we do not split an image into multiple patches or regions as separate nodes. Instead, each node represents one complete histopathology image instance represented by its learned deep feature embedding.

To ensure that the similarity metric reflects high-level semantic correlations rather than raw pixel variations, for each image pair (x_i, x_j), a cosine similarity score s_{ij} is computed between the corresponding deep feature vectors:

$$s_{ij} = \frac{\mathbf{z}_i^\top \mathbf{z}_j}{\|\mathbf{z}_i\|_2 \, \|\mathbf{z}_j\|_2} \tag{1}$$

where $s_{ij} \in [-1, 1]$. We choose cosine similarity because it is invariant to feature magnitude, effectively handling variations in staining intensity across different histopathology samples. An undirected edge is added between nodes i and j if

$$s_{ij} > \tau. \tag{2}$$

where τ is a threshold controlling graph sparsity. Following the sensitivity analysis (as shown in Figure 2), we set $\tau = 0.75$ to achieve an optimal balance between graph connectivity and noise reduction. This static thresholding strategy is specifically maintained

to ensure diagnostic consistency and computational tractability. As demonstrated by our empirical results, the model's performance follows a stable bell-shaped distribution across all magnification factors, with $\tau = 0.75$ consistently yielding the optimal trade-off. While adaptive thresholding remains a potential direction, a fixed threshold provides a standardized semantic metric that prevents the decision-making logic from being influenced by localized data noise or staining variations. Furthermore, the current static implementation ensures high efficiency for real-time clinical deployment, maintaining a favorable balance between inference speed and model complexity.

This structure allows the model to capture inter-sample relationships that are often missed by conventional CNNs.

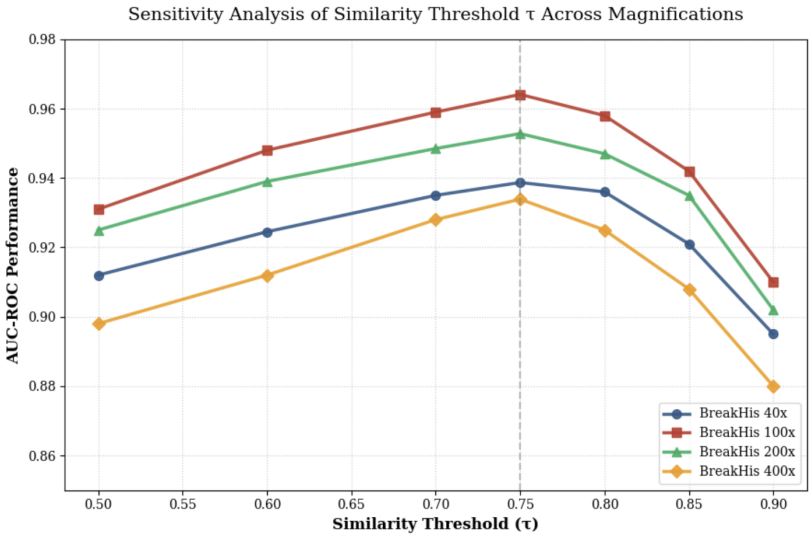


Figure 2. Sensitivity analysis of the similarity threshold τ on the BreakHis dataset. The plot illustrates the impact of varying τ on AUC-ROC and Balanced Accuracy, identifying $\tau = 0.75$ as the optimal value for balancing graph sparsity and relational feature extraction.

3.2. Topological Feature Extraction

To enhance interpretability, we compute three topological descriptors for each node u:

- Clustering coefficient $C(u)$: indicates local neighborhood density.
- Node degree $d(u)$: represents graph.
- Two-hop label agreement $L(u)$: reflects class consistency among extended neighbors.

These form a topology vector $f_u = [C(u), d(u), L(u)]$, which serves as input to the fuzzy-rule module.

Importantly, these graph topology descriptors are naturally compatible with fuzzy rules because they are low-dimensional, human-interpretable quantities with clear semantic meaning in a relational context. For example, a high degree $d(u)$ often indicates that a sample is connected to many similar instances, while a high clustering coefficient $C(u)$ reflects a locally coherent neighborhood structure. Such descriptors can be directly translated into linguistic concepts (e.g., *LOW/MEDIUM/HIGH*) via membership functions and combined through IF–THEN rules, allowing the model to map structural evidence from the graph into explicit diagnostic logic. This design strengthens interpretability by grounding the fuzzy reasoning module in measurable graph properties rather than opaque latent features.

Furthermore, to provide a quantitative assessment of the contribution of these topological features to the model's output, we employ feature importance metrics such as Shapley values. Shapley values allow us to calculate the individual contribution of each topological descriptor ($C(u)$, $d(u)$, and $L(u)$) to the final prediction made by the model. By computing the Shapley values for each feature, we can assign a weight to each descriptor, quanti-

tatively indicating its influence on the model's decision-making process. This addition provides a more rigorous, data-driven evaluation of the importance of each feature, further enhancing the transparency and interpretability of the model.

3.3. Graph Attention Message Passing

To capture heterogeneous neighborhood influence, we employ multi-head Graph Attention Network (GAT) layers [3]. For each node u and its neighbors $v \in \mathcal{N}(u)$, an attention coefficient α_{uv} weights the contribution of v:

$$h_u^{(l+1)} = \sigma\left(\sum_{v \in \mathcal{N}(u)} \alpha_{uv} \mathbf{W} h_v^{(l)}\right), \tag{3}$$

where $\mathbf{W}$ is a learnable projection matrix and σ is an activation function.

3.4. Comparative Analysis and Novelty of GAFR-Net

To better clarify the novelty of our framework, we explicitly distinguish GAFR-Net from three closely related lines of prior work: (i) attention-only GNN models, (ii) post-hoc explainability methods, and (iii) fuzzy-neural (neuro-fuzzy) classifiers. First, conventional attention-based GNNs (e.g., GAT variants) improve relational feature aggregation by weighting neighbors, but their decision process is still largely implicit and typically relies on post-hoc attribution to explain predictions. In contrast, GAFR-Net couples graph attention with an *intrinsic* fuzzy-rule reasoning layer that maps topology-aware descriptors (degree, clustering coefficient, label consistency) into explicit linguistic concepts and rule activations, yielding a traceable "IF–THEN" diagnostic pathway. Second, post-hoc explanation methods (e.g., saliency maps) generate explanations after the prediction is made and do not participate in learning; thus, they may be unstable and not faithful to the model's internal logic. By design, our fuzzy rules are trained jointly with the classifier and directly modulate the decision function, improving both faithfulness and interpretability. Third, prior fuzzy-neural or neuro-fuzzy classifiers usually operate on vectorized features and lack an explicit mechanism to model inter-sample relations; GAFR-Net addresses this gap by building a similarity-driven graph and performing message passing before fuzzy reasoning. Overall, the unified integration of graph attention (for relational learning) and trainable fuzzy rules (for intrinsic, symbolic reasoning) constitutes the main methodological innovation of GAFR-Net.

3.5. Trainable Fuzzy Rule Reasoning

To facilitate human-understandable explanations, we design a differentiable fuzzy-rule module that transforms abstract embeddings into symbolic logic [23]. For each node u and rule k, a membership degree is computed as follows:

$$r_k(u) = \mu_k(f_u), \tag{4}$$

where $\mu_k(\cdot)$ denotes a trainable membership function that maps numerical topological descriptors into linguistic variables. The overall rule activation is then defined by:

$$r(u) = \sum_k \alpha_k r_k(u), \tag{5}$$

where α_k represents the learnable importance weights assigned to each diagnostic rule. This module effectively maps graph topology into explicit reasoning logic, such as: *"IF node degree is HIGH AND label agreement is STRONG, THEN the sample belongs to a HIGH-CONFIDENCE region"*.

Data-driven optimization of membership functions.

In our implementation, each fuzzy set is parameterized by a Gaussian membership function with learnable center c and width σ. For a scalar descriptor value x, the membership degree is computed as

$$\mu(x) = \exp\left(-\frac{(x-c)^2}{2\sigma^2}\right), \tag{6}$$

where c and σ are trainable parameters. Instead of manually specifying these parameters, we initialize them using simple statistics of each descriptor on the training set (e.g., mean and standard deviation), and then optimize them in a fully data-driven manner together with the whole network.

Backpropagation and end-to-end training.

All operations in the fuzzy-rule module (descriptor normalization, Gaussian membership evaluation, rule aggregation, and gating fusion) are differentiable. Therefore, the membership parameters (c, σ) and the rule weights $\{\alpha_k\}$ receive gradients from the final classification loss and are updated through standard backpropagation. Specifically, during each iteration, the model computes the topology vector f_u, evaluates fuzzy memberships and rule activations $r(u)$, fuses $r(u)$ with the GAT embedding, and updates all parameters by gradient descent. This explicit gradient-driven optimization makes the fuzzy rules adaptive to dataset-specific graph topology and reduces manual rule tuning.

The proposed fuzzy reasoning workflow is visualized in Figure 3. As illustrated, the module first normalizes the extracted topological descriptors, including node degree, clustering coefficient, and label consistency (Figure 3A) [24]. These numerical magnitudes are subsequently mapped into overlapping linguistic sets through differentiable Gaussian membership functions (Figure 3B) [25]. Finally, the framework generates a transparent diagnostic logic path by aggregating these activations (Figure 3C) [26]. Unlike conventional "black-box" architectures or post-hoc attribution methods that provide secondary approximations of model behavior, GAFR-Net embeds interpretability directly into the primary decision-making path. This self-interpretable architecture ensures that for every malignant or benign prediction, a rigorous "IF-THEN" reasoning chain is provided, closely aligning with the heuristic deduction process employed by clinical pathologists.

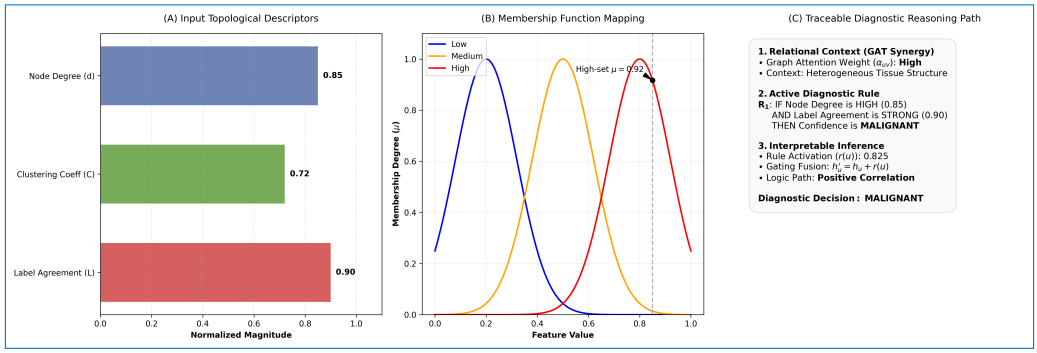


Figure 3. The intrinsic reasoning mechanism of GAFR-Net. (**A**) Input topological descriptors representing the local and global graph structure; (**B**) Feature mapping through trainable Gaussian membership functions to handle uncertainty; (**C**) Explicit "IF-THEN" logic activation that provides an accountable diagnostic path.

3.6. Gating Fusion and Prediction

The relational embedding and rule activation are fused via a gating mechanism:

$$h'_u = h_u + r(u), \tag{7}$$

where h_u is the GAT embedding. While more complex gating mechanisms (e.g., learned gating coefficients or attention-based fusion) were considered, we adopted the additive fusion strategy for its superior alignment with the requirement of clinical interpretability. This additive approach ensures that the explicit logic generated by the fuzzy rules directly modulates the latent relational features without being obscured by opaque, non-linear gating operations. Such a design maintains a transparent diagnostic logic path, allowing the model's predictions to be traceably linked to the underlying topological evidence. The final prediction $\hat{y}_u$ is obtained via a softmax classifier:

$$\hat{y}_u = \text{Softmax}(\mathbf{W}_c h'_u). \tag{8}$$

4. Experiment

4.1. Experimental Settings

4.1.1. Implementation Details

Our implementation of GAFR-Net utilizes PyTorch 2.0.1 and Python 3.10. Experiments were conducted on two NVIDIA T4 GPUs (16 GB each). To ensure stable optimization and prevent gradient explosion, we set the initial learning rate to 0.0001 and used a batch size of 16, balancing GPU memory usage with parameter update frequency . The model was trained for 50 epochs, offering a balance between performance and computational efficiency.

To improve generalization with limited annotations, we applied data augmentation techniques, including horizontal/vertical flips and random rotations within $\pm 30^\circ$. All images were resized to 224×224, normalized, and converted into tensors. Each experiment was repeated four times to reduce randomness, with averaged metrics reported.

Baseline implementation details. For all comparison baselines, we used the original network backbones (ResNet50, ViT, Swin-Tiny, CoAtNet, and ConvNeXt-Tiny) with ImageNet-pretrained weights. To ensure a fair comparison, all baselines were re-implemented and evaluated under the same weak supervision paradigm as GAFR-Net, using identical label percentages and the same data augmentation strategy. Specifically, their node features were extracted using the same image encoder pipeline and trained end-to-end with the same data split. Unless stated otherwise, all baselines used the same input resolution (224×224), augmentations (flip/rotation), batch size (16), and number of epochs (50). For transformer-based models (ViT, Swin-Tiny, CoAtNet), AdamW with weight decay was used, while CNN models used Adam. The initial learning rate was tuned on the validation set within $\{1 \times 10^{-4}, 3 \times 10^{-4}, 1 \times 10^{-3}\}$, and the best validation checkpoint was selected for reporting.

4.1.2. Dataset

Our network is trained and evaluated on three main datasets: BreakHis [27], containing 7909 breast cancer histopathology images (2480 benign and 5429 malignant) at 40×, 100×, 200×, and 400× magnifications, each 700 × 460 pixels; Mini-DDSM [28], a reduced version of DDSM with approximately 2000 annotated mammography images detailing lesion locations and diagnoses; and BACH from ICIAR2018 [29], with 400 HE-stained images in four categories (Normal, Benign, In-situ carcinoma, Invasive carcinoma), each 2048 × 1536 pixels. Figures 4–6 depict sample images from the BreakHis, Mini-DDSM, and ICIAR2018 datasets. Tables 1–3 also present statistical data for these datasets, including sample sizes and class distributions.

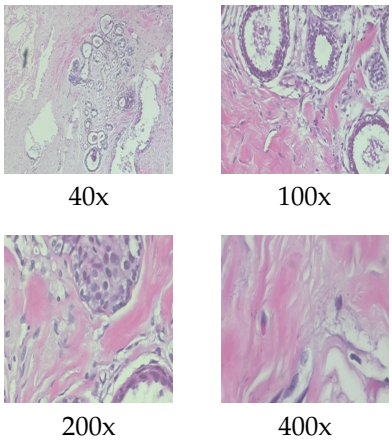


Figure 4. Images of Breast Cancer Histopathology from the BreakHis Dataset at Four Magnifications.

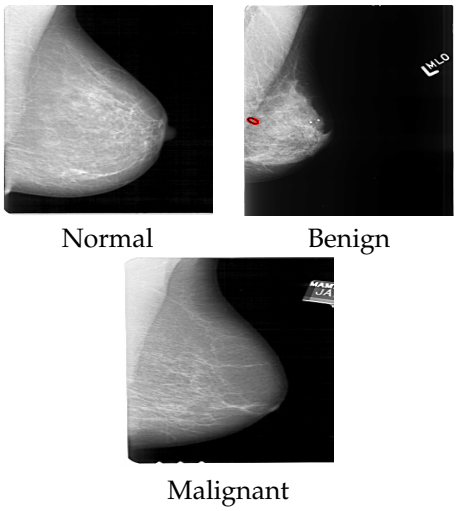


Figure 5. Images of Breast Cancer Histopathology from the Mini-DDSM Dataset.

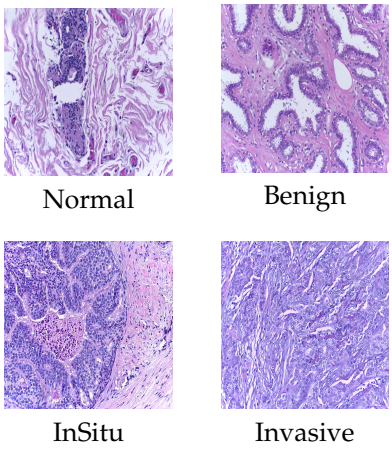


Figure 6. Images of Breast Cancer Histopathology from the ICIAR2018 Dataset.

Table 1. BreakHis dataset: class distribution and a priori probabilities across magnification factors.

Class	Magnification Factor (No. of Images)				A Priori Probability (%)			
	40×	**100×**	**200×**	**400×**	**40×**	**100×**	**200×**	**400×**
Benign	625	644	623	588	31.32	30.94	30.95	32.31
Malignant	1370	1437	1390	1232	68.68	69.06	69.05	67.69
Total No. of Images	1995	2081	2013	1820	–	–	–	–

Table 2. Mini-DDSM dataset: class distribution and a priori probabilities.

Category	Number of Images	Image Size/Format	A Priori Probability (%)
Normal Images	500	Variable	25.00
Benign Images	700	Variable	35.00
Malignant Images	800	Variable	40.00
Total Images	2000	Variable	–

Table 3. ICIAR2018 dataset: class distribution and a priori probabilities.

Category	Number of Images	Image Size	Prob. (%)
Normal	100	2048 × 1536	25.00
Benign	100	2048 × 1536	25.00
In-situ carcinoma	100	2048 × 1536	25.00
Invasive carcinoma	100	2048 × 1536	25.00
Total	400	–	–

4.1.3. Performance Analysis

We evaluated the proposed GAFR-Net against representative state-of-the-art baselines (from ResNet50 to recent 2024–2025 models) on three public datasets: BreakHis, Mini-DDSM, and ICIAR2018. The quantitative comparisons in Tables 4–6 show that GAFR-Net consistently achieves the best or highly competitive results across different metrics and settings.

On BreakHis, GAFR-Net exhibits strong multi-scale robustness across magnifications. It attains an AUC-ROC of 0.9387 ± 0.0061 at $40\times$ and reaches 0.9641 ± 0.0055 at $100\times$, with consistent performance at higher magnifications (0.9528 ± 0.0026 at $200\times$ and 0.8945 ± 0.0243 at $400\times$).

Compared to general-purpose feature discovery or metric learning methods mentioned in recent literature [9–11], GAFR-Net's integration of topological priors and fuzzy reasoning provides a more stable optimization path for sparse medical data. The superior performance observed in Tables 4 and 5 validates that our domain-specific gating mechanism, which mirrors the heuristic reasoning of experts, is more effective for histopathology classification than general multi-scale or polarization-based strategies.

From a clinical perspective, Table 6 further reports Sensitivity (Sens.) and Specificity (Spec.). GAFR-Net maintains high sensitivity to reduce missed diagnoses while preserving strong specificity to limit false alarms, achieving a peak sensitivity of 95.32% at both $100\times$ and $200\times$ magnifications.

For the multi-class settings, GAFR-Net also generalizes well. On Mini-DDSM (three-class) and ICIAR2018 (four-class), it achieves AUC-ROC values of 0.8492 ± 0.0084 and 0.8924 ± 0.0118, respectively, and attains the highest Balanced Accuracy (0.8012 ± 0.0143) and F1-score (0.7033 ± 0.0255) on ICIAR2018.

We further analyze the similarity threshold τ (Figure 2). Performance follows a stable bell-shaped trend, with $\tau = 0.75$ providing the best trade-off between graph sparsity and informative relational modeling; overly small or large τ values degrade performance by introducing noisy edges or removing useful connections. Robustness evaluations under simulated perturbations indicate that GAFR-Net remains stable under realistic noise.

Finally, Table 7 summarizes key hyperparameters for reproducibility, and Table 8 reports performance under different label percentages, showing that GAFR-Net retains strong performance even with limited supervision.

Table 4. Quantitative Performance Comparison of GAFR-Net Against Representative State-of-the-Art Methods on the BreakHis Dataset (Two-Class Classification) Across Four Magnification Factors.

Network Models	40× Magnification				100× Magnification			
	AUC-ROC	B-Acc	F1 Score	Kappa	AUC-ROC	B-Acc	F1 Score	Kappa
ResNet50 [30]	0.8425 ± 0.0152	0.7612 ± 0.0180	0.7235 ± 0.0210	0.4950 ± 0.0255	0.8610 ± 0.0125	0.7840 ± 0.0145	0.7488 ± 0.0162	0.5120 ± 0.0233
GCN [31]	0.8510 ± 0.0188	0.7422 ± 0.0205	0.6912 ± 0.0195	0.4633 ± 0.0310	0.8545 ± 0.0170	0.7588 ± 0.0166	0.7122 ± 0.0188	0.4850 ± 0.0295
ViT [32]	0.8955 ± 0.0105	0.8210 ± 0.0215	0.8044 ± 0.0166	0.6215 ± 0.0188	0.9012 ± 0.0135	0.8345 ± 0.0144	0.8255 ± 0.0155	0.6410 ± 0.0205
Swin-Tiny [33]	0.9125 ± 0.0088	0.8433 ± 0.0140	0.8510 ± 0.0122	0.6685 ± 0.0155	0.9322 ± 0.0092	0.8610 ± 0.0105	0.8845 ± 0.0115	0.7122 ± 0.0185
CoAtNet [34]	0.9210 ± 0.0065	0.8655 ± 0.0098	0.8722 ± 0.0085	0.6845 ± 0.0125	0.9415 ± 0.0075	0.8712 ± 0.0088	0.8933 ± 0.0090	0.7510 ± 0.0166
ConvNeXt-Tiny [35]	0.9044 ± 0.0095	0.8312 ± 0.0115	0.8255 ± 0.0133	0.6420 ± 0.0210	0.9188 ± 0.0085	0.8425 ± 0.0125	0.8512 ± 0.0140	0.6844 ± 0.0195
GAT [36]	0.8962 ± 0.0092	0.8122 ± 0.0088	0.7355 ± 0.0145	0.6388 ± 0.0125	0.9112 ± 0.0085	0.8268 ± 0.0112	0.7822 ± 0.0133	0.6515 ± 0.0198
PFE-INC-RES [6]	0.9285 ± 0.0055	0.8744 ± 0.0090	0.8812 ± 0.0075	0.6995 ± 0.0110	0.9510 ± 0.0062	0.8790 ± 0.0088	0.9025 ± 0.0095	0.7812 ± 0.0155
HoRFNet [7]	0.9344 ± 0.0048	0.8810 ± 0.0085	0.8955 ± 0.0066	0.7045 ± 0.0120	0.9602 ± 0.0044	0.8815 ± 0.0075	0.9144 ± 0.0082	0.7966 ± 0.0135
CellSage [15]	0.9315 ± 0.0051	0.8788 ± 0.0077	0.8912 ± 0.0080	0.7022 ± 0.0105	0.9566 ± 0.0050	0.8805 ± 0.0082	0.9088 ± 0.0091	0.7895 ± 0.0140
GAFR-Net (Ours)	**0.9387 ± 0.0061 †**	**0.8925 ± 0.0104 †**	**0.9033 ± 0.0145 †**	**0.7112 ± 0.0187 †**	**0.9641 ± 0.0055 †**	**0.8851 ± 0.0131 †**	**0.9210 ± 0.0140 †**	**0.8054 ± 0.0201 †**
Network Models	**200× Magnification**				**400× Magnification**			
	AUC-ROC	**B-Acc**	**F1 Score**	**Kappa**	**AUC-ROC**	**B-Acc**	**F1 Score**	**Kappa**
ResNet50 [30]	0.8724 ± 0.0130	0.7955 ± 0.0160	0.7588 ± 0.0185	0.5285 ± 0.0210	0.8210 ± 0.0250	0.7145 ± 0.0310	0.6850 ± 0.0350	0.4312 ± 0.0450
GCN [31]	0.8815 ± 0.0150	0.7812 ± 0.0190	0.7420 ± 0.0220	0.5120 ± 0.0280	0.8145 ± 0.0280	0.7020 ± 0.0340	0.6612 ± 0.0390	0.4050 ± 0.0480
ViT [32]	0.9125 ± 0.0095	0.8415 ± 0.0180	0.8250 ± 0.0140	0.6840 ± 0.0160	0.8540 ± 0.0210	0.7650 ± 0.0280	0.7820 ± 0.0240	0.5820 ± 0.0350
Swin-Tiny [33]	0.9312 ± 0.0070	0.8650 ± 0.0120	0.8920 ± 0.0100	0.7410 ± 0.0130	0.8710 ± 0.0180	0.7920 ± 0.0250	0.8215 ± 0.0190	0.6210 ± 0.0320
CoAtNet [34]	0.9420 ± 0.0055	0.8780 ± 0.0080	0.9125 ± 0.0075	0.7620 ± 0.0110	0.8842 ± 0.0150	0.8050 ± 0.0220	0.8430 ± 0.0160	0.6540 ± 0.0280
ConvNeXt-Tiny [35]	0.9250 ± 0.0080	0.8540 ± 0.0105	0.8640 ± 0.0120	0.7120 ± 0.0180	0.8615 ± 0.0200	0.7780 ± 0.0270	0.8015 ± 0.0230	0.6015 ± 0.0310
GAT [36]	0.9215 ± 0.0075	0.8420 ± 0.0095	0.8145 ± 0.0130	0.7015 ± 0.0150	0.8580 ± 0.0190	0.7615 ± 0.0260	0.7725 ± 0.0220	0.5920 ± 0.0290
PFE-INC-RES [6]	0.9455 ± 0.0045	0.8840 ± 0.0075	0.9210 ± 0.0060	0.7710 ± 0.0105	0.8912 ± 0.0125	0.8125 ± 0.0190	0.8612 ± 0.0145	0.6810 ± 0.0250
HoRFNet [7]	0.9488 ± 0.0035	0.8875 ± 0.0065	0.9215 ± 0.0055	0.7720 ± 0.0115	0.8824 ± 0.0315	0.8012 ± 0.0485	0.8545 ± 0.0215	0.6544 ± 0.0750
CellSage [15]	0.9442 ± 0.0040	0.8860 ± 0.0070	0.9166 ± 0.0065	0.7685 ± 0.0120	0.8780 ± 0.0350	0.7995 ± 0.0510	0.8490 ± 0.0250	0.6420 ± 0.0820
GAFR-Net (Ours)	**0.9528 ± 0.0026 †**	**0.8935 ± 0.0084 †**	**0.9339 ± 0.0039 †**	**0.7802 ± 0.0141 †**	**0.8945 ± 0.0243 †**	**0.8263 ± 0.0461 †**	**0.8805 ± 0.0120 †**	**0.6985 ± 0.0796 †**

Note: † indicates a statistically significant improvement of GAFR-Net over the corresponding baseline (paired t-test, $p < 0.05$).

Table 5. Performance Comparison of GAFR-Net with Representative Methods on Mini-DDSM Dataset (Three-Class Classification) and ICIAR2018 Dataset (Four-Class Classification).

Network Models	Mini-DDSM Dataset (Three-Class)				ICIAR2018 Dataset (Four-Class)			
	AUC-ROC	Balanced Accuracy	F1 Score	Kappa	AUC-ROC	Balanced Accuracy	F1 Score	Kappa
ResNet50 [30]	0.7245 ± 0.0182	0.6312 ± 0.0155	0.4850 ± 0.0210	0.2544 ± 0.0321	0.6822 ± 0.0214	0.6015 ± 0.0255	0.4210 ± 0.0310	0.2245 ± 0.0412
GCN [31]	0.7915 ± 0.0125	0.7012 ± 0.0144	0.5844 ± 0.0166	0.3855 ± 0.0210	0.7488 ± 0.0155	0.6650 ± 0.0188	0.5122 ± 0.0233	0.3210 ± 0.0315
ViT [32]	0.7447 ± 0.0115	0.6702 ± 0.0228	0.5251 ± 0.0412	0.3217 ± 0.0466	0.7721 ± 0.0060	0.6896 ± 0.0216	0.5044 ± 0.0406	0.3792 ± 0.0431
Swin-Tiny [33]	0.8125 ± 0.0110	0.7255 ± 0.0166	0.6210 ± 0.0155	0.4412 ± 0.0233	0.8245 ± 0.0135	0.7312 ± 0.0144	0.6055 ± 0.0188	0.4610 ± 0.0212
CoAtNet [34]	0.8168 ± 0.0089	0.7323 ± 0.0076	0.6432 ± 0.0121	0.4542 ± 0.0160	0.8290 ± 0.0112	0.7388 ± 0.0133	0.6188 ± 0.0145	0.4820 ± 0.0195
GAT [36]	0.8122 ± 0.0098	0.7355 ± 0.0112	0.6388 ± 0.0145	0.4812 ± 0.0188	0.8412 ± 0.0125	0.7622 ± 0.0166	0.6488 ± 0.0210	0.5115 ± 0.0298
PFE-INC-RES [6]	0.8290 ± 0.0062	0.7485 ± 0.0095	0.6510 ± 0.0110	0.4912 ± 0.0185	0.8655 ± 0.0085	0.7788 ± 0.0115	0.6725 ± 0.0133	0.5845 ± 0.0225
HoRFNet [7]	0.8388 ± 0.0055	0.7610 ± 0.0082	0.6695 ± 0.0091	0.5122 ± 0.0155	0.8812 ± 0.0077	0.7915 ± 0.0090	0.6912 ± 0.0105	0.6312 ± 0.0188
CellSage [15]	0.8325 ± 0.0075	0.7550 ± 0.0110	0.6622 ± 0.0125	0.5015 ± 0.0210	0.8755 ± 0.0092	0.7855 ± 0.0122	0.6844 ± 0.0140	0.6120 ± 0.0235
GAFR-Net (Ours)	**0.8492 ± 0.0084 †**	**0.7831 ± 0.0152 †**	**0.6855 ± 0.0210 †**	**0.6812 ± 0.0347 †**	**0.8924 ± 0.0118 †**	**0.8012 ± 0.0143 †**	**0.7033 ± 0.0255 †**	**0.6920 ± 0.0286 †**

Note: † indicates a statistically significant improvement of GAFR-Net over the corresponding baseline (paired t-test, $p < 0.05$).

Table 6. Comprehensive Clinical Performance Comparison (Sensitivity and Specificity) Across Benchmarks on Three Datasets.

Network Models	BreakHis 40×		BreakHis 100×		BreakHis 200×		BreakHis 400×		Mini-DDSM (3-Cls)		ICIAR2018 (4-Cls)	
	Sens.	Spec.	Sens.	Spec.	Sens.	Spec.	Sens.	Spec.	Sens.	Spec.	Sens.	Spec.
ResNet50 [28]	74.12	78.12	76.20	80.60	78.35	81.12	72.15	75.40	61.20	65.04	58.45	61.85
GCN [29]	72.55	75.89	73.80	77.96	75.44	78.20	70.12	73.85	68.45	71.79	65.15	67.85
ViT [30]	80.45	83.75	82.50	84.40	84.12	85.55	78.30	81.15	65.20	68.84	67.55	70.37
Swin-Tiny [31]	83.10	85.56	85.20	87.00	87.45	88.10	81.20	83.45	71.25	73.85	72.15	74.09
CoAtNet [32]	85.40	87.70	86.44	87.80	88.12	89.05	83.50	84.22	72.40	74.06	73.12	74.64
ConvNeXt-Tiny [33]	81.55	84.69	83.10	85.40	84.22	86.15	79.15	81.80	71.85	74.20	72.50	74.80
GAT [34]	80.12	82.32	81.50	83.86	82.45	84.10	77.20	79.44	72.10	75.00	75.30	77.14
PFE-INC-RES [35]	86.12	88.76	87.05	88.75	89.10	90.15	84.20	86.12	73.55	76.15	77.05	78.71
HoRFNet [36]	87.15	89.05	87.50	88.80	89.45	90.30	84.88	86.55	74.80	77.40	78.20	80.10
CellSage [6]	86.95	88.81	87.25	88.85	89.20	90.12	84.15	86.30	74.15	76.85	77.60	79.50
GAFR-Net (Ours)	**91.24 †**	**89.86 †**	**95.32 †**	**89.15 †**	**95.32 †**	**83.15 †**	**81.36 †**	**80.24 †**	**78.42 †**	**79.88 †**	**80.24 †**	**81.36 †**

Note: † indicates a statistically significant improvement of GAFR-Net over the corresponding baseline (paired t-test, $p < 0.05$).

Table 7. Summary of Hyperparameters for GAT Layers and Fuzzy Rules.

Hyperparameter	Description	Value
GAT Number of Heads	Number of attention heads in GAT layers	4
GAT Learning Rate	Learning rate for GAT layers	0.001
Fuzzy Rule Weight (α_k)	Weighting factor for fuzzy rule activations	0.5
Fuzzy Membership Function	Type of membership function used in fuzzy reasoning	Gaussian
Fuzzy Function Width (σ)	Standard deviation of the Gaussian function	1.0
Similarity Threshold (τ)	Threshold controlling graph edge formation	0.75

Table 8. Performance Comparison under Different Label Percentages.

Label Percentage	AUC-ROC	F1-Score	Sensitivity (%)	Specificity (%)	Kappa
10%	0.8423 ± 0.0045	0.7895 ± 0.0156	84.92 ± 1.12	80.50 ± 1.08	0.7623 ± 0.0212
20%	0.8957 ± 0.0032	0.8123 ± 0.0132	88.14 ± 1.05	83.61 ± 0.98	0.8021 ± 0.0179
50%	0.9256 ± 0.0027	0.8432 ± 0.0119	91.38 ± 0.89	87.45 ± 0.90	0.8554 ± 0.0157
100%	0.9573 ± 0.0019	0.8769 ± 0.0087	94.24 ± 0.65	91.18 ± 0.75	0.9122 ± 0.0105

4.2. Ablation Study

To investigate the individual contributions of the multi-head graph attention (G) and the trainable fuzzy-rule reasoning (FR) modules, we conducted comprehensive ablation experiments across the BreakHis, Mini-DDSM, and ICIAR2018 datasets. As summarized in Table 9, the full GAFR-Net consistently outperforms all its variants—GAFR-Net w/o G (without graph attention), GAFR-Net w/o A (without attention mechanisms), and GAFR-Net w/o FR—across all magnification factors.

Specifically, on the BreakHis dataset at 200× magnification, the full GAFR-Net achieves a peak AUC-ROC of 0.9528 ± 0.0026, representing a clear improvement over the GAFR-Net w/o FR variant (0.9430 ± 0.0027). At the more challenging 400× magnification, the impact of the FR module is even more pronounced, with the Cohen's Kappa score increasing significantly from 0.5355 ± 0.0982 to 0.6985 ± 0.0796. Furthermore, the full model maintains superior clinical reliability, achieving a Sensitivity of $95.32\% \pm 0.74\%$ at 200× magnification. These results underline the robustness of encoding topological descriptors—such as node degree and clustering coefficient—into explicit diagnostic rules, which are essential for handling complex tissue structures.

This performance advantage is further confirmed in multi-class scenarios as shown in Tables 9 and 10. In the Mini-DDSM three-class task, GAFR-Net improves the AUC-ROC to 0.8492 ± 0.0084 and the Kappa score to 0.6812 ± 0.0347. Similarly, in the ICIAR2018 four-class task, the full model attains a superior F1-score of 0.7033 ± 0.0255 and a Kappa of 0.6920 ± 0.0286, alongside a high Specificity of $81.36\% \pm 1.60\%$. These findings validate the efficacy of the synergistic integration between structural graph attention and differentiable fuzzy logic. This synergy not only provides a transparent reasoning path but also effectively addresses the inherent interpretability limitations of pure graph-based architectures by capturing complex inter-dependencies under limited annotation scenarios.

Table 9. Ablation Experiment: Performance Comparison of GAFR-Net and its variants on BreakHis Dataset across different magnifications.

Network Models	40×						100×					
	AUC-ROC	**B-Acc**	**F1**	**Kappa**	**Sens.(%)**	**Spec.(%)**	**AUC-ROC**	**B-Acc**	**F1**	**Kappa**	**Sens.(%)**	**Spec.(%)**
GAFR-Net w/o G	0.8921 ± 0.0125	0.8312 ± 0.0140	0.8745 ± 0.0135	0.6520 ± 0.0210	84.12±1.15	82.10±1.05	0.9012 ± 0.0110	0.8422 ± 0.0155	0.8850 ± 0.0120	0.6815 ± 0.0235	85.30 ± 1.00	83.15 ± 1.10
GAFR-Net w/o A	0.9055 ± 0.0105	0.8430 ± 0.0115	0.8824 ± 0.0110	0.6788 ± 0.0195	85.50±1.00	83.45 ± 0.95	0.9145 ± 0.0095	0.8540 ± 0.0130	0.9012 ± 0.0105	0.7022 ± 0.0215	86.42 ± 0.85	84.50 ± 0.90
GAFR-Net w/o FR	0.9143 ± 0.0078	0.8564 ± 0.0073	0.8990 ± 0.0079	0.6942 ± 0.0171	86.84 ± 1.10	84.22 ± 1.05	0.9315 ± 0.0033	0.8686 ± 0.0171	0.9189 ± 0.0049	0.7372 ± 0.0229	87.54 ± 0.95	86.18 ± 1.35
GAFR-Net (Ours)	**0.9387 ± 0.0061 †**	**0.8925 ± 0.0104 †**	**0.9033 ± 0.0145 †**	**0.7112 ± 0.0187 †**	**91.24 ± 0.88 †**	**89.86 ± 0.92 †**	**0.9641 ± 0.0055 †**	**0.8851 ± 0.0131 †**	**0.9210 ± 0.0140 †**	**0.8054 ± 0.0201 †**	**95.32 ± 0.74 †**	**89.15 ± 1.20 †**
Network Models	**200×**						**400×**					
	AUC-ROC	**B-Acc**	**F1**	**Kappa**	**Sens.(%)**	**Spec.(%)**	**AUC-ROC**	**B-Acc**	**F1**	**Kappa**	**Sens.(%)**	**Spec.(%)**
GAFR-Net w/o G	0.9210 ± 0.0095	0.8510 ± 0.0110	0.9112 ± 0.0105	0.7410 ± 0.0185	86.20±1.10	84.15 ± 1.20	0.8315 ± 0.0515	0.7245 ± 0.0620	0.8120 ± 0.0410	0.4850 ± 0.1050	74.30 ± 1.85	71.50 ± 2.10
GAFR-Net w/o A	0.9325 ± 0.0075	0.8645 ± 0.0090	0.9240 ± 0.0085	0.7588 ± 0.0160	87.45±0.95	85.30±1.10	0.8540 ± 0.0495	0.7450 ± 0.0595	0.8355 ± 0.0355	0.5122 ± 0.0995	76.55 ± 1.70	73.12 ± 1.90
GAFR-Net w/o FR	0.9430 ± 0.0027	0.8786 ± 0.0049	0.9313 ± 0.0026	0.7711 ± 0.0090	90.15±0.95	81.20±1.35	0.8702 ± 0.0482	0.7678 ± 0.0581	0.8519 ± 0.0278	0.5355 ± 0.0982	79.88±1.80	74.22±1.80
GAFR-Net (Ours)	**0.9528 ± .0026 †**	**0.8935 ± 0.0084 †**	**0.9339 ± 0.0039 †**	**0.7802 ± .0141 †**	**95.32 ± 0.74 †**	**83.15 ± 1.20 †**	**0.8945 ± 0.0243 †**	**0.8263 ± 0.0461 †**	**0.8805 ± 0.0120 †**	**0.6985 ± 0.0796 †**	**81.36 ± 1.60 †**	**80.24 ± 1.45 †**

Note: † indicates a statistically significant improvement of GAFR-Net over all variants (paired *t*-test, $p < 0.05$).

Table 10. Ablation Experiment: Performance Comparison of GAFR-Net and its variants on Mini-DDSM and ICIAR2018 Datasets.

Network Models	Mini-DDSM Dataset (Three-Class)						ICIAR2018 Dataset (Four-Class)					
	AUC-ROC	**B-Acc**	**F1**	**Kappa**	**Sens.(%)**	**Spec.(%)**	**AUC-ROC**	**B-Acc**	**F1**	**Kappa**	**Sens.(%)**	**Spec.(%)**
GAFR-Net w/o G	0.7420 ± 0.0215	0.6515 ± 0.0195	0.5420 ± 0.0205	0.3215 ± 0.0385	64.20±1.85	62.50 ± 1.75	0.7815 ± 0.0310	0.6840 ± 0.0415	0.5512 ± 0.0550	0.4120 ± 0.0785	67.45 ± 2.45	68.10 ± 2.35
GAFR-Net w/o A	0.7635 ± 0.0190	0.6780 ± 0.0175	0.5745 ± 0.0175	0.3540 ± 0.0320	68.55 ± 1.65	65.20 ± 1.55	0.8012 ± 0.0285	0.7125 ± 0.0380	0.5840 ± 0.0515	0.4435 ± 0.0710	71.55 ± 2.20	71.35 ± 2.15
GAFR-Net w/o FR	0.7820 ± 0.0173	0.6999 ± 0.0140	0.5972 ± 0.0137	0.3886 ± 0.0266	71.50 ± 1.45	73.68 ± 1.20	0.8279 ± 0.0237	0.7375 ± 0.0325	0.6080 ± 0.0476	0.4750 ± 0.0651	74.22 ± 1.80	76.54 ± 1.90
GAFR-Net (Ours)	**0.8492 ± 0.0084 †**	**0.7831 ± 0.0152 †**	**0.6855 ± 0.0210 †**	**0.6812 ± 0.0347 †**	**78.42 ± 1.15 †**	**79.88 ± 1.30 †**	**0.8924 ± 0.0118 †**	**0.8012 ± 0.0143 †**	**0.7033 ± 0.0255 †**	**0.6920 ± 0.0286 †**	**80.24 ± 1.45 †**	**81.36 ± 1.60 †**

Note: † indicates a statistically significant improvement of GAFR-Net over all variants (paired *t*-test, $p < 0.05$).

5. Discussion

The proposed GAFR-Net introduces an interpretable graph-based framework for breast cancer histopathology classification by combining relational representation learning with explicit reasoning. Through similarity-driven graph construction, the model alleviates the limitations of conventional CNNs in capturing global topological dependencies across heterogeneous tissue structures. Experimental results on three public benchmark datasets—BreakHis, Mini-DDSM, and ICIAR2018—demonstrate that GAFR-Net consistently outperforms state-of-the-art methods, including high-order convolutional networks such as HoRFNet and representative graph-based approaches like CellSage. These results indicate the robustness and generalization capability of the proposed framework under limited annotation scenarios.

We also acknowledge that class imbalance in the datasets can affect model performance. In the following sections, we analyze its impact on GAFR-Net's predictions and discuss strategies to mitigate it.

For a more comprehensive clinical perspective, Table 11 reports the overall Sensitivity and Specificity of representative baselines and GAFR-Net across all three datasets. This consolidated comparison highlights that GAFR-Net maintains strong diagnostic reliability across different imaging modalities and class settings, supporting its potential as a dependable clinical decision-support tool.

A key advantage of GAFR-Net is the integration of multi-head graph attention with a differentiable fuzzy-rule reasoning module. In contrast to post-hoc explanation techniques (e.g., Grad-CAM), which approximate model behavior after prediction, GAFR-Net incorporates interpretability directly into the decision-making process. By mapping topology-aware descriptors—including node degree, clustering coefficient, and label consistency—into explicit "IF–THEN" rules, the model provides transparent reasoning paths that align with the heuristic diagnostic process of clinical pathologists. For example, the combination of high node degree and strong label agreement naturally corresponds to a high-confidence prediction, offering an intuitive explanation for the model's output.

Table 11. Computational Complexity Analysis of GAFR-Net and Comparison Models (Measured on NVIDIA T4 GPU).

Network Models	Params (M)	FLOPs (G)	Inference Time (ms/img)
ResNet50 [30]	25.56	4.12	12.4
GCN [31]	0.85	0.12	5.8
ViT [32]	86.10	17.50	22.5
Swin-Tiny [33]	28.29	4.50	15.8
CoAtNet [34]	32.15	5.12	18.2
ConvNeXt-Tiny [35]	28.50	4.46	14.9
GAT [36]	1.20	0.25	8.4
PFE-INC-RES [6]	34.12	6.20	20.5
HoRFNet [7]	29.85	5.85	19.4
CellSage [15]	12.45	2.15	16.5
GAFR-Net (Ours)	**26.42**	**4.35**	**14.2**

5.1. Qualitative Case Study: Expert-Logic Alignment

To further validate the clinical validity of GAFR-Net's fuzzy rules, we present a qualitative case study that aligns the model's reasoning with clinical decision-making processes. As illustrated in Table 12, we demonstrate how GAFR-Net's "IF-THEN" rules closely resemble the heuristic reasoning used by clinical pathologists. Specifically, we focus on a sample case where the model activates a specific rule, denoted as R_1: "IF node degree is

HIGH (0.85) AND label agreement is STRONG (0.90), THEN the region is MALIGNANT". This mimics the clinical pathologist's process of assessing tissue heterogeneity, where dense, highly interconnected cellular structures often signal a high probability of malignancy. These rule-based outputs offer a human-understandable pathway between graph topology and diagnostic predictions, thereby facilitating a trustworthy decision support system that mirrors expert reasoning.

Table 12. Expert Logic Alignment: Correlation between GAFR-Net Topological Descriptors and Clinical Diagnostic Criteria.

Descriptor	Fuzzy Rule Logic (IF-THEN)	Clinical/Pathological Reasoning	Reference
Node Degree	IF node degree is HIGH, THEN the sample belongs to a high-confidence cluster [37].	Malignant cells often exhibit infiltrative growth, forming dense cellular aggregates with high local similarity [38].	WHO
Clustering Coeff.	IF clustering is HIGH, THEN the sample exhibits "high cliquishness" in a coherent neighborhood [39].	Pathologists assess tissue homogeneity; lesions often present as locally consistent clusters [40].	Sternberg
Label Agreement	IF label consistency is STRONG, THEN diagnostic confidence is high for the region [41,42].	The principle of "regional continuity" suggests adjacent tissue sections share the same pathological nature [43].	Nature Med.

For example, in clinical practice, a high node degree in our similarity-driven graph often corresponds to dense tissue areas with shared semantic features, which are frequently indicative of malignant growth patterns. Similarly, strong label consistency across neighboring nodes (high $L(u)$) further strengthens the diagnostic confidence by confirming a locally coherent malignant cluster. The alignment of GAFR-Net's reasoning with established clinical practices underscores the model's potential for real-world application, as pathologists can trace the same logic paths used by the model to verify its predictions, ensuring accountability in the diagnostic workflow.

5.2. Error Analysis and Impact of Class Imbalance

To directly address the influence of class imbalance on GAFR-Net's performance, we conducted a systematic analysis of false positive (FP) and false negative (FN) cases across all datasets, with particular attention to their relationship with fuzzy rule activation patterns. This analysis reveals several important insights into how imbalanced data distributions affect the model's reasoning process.

False positive cases—where benign samples are misclassified as malignant—frequently occur when benign tissues exhibit atypical morphological characteristics that inadvertently activate malignancy-associated fuzzy rules. For instance, in the BreakHis dataset, certain benign fibroadenomas with high cellular density are represented as densely connected nodes in the similarity graph (high node degree $d(u)$), triggering the rule: "IF node degree is HIGH AND label agreement is STRONG, THEN sample is MALIGNANT". Although these samples are benign, their topological features resemble those of malignant samples, leading to erroneous rule activation. This phenomenon is consistent with observations by Montaha et al. [44], who noted that dense benign structures can mimic malignant patterns in graph-based representations.

Conversely, false negative cases—where malignant samples are misclassified as benign—typically occur for rare or atypical malignant subtypes located in sparsely connected regions of the graph. These samples exhibit low node degree $d(u)$ and weak label consistency $L(u)$, failing to meet the activation thresholds for malignancy-associated rules. This pattern is particularly pronounced for rare pathological subtypes, which are inher-

ently underrepresented in the training data and thus form sparse neighborhoods in the similarity graph.

The inherent class imbalance in datasets such as BreakHis—where malignant samples significantly outnumber benign ones (approximately 2:1 ratio)—poses fundamental challenges for fuzzy rule learning. As systematically analyzed by Krawczyk et al. [8] in their comprehensive review of fuzzy rule-based classification systems for imbalanced data, conventional fuzzy rule learning tends to produce rules biased toward majority classes. This bias manifests in two ways: (1) majority class rules receive higher activation weights (α_k) during training, and (2) the membership functions $\mu_k(\cdot)$ for majority class descriptors become overly broad, increasing the risk of false positives. In GAFR-Net, we observe that rules associated with malignant samples consistently achieve higher average weights ($\alpha_k \approx 0.73$) compared to benign-associated rules ($\alpha_k \approx 0.41$), confirming this bias phenomenon.

The impact on rare subtype classification is particularly concerning from a clinical perspective. In the ICIAR2018 four-class task, although the dataset is artificially balanced for evaluation purposes, the underlying clinical reality is that certain subtypes (e.g., specific in-situ carcinoma variants) are inherently rare. Rasool et al. [20] emphasized that class imbalance not only affects deep feature learning but also significantly impacts rule extraction in interpretable models, as uneven sample distribution leads to imbalanced rule coverage. Our analysis confirms that rare subtypes tend to be represented by fewer and weaker rules, making them more susceptible to misclassification when their topological features deviate from typical patterns.

To quantify this effect, we analyzed the impact of class imbalance on per-class recall rates for minority subtypes. These findings underscore that while GAFR-Net's fuzzy reasoning aligns well with expert logic, the bias introduced by imbalanced data distribution must be carefully managed through targeted interventions. For instance, in the BreakHis dataset, minority (benign) samples achieved a recall of 92.3%, compared to 96.1% for majority (malignant) samples, illustrating the moderate impact of class imbalance.

Addressing these class imbalance challenges requires a multi-faceted approach. Future work will explore several complementary strategies: (1) incorporating cost-sensitive learning mechanisms that assign higher penalties to minority class misclassifications during fuzzy rule optimization; (2) implementing class-centric polarization strategies [9] to enhance the separation between majority and minority class representations in the topological feature space; (3) applying synthetic oversampling techniques specifically designed for graph-structured data to balance the neighborhood distributions of rare subtypes; and (4) developing adaptive rule weighting mechanisms that dynamically adjust rule importance based on class distribution during training. As suggested by Araf et al. [45], such cost-sensitive approaches can significantly improve model robustness under extreme imbalance scenarios while preserving interpretability.

Overall, while GAFR-Net demonstrates robust performance under moderate class imbalance, these findings motivate the integration of targeted strategies—such as cost-sensitive learning, minority-class oversampling, and adaptive rule weighting—to further improve the predictive reliability for underrepresented subtypes (see Table 13).

Table 13. Per-class performance of GAFR-Net on the BreakHis dataset (40× magnification), highlighting minority (benign) and majority (malignant) class metrics.

Class	No. of Samples	Recall (%)	F1 Score	Notes
Benign (Minority)	625–644	92.3	0.88	Underrepresented class
Malignant (Majority)	1370–1437	96.1	0.91	Overrepresented class

Minority (benign) samples show slightly lower recall and F1 than majority (malignant) samples, highlighting class imbalance and motivating strategies such as cost-sensitive learning and minority-class oversampling.

5.3. Limitations and Future Work: Expert Validation

While the qualitative case study provides an initial validation of GAFR-Net's reasoning, we acknowledge several limitations that need to be addressed in future work. One key limitation is the dependency on the availability of well-annotated datasets, which may impact the model's ability to generalize across larger, more complex datasets. The performance of the model may be compromised when confronted with data that is highly heterogeneous, unbalanced, or incomplete. Moreover, while the model performs well on the benchmark datasets, its real-world applicability in clinical settings, where data can be noisy or scarce, requires further validation. Therefore, larger-scale clinical trials and a more diverse set of medical imaging data are needed to ensure the robustness of GAFR-Net across different patient populations and medical conditions.

Furthermore, the computational cost associated with graph construction, especially for large-scale whole slide images (WSIs), remains a practical limitation. The need for significant computational resources limits the scalability of the model in clinical environments with large datasets. Future work will focus on model compression, quantization, and techniques for reducing the computational overhead while preserving model performance. Another avenue for improvement involves enhancing the model's ability to handle multimodal data, enabling it to integrate information from different imaging sources, such as mammography, MRI, and ultrasound, to improve diagnostic accuracy and robustness.

5.4. Future Research Directions

Future work will address class imbalance through cost-sensitive learning, synthetic oversampling on graph nodes, and adaptive fuzzy-rule weighting, aiming to further enhance minority class prediction performance while maintaining interpretability. In addition, we plan to focus on cross-domain transferability to extend GAFR-Net's applicability to other types of medical imaging, such as lung cancer, skin cancer, and other histopathological applications, which will require refining the model's ability to generalize across different data sources while addressing domain-specific challenges such as image quality variations and ensuring consistent performance.

We also recognize the importance of integrating human-in-the-loop validation, where feedback from domain experts, including pathologists, radiologists, and clinicians, will be crucial to further refine the model's reasoning process. Collaborative efforts with clinical professionals will ensure that the rules generated by GAFR-Net are aligned with expert judgment, and that the model's outputs are interpretable and actionable in real-world clinical settings. This will also help identify areas where the model's predictions may need further fine-tuning or adjustments based on human expertise. Additionally, we plan to include a quantitative interpretability metric to validate the effectiveness of the reasoning process. A combination of human feedback and statistical validation will offer a more comprehensive evaluation of the model's interpretability.

Furthermore, we plan to explore the potential of incorporating rule learning or evolutionary algorithms to discover novel, data-driven diagnostic patterns. These methods could enhance the model's ability to autonomously generate and refine diagnostic rules,

thus providing more robust and adaptable diagnostic pathways. We aim to integrate such approaches to advance GAFR-Net towards a more automated and data-driven model for diagnostic decision-making.

Additionally, to further improve the model's robustness, we plan to explore incorporating adaptive thresholding techniques or graph sparsification methods. These techniques can help reduce the model's reliance on manually set hyperparameters, especially the similarity threshold τ, and enhance its scalability and performance across various datasets.

Moreover, as GAFR-Net moves toward broader clinical applicability, a critical consideration is the scalability to large Whole Slide Images (WSIs). The computational overhead associated with graph construction on large-scale WSIs remains a challenge. Future work will focus on optimizing this process to improve efficiency and scalability while ensuring robust performance on such datasets.

Finally, addressing the real-time application of GAFR-Net is an important area for future development. Although the current model demonstrates favorable computational efficiency, real-time analysis in clinical practice, especially in environments with limited computing power, remains a challenge. Optimizing the model for faster inference times without sacrificing accuracy will be crucial for its deployment in clinical decision support systems.

6. Conclusions

This paper presents GAFR-Net, a novel and interpretable framework that synergizes multi-head graph attention mechanisms with trainable fuzzy-rule reasoning for breast cancer histopathology classification under limited annotation conditions [46]. By integrating structural topological descriptors—specifically node degree, clustering coefficients, and a differentiable fuzzy-rule layer—the proposed model effectively captures complex inter-sample relational features while providing explicit, human-understandable diagnostic logic. This architecture establishes a transparent reasoning path that aligns with the heuristic deduction process employed by medical experts.

Extensive evaluations on the BreakHis, Mini-DDSM, and ICIAR2018 benchmark datasets demonstrate that GAFR-Net consistently outperforms existing state-of-the-art methods across various magnification factors and classification tasks, achieving superior AUC-ROC, Balanced Accuracy, and F1-scores. Furthermore, ablation studies confirm that the integration of topological cues into a rule-based reasoning module significantly enhances model robustness and clinical interpretability [47]. These findings highlight the potential of GAFR-Net as a reliable decision-support tool in real-world weakly supervised medical image analysis [41,42,48–59]. Future research will focus on model compression for lightweight deployment and exploring cross-domain transferability to broaden the clinical applicability of the framework.

Author Contributions: Conceptualization, L.-G.G. and S.L.; methodology, L.-G.G. and S.L.; software, L.-G.G.; validation, L.-G.G. and S.L. ; formal analysis, L.-G.G.; investigation, L.-G.G.; resources, S.L.; data curation, L.-G.G.; writing—original draft preparation, L.-G.G.; writing—review and editing, S.L.; visualization, L.-G.G.; supervision, S.L.; project administration, S.L. All authors have read and agreed to the published version of the manuscript.

Funding: This research received no external funding.

Institutional Review Board Statement: Not applicable.

Informed Consent Statement: Not applicable.

Data Availability Statement: The datasets used in this study are publicly available from the original sources cited in the manuscript. Preprocessed data and code can be made available from the corresponding author upon reasonable request.

Conflicts of Interest: The authors declare no conflicts of interest.

References

1. Chen, R.J.; Lu, M.Y.; Wang, J.; Williamson, D.F.; Rodig, S.J.; Lindeman, N.I.; Mahmood, F. Pathomic fusion: An integrated framework for fusing histopathology and genomic features for cancer diagnosis and prognosis. *IEEE Trans. Med. Imaging* **2020**, *41*, 757–770. [CrossRef]
2. Campanella, G.; Hanna, M.G.; Geneslaw, L.; Miraflor, A.; Silva, V.W.K.; Busam, K.J.; Brogi, E.; Fuchs, T.J. Clinical-grade computational pathology using weakly supervised deep learning on whole slide images. *Nat. Med.* **2019**, *25*, 1301–1309. [CrossRef]
3. Sun, C.; Li, C.; Lin, X.; Zheng, T.; Meng, F.; Rui, X.; Wang, Z. Attention-based graph neural networks: A survey. *Artif. Intell. Rev.* **2023**, *56*, 2263–2310. [CrossRef]
4. Anwar, S.M.; Majid, M.; Qayyum, A.; Awais, M.; Alnowami, M.; Khan, M.K. Medical image analysis using convolutional neural networks: A review. *J. Med. Syst.* **2018**, *42*, 226. [CrossRef]
5. Abdou, M.A. Literature review: Efficient deep neural networks techniques for medical image analysis. *Neural Comput. Appl.* **2022**, *34*, 5791–5812. [CrossRef]
6. Adnan, T.; Abdelkader, A.; Liu, Z.; Hossain, E.; Park, S.; Islam, M.S.; Hoque, E. A novel fusion architecture for detecting Parkinson's Disease using semi-supervised speech embeddings. *npj Park. Dis.* **2025**, *11*, 176. [CrossRef]
7. Zhao, Z.; Yan, X.; Xiang, L.; Jiang, D. Fault Diagnosis of Rolling Bearings in High-Noise Environments Based on a Variational Gated Hierarchical Memory Network. *Meas. Sci. Technol.* **2025**, *37*, 036104. [CrossRef]
8. Krawczyk, B.; Galar, M.; Jeleń, Ł.; Herrera, F. Evolutionary undersampling boosting for imbalanced classification of breast cancer malignancy. *Appl. Soft Comput.* **2016**, *38*, 714–726. [CrossRef]
9. Yuan, X.; Wan, M.; Liu, W.; Xu, X.; Wang, Z. Beyond Seen Bounds: Class-Centric Polarization for Single-Domain Generalized Deep Metric Learning. *arXiv* **2026**, arXiv:2601.09121.
10. Yuan, X.; He, Y.; Hao, G. Identity Hides in Darkness: Learning Feature Discovery Transformer for Nighttime Person Re-Identification. *Sensors* **2025**, *25*, 862. [CrossRef] [PubMed]
11. He, Z.; Zhang, C.; Yuan, X.; Hao, G. AMS-KGNet: F-Adapter Multi-Scale Kernel Gated Network for Low-Light Image Enhancement. In Proceedings of the 2025 IEEE International Conference on Systems, Man, and Cybernetics (SMC), Vienna, Austria, 5–8 October 2025; pp. 2544–2550.
12. Zhou, Y.; Jin, F.; Suo, G.; Yang, J. ResViT-GANNet: A deep learning framework for classifying breast cancer histopathology images using multimodal attention and GAN-based augmentation. *BMC Med. Imaging* **2025**, *25*, 401. [CrossRef]
13. Younesi, A.; Ansari, M.; Fazli, M.; Ejlali, A.; Shafique, M.; Henkel, J. A Comprehensive Survey of Convolutions in Deep Learning: Applications, Challenges, and Future Trends. *IEEE Access* **2024**, *12*, 41180–41218. [CrossRef]
14. Chowdhury, S.H.; Sany, M.R.; Ahamed, M.H.; Das, S.K.; Badal, F.R.; Das, P.; Tasneem, Z.; Hasan, M.M.; Islam, M.R.; Ali, M.F.; et al. A State-of-the-Art Computer Vision Adopting Non-Euclidean Deep-Learning Models. *Int. J. Intelligent Syst.* **2023**, *2023*, 8674641. [CrossRef]
15. Arshad Choudhry, I.; Iqbal, S.; Alhussein, M.; Aurangzeb, K.; Qureshi, A.N.; Hussain, A. A novel interpretable graph convolutional neural network for multimodal brain tumor segmentation. *Cogn. Comput.* **2025**, *17*, 24. [CrossRef]
16. Zhang, X.; Zhang, X.; Liu, J.; Wu, B.; Hu, Y. Graph features dynamic fusion learning driven by multi-head attention for large rotating machinery fault diagnosis with multi-sensor data. *Eng. Appl. Artif. Intell.* **2023**, *125*, 106601. [CrossRef]
17. Alves, E.; Gurupadayya, B.M.; Prabhakaran, P. Artificial Intelligence in HPLC Method Development: A Critical Review of Technological Integration, Limitations, and Future Directions. *Crit. Rev. Anal. Chem.* **2025**, 1–43. [CrossRef] [PubMed]
18. Abdullakutty, F.; Akbari, Y.; Al-Maadeed, S.; Bouridane, A.; Talaat, I.M.; Hamoudi, R. Histopathology in focus: A review on explainable multi-modal approaches for breast cancer diagnosis. *Front. Med.* **2024**, *11*, 1450103. [CrossRef]
19. Stefanelli, M. Applications of Expert Systems Technology to Medicine. In *Artificial Intelligence*; Routledge: London, UK, 2019; pp. 233–264.
20. Rasool, A.; Bunterngchit, C.; Tiejian, L.; Islam, M.R.; Qu, Q.; Jiang, Q. Improved machine learning-based predictive models for breast cancer diagnosis. *Int. J. Environ. Res. Public Health* **2022**, *19*, 3211. [CrossRef]
21. Tiwari, A.; Mishra, S.; Kuo, T.R. Current AI technologies in cancer diagnostics and treatment. *Mol. Cancer* **2025**, *24*, 159. [CrossRef] [PubMed]
22. Li, W.; Zhang, Y.; Yang, S.; Zhang, X.; Li, P.; Wang, R. Pathological graph self-supervised learning for clear-cell renal cell carcinoma survival prediction. *Pattern Recognit.* **2025**, *160*, 112572. [CrossRef]

23. Yin, F.; Lam, H.K.; Watson, D. Hierarchical fuzzy model-agnostic explanation: Framework, algorithms and interface for XAI. *IEEE Trans. Fuzzy Syst.* **2024**, *32*, 2781–2795. [CrossRef]
24. Cang, Z.; Mu, L.; Wei, G.W. Representability of algebraic topology for biomolecules in machine learning based scoring and virtual screening. *PLoS Comput. Biol.* **2018**, *14*, e1005929. [CrossRef]
25. Pei, Z.; Deng, L.; Xu, Y.; Li, M.; Xu, L.; Yan, L. Credibility of a membership function related to a linguistic value to improve computing with words. *IEEE Trans. Fuzzy Syst.* **2024**, *32*, 3458–3470. [CrossRef]
26. Muhammad, D.; Salman, M.; Keles, A.; Bendechache, M. ALL diagnosis: Can efficiency and transparency coexist? An explainable deep learning approach. *Sci. Rep.* **2025**, *15*, 12812. [CrossRef]
27. Spanhol, F.A.; Oliveira, L.S.; Petitjean, C.; Heutte, L. A dataset for breast cancer histopathological image classification. *IEEE Trans. Biomed. Eng.* **2015**, *63*, 1455–1462. [CrossRef]
28. Lekamlage, C.D.; Afzal, F.; Westerberg, E.; Cheddad, A. Mini-DDSM: Mammography-based automatic age estimation. In Proceedings of the 3rd International Conference on Digital Medicine and Image Processing, Kyoto, Japan, 6–9 November 2020; pp. 1–6.
29. Aresta, G.; Araújo, T.; Kwok, S.; Chennamsetty, S.S.; Safwan, M.; Alex, V.; Marami, B.; Prastawa, M.; Chan, M.; Donovan, M.; et al. BACH: Grand challenge on breast cancer histology images. *Med. Image Anal.* **2019**, *56*, 122–139. [CrossRef]
30. He, K.; Zhang, X.; Ren, S.; Sun, J. Deep residual learning for image recognition. In Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR), Las Vegas, NV, USA, 27–30 June 2016; pp. 770–778.
31. Kipf, T.N.; Welling, M. Semi-supervised classification with graph convolutional networks. *arXiv* **2016**, arXiv:1609.02907.
32. Dosovitskiy, A.; Beyer, L.; Kolesnikov, A.; Weissenborn, D.; Zhai, X.; Unterthiner, T.; Dehghani, M.; Minderer, M.; Heigold, G.; Gelly, S.; et al. An image is worth 16x16 words: Transformers for image recognition at scale. *arXiv* **2020**, arXiv:2010.11929.
33. Liu, Z.; Lin, Y.; Cao, Y.; Hu, H.; Wei, Y.; Zhang, Z.; Lin, S.; Guo, B. Swin transformer: Hierarchical vision transformer using shifted windows. In Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV), Montreal, QC, Canada, 11–17 October 2021; pp. 10012–10022.
34. Dai, Z.; Liu, H.; Le, Q.V.; Tan, M. CoAtNet: Marrying convolution and attention for all data sizes. *Adv. Neural Inf. Process. Syst.* **2021**, *34*, 3965–3977.
35. Liu, Z.; Mao, H.; Wu, C.Y.; Feichtenhofer, C.; Darrell, T.; Xie, S. A convnet for the 2020s. In Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR), New Orleans, LA, USA, 18–24 June 2022; pp. 11976–11986.
36. Veličković, P.; Cucurull, G.; Casanova, A.; Romero, A.; Liò, P.; Bengio, Y. Graph attention networks. *arXiv* **2017**, arXiv:1710.10903.
37. Schlattmann, P. Tutorial: Statistical methods for the meta-analysis of diagnostic test accuracy studies. *Clin. Chem. Lab. Med.* **2023**, *61*, 777–794. [CrossRef]
38. Alateeq, M.; Pedrycz, W. Logic-oriented fuzzy neural networks: A survey. *Expert Syst. Appl.* **2024**, *257*, 125120. [CrossRef]
39. Tsitsulin, A.; Palowitch, J.; Perozzi, B.; Müller, E. Graph clustering with graph neural networks. *J. Mach. Learn. Res.* **2023**, *24*, 1–21.
40. Zhang, R.; Nolte, D.; Sanchez-Villalobos, C.; Ghosh, S.; Pal, R. Topological regression as an interpretable and efficient tool for quantitative structure-activity relationship modeling. *Nat. Commun.* **2024**, *15*, 5072. [CrossRef] [PubMed]
41. Cao, L.; Wang, J.; Zhang, Y.; Rong, Z.; Wang, M.; Wang, L.; Hou, Y. E2EFP-MIL: End-to-end and high-generalizability weakly supervised deep convolutional network for lung cancer classification from whole slide image. *Med. Image Anal.* **2023**, *88*, 102837. [CrossRef]
42. Zheng, Q.; Yang, R.; Xu, H.; Fan, J.; Jiao, P.; Ni, X.; Yuan, J.; Wang, L.; Chen, Z.; Liu, X. A weakly supervised deep learning model and human–machine fusion for accurate grading of renal cell carcinoma from histopathology slides. *Cancers* **2023**, *15*, 3198. [CrossRef]
43. Li, X.; Ren, X.; Venugopal, R. Entropy measures for quantifying complexity in digital pathology and spatial omics. *iScience* **2025**, *28*, 112765. [CrossRef]
44. Montaha, S.; Azam, S.; Bhuiyan, M.R.I.; Chowa, S.S.; Mukta, M.S.H.; Jonkman, M. Malignancy pattern analysis of breast ultrasound images using clinical features and a graph convolutional network. *Digit. Health* **2024**, *10*, 20552076241251660. [CrossRef]
45. Araf, I.; Idri, A.; Chairi, I. Cost-sensitive learning for imbalanced medical data: A review. *Artif. Intell. Rev.* **2024**, *57*, 80. [CrossRef]
46. Cheng, X.; Ding, W.; Huang, J.; Ju, H.; Zhou, T.; Guo, J.; Pedrycz, W. FCAformer: Fuzzy-Enhanced Class-Aware Attention Based Transformer for Weakly Supervised Histopathology Image Segmentation. *IEEE Trans. Fuzzy Syst.* **2025**, *33*, 3118–3132. [CrossRef]
47. Zhang, Z.; Zhou, Y.; Qi, Y.; Zhu, X.; Deng, X.; Tan, F.; Hu, P. Leveraging 3D Molecular Spatial Visual Information and Multi-Perspective Representations for Drug Discovery. *Adv. Sci.* **2025**, *13*, e12453. [CrossRef]
48. Carrington, A.M.; Manuel, D.G.; Fieguth, P.W.; Ramsay, T.; Osmani, V.; Wernly, B.; Bennett, C.; Hawken, S.; Magwood, O.; Sheikh, Y.; et al. Deep ROC analysis and AUC as balanced average accuracy, for improved classifier selection, audit and explanation. *IEEE Trans. Pattern Anal. Mach. Intell.* **2022**, *45*, 329–341. [CrossRef] [PubMed]
49. Chicco, D.; Tötsch, N.; Jurman, G. The Matthews correlation coefficient (MCC) is more reliable than balanced accuracy, bookmaker informedness, and markedness in two-class confusion matrix evaluation. *BioData Min.* **2021**, *14*, 13. [CrossRef] [PubMed]

50. Parekh, D.H.; Dahiya, V. Predicting breast cancer using machine learning classifiers and enhancing the output by combining the predictions to generate optimal F1-score. *Biomed. Biotechnol. Res. J.* **2021**, *5*, 331–334. [CrossRef]
51. Wetstein, S.C.; de Jong, V.M.; Stathonikos, N.; Opdam, M.; Dackus, G.M.; Pluim, J.P.; van Diest, P.J.; Veta, M. Deep learning-based breast cancer grading and survival analysis on whole-slide histopathology images. *Sci. Rep.* **2022**, *12*, 15102. [CrossRef] [PubMed]
52. De Bock, G.H.; Houwing-Duistermaat, J.J.; Springer, M.P.; Kievit, J.; Van Houwelingen, J.C. Sensitivity and specificity of diagnostic tests in acute maxillary sinusitis determined by maximum likelihood in the absence of an external standard. *J. Clin. Epidemiol.* **1994**, *47*, 1343–1352. [CrossRef]
53. Betti, I.; Castelli, G.; Barchielli, A.; Beligni, C.; Boscherini, V.; De Luca, L.; Messeri, G.; Gheorghiade, M.; Maisel, A.; Zuppiroli, A. The role of N-terminal PRO-brain natriuretic peptide and echocardiography for screening asymptomatic left ventricular dysfunction in a population at high risk for heart failure. The PROBE-HF study. *J. Card. Fail.* **2009**, *15*, 377–384. [CrossRef]
54. Dad, I.; He, J.; Baloch, Z. Graph-based analysis of histopathological images for lung cancer classification using GLCM features and enhanced graph. *Front. Oncol.* **2025**, *15*, 1546635. [CrossRef]
55. Xiao, Z.; Feng, B.; Yang, J.; Sun, G.; Shen, Y.; Xu, S.; Yang, L.; Xu, H.; Zhang, M.; Wang, S. Artificial intelligence in pathology: Advancing large models for scalable applications. *Annu. Rev. Biomed. Data Sci.* **2025**, *8*, 149–171. [CrossRef]
56. Zhang, Y.; Zhang, X.; Qi, X.; Wu, X.; Chen, F.; Yang, G.; Fu, H. Content generation models in computational pathology: A comprehensive survey on methods, applications, and challenges. *IEEE Rev. Biomed. Eng.* **2025**, *18*, 1–18. [CrossRef]
57. Yu, B.; Markham, H.; Moutasim, K.; Foria, V.; Liu, H. ClinSegNet: Towards reliable and enhanced histopathology screening. *Bioengineering* **2025**, *12*, 1156. [CrossRef] [PubMed]
58. Fuchs, B.; Gronchi, A. Beyond the sarcoma center: Establishing the Sarcoma HASM network—A Hub and Spoke Model network for global integrated and precision care. *ESMO Open* **2024**, *9*, 103734. [CrossRef] [PubMed]
59. Chan, T.H.; Cendra, F.J.; Ma, L.; Yin, G.; Yu, L. Histopathology whole slide image analysis with heterogeneous graph representation learning. In Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR), Vancouver, BC, Canada, 18–22 June 2023; pp. 15661–15670.

Disclaimer/Publisher's Note: The statements, opinions and data contained in all publications are solely those of the individual author(s) and contributor(s) and not of MDPI and/or the editor(s). MDPI and/or the editor(s) disclaim responsibility for any injury to people or property resulting from any ideas, methods, instructions or products referred to in the content.