Journal of Artificial Intelligence Practice (2026)
Clausius Scientific Press, Canada

DOI: 10.23977/jaip.2026.090108
ISSN 2371-8412 Vol. 9 Num. 1

A Pluggable Hybrid Generative Model for MVI Classification of Hepatocellular Carcinoma Based on ICG Fluorescence Imaging

Zheng Tao

School of Health Science and Engineering, University of Shanghai for Science and Technology, Shanghai, 200093, China

Keywords: Hepatocellular Carcinoma; Microvascular Invasion; ICG Fluorescence Imaging; Generative Data Augmentation; Hybrid Generative Model; Deep Learning

Abstract: Microvascular invasion (MVI) of hepatocellular carcinoma (HCC) is a key factor affecting clinical diagnosis and prognosis, but the classification of MVI based on indocyanine green (ICG) fluorescence images is limited by insufficient clinical data (a total of 213 ICG laparoscopic fluorescence images were used in the experiment) and difficulty in capturing subtle features. To address these problems, this study proposes a pluggable hybrid generative model for MVI classification and verifies the value of generative data augmentation in small-sample medical image classification through multiple experiments. The proposed Hybrid module integrates the detailed fitting ability of LightGAN and the generation stability of DDIM, with a closed-loop iteration optimal interval of 5-10 rounds (excessive iteration leads to performance degradation). Experimental results show that when iterated 5 times on ResNet18, the model achieves an accuracy of 86.92% and an AUC value of 0.8846, which are 4.68% and 0.0546 higher than the peaks of LightGAN and DDIM respectively; the module not only improves the peak performance of a single backbone network but also has good generalization to different backbone networks (e.g., ConvTransResNet achieves an accuracy improvement of 7.48%, ResNet50 achieves an F1-score improvement of 7.58%), narrowing the performance gap between networks; it can effectively improve model accuracy, F1-score and AUC without additional clinical data collection, has a greater improvement effect on weak-performance networks, and reduces the dependence on high-end computing resources in medical scenarios. This study provides a feasible technical framework for small-sample medical image classification tasks and experimental support for the development of subsequent medical image auxiliary diagnosis systems.

1. Introduction

Hepatocellular carcinoma (HCC) ranks as the sixth most commonly diagnosed type of cancer globally, with the third highest mortality rate[1]. Despite the rapid developments of medical technology, that have improved the diagnosis and prognosis of patients with HCC, the prognosis and long-term survival of these patients remain unfavorable[2-4]. One of the primary reasons for this is

the existence of microvascular invasion (MVI). MVI is a factor related to the early recurrence of HCC[5], which leads to the poor long-term survival of patients after radical hepatectomy. With the rapid development of deep learning technology in the field of medical image analysis, end-to-end automatic classification methods[6, 7] have provided new solutions for the diagnosis of MVI in HCC. Existing studies mostly construct classification models based on convolutional neural networks[8-10] (CNNs) and attention mechanisms[11-13], realizing the judgment of MVI status by learning high-dimensional features in Indocyanine Green (ICG) fluorescence images. However, these methods still face two core challenges that restrict their clinical application and promotion.

The first challenge is the prominent problem of data scarcity. The acquisition of ICG fluorescence images relies on professional surgical equipment and standardized operating procedures, and MVI labels must be confirmed by combining postoperative pathological sections. The high labeling cost and long cycle lead to the scarcity of publicly available large-scale datasets. In small-sample scenarios, models are prone to overfitting, resulting in insufficient generalization ability, which makes it difficult to adapt to imaging differences among different hospitals and equipment.

The second challenge is the difficulty in capturing subtle features. MVI-related imaging features (such as tumor emboli in tiny blood vessels, local abnormalities in fluorescence intensity, and irregular changes in vascular branch morphology) are usually weak and easily affected by factors such as tumor heterogeneity, imaging noise, and surgical operation interference. Traditional models are difficult to accurately capture and distinguish these key features from background noise, which directly affects the accuracy of MVI classification.

To address the above problems, this study proposes a pluggable hybrid generative model for MVI classification of HCC using ICG fluorescence images. By integrating the advantages of Generative Adversarial Networks (GAN)[14-16] and Denoising Diffusion Implicit Models[17, 18] (DDIM), the model achieves training data augmentation and feature extraction optimization, and combines a weighted score fusion strategy to improve classification accuracy and clinical applicability. The core contributions of this study are as follows:

(1) A hybrid generative model adapted to the characteristics of ICG fluorescence images is constructed to generate MVI-positive/negative samples in a targeted manner, effectively alleviating the small-sample dilemma. The generated samples are highly consistent with real images in key features such as fluorescence intensity distribution and microvascular morphology.

(2) A closed-loop iterative training mechanism is designed. The classification feedback of the backbone network is used to dynamically optimize the generator parameters and generation strategy, ensuring that the generated samples accurately meet the needs of MVI feature learning, and forming a positive cycle of "generation-training-feedback-optimization-generation".

(3) A multi-model weighted fusion strategy is proposed to integrate the optimal model results of multiple rounds of closed-loop training, balance the classification sensitivity and specificity, reduce the prediction risk of a single model, and improve the stability of results and clinical credibility.

2. Materials and Methods

2.1. Laparoscopic Fluorescence Images Dataset

Laparoscopic fluorescence images were collected from Eastern Hepatobiliary Surgery Hospital between August 2022 and April 2024 using a fluorescence laparoscope[19-21]. The images acquired were cross-sectional fluorescence views obtained intraoperatively from patients diagnosed with HCC. A total of 213 laparoscopic fluorescence images were included in this study: 115 were pathologically classified as M0, and 98 were pathologically classified as non-M0. The sample size was balanced across the two categories to ensure the reliability of subsequent analyses.

2.2. Hybrid Generator Network Structure

This study proposes a pluggable hybrid generative model for MVI classification of HCC using ICG fluorescence images, with the overall workflow illustrated in Fig. 1. The framework fully considers the characteristics of ICG fluorescence images and the clinical requirements of MVI diagnosis, and its specific implementation process is described as follows: first, preprocessing operations including format conversion, denoising enhancement, and size standardization are performed on the original ICG fluorescence images to ensure the consistency and effectiveness of the input data, which are then divided into training set and validation set; second, the pluggable hybrid generative model is pre-trained once, and the generator is dynamically optimized through a closed-loop iterative mechanism that receives classification feedback from the backbone network to generate high-quality MVI-positive/negative samples; these generated samples are mixed with the original training set to construct an augmented dataset, enriching the diversity and representativeness of the training data; subsequently, ResNet18 with a hybrid attention mechanism is adopted as the backbone network to extract MVI-related features from the augmented dataset and complete the binary classification task; finally, the optimal models from multiple rounds of closed-loop training are selected, and their prediction results are integrated through a weighted fusion strategy to obtain the final classification decision. In this way, the hybrid generative model solves the problem of data scarcity, and the attention-enhanced backbone network improves the ability to capture subtle features, realizing more accurate binary classification of HCC MVI.

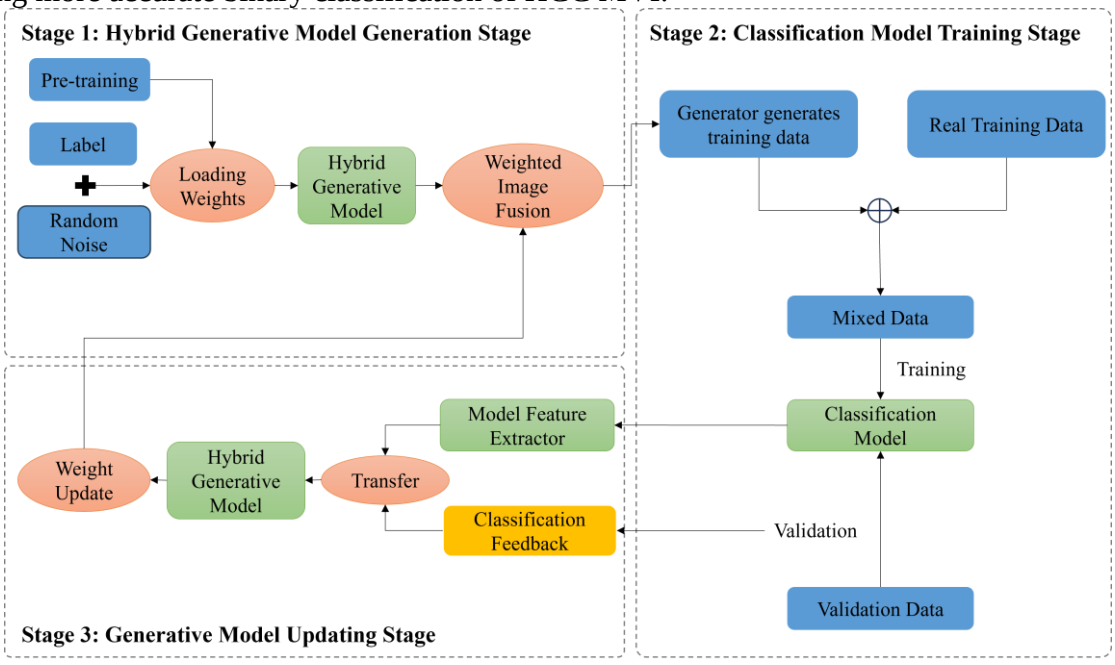

Fig. 1 Flowchart of the Proposed Method

In this experiment, a Hybrid Generator was designed as the generative model for generating target category images in the MVI binary classification task of ICG fluorescence images. A collaborative optimization strategy combining feature constraints and classification guidance was adopted to iteratively optimize the model. The hybrid generative model consists of two sub-modules: a lightweight LightGAN generator and a lightweight DDIM generator, which integrates the advantages of LightGAN (fast generation speed) and DDIM (stable generation quality), and its network structure is shown in Fig. 2.

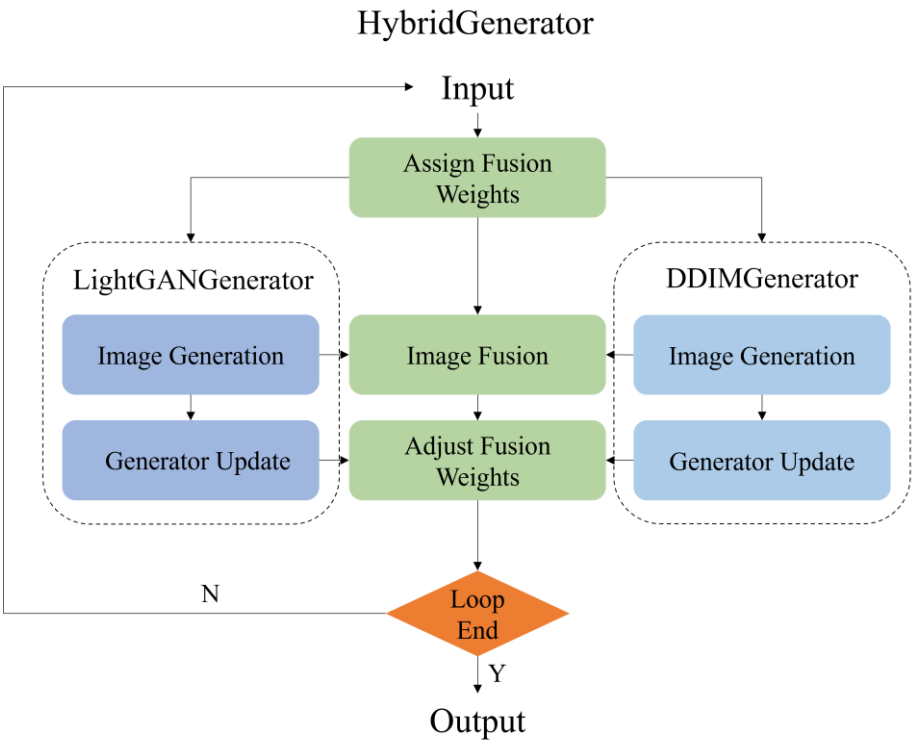

Fig. 2 Network Structure of the Hybrid Generation Module

The core advantage of the hybrid generative model lies in its dynamic fusion weight mechanism. The model can real-time adjust the image generation proportion of the two sub-generators according to the category feature distribution and classification accuracy fed back by the classification network: when the classification accuracy is low, the fusion weight of LightGAN is increased, and the flexibility of its parameter-level gradient update is used to quickly align the mean value of target category features; when the classification accuracy is high, the weight of DDIM is increased, and the stability of its strategy-level hyperparameter adjustment is used to improve the overall quality of generated images.

Based on the above structure, a closed-loop optimization process was designed for the hybrid generative model: first, the model receives category labels, and the two sub-generators generate images of the corresponding categories respectively and complete the fusion; then, using the feedback information from the classification network, gradient backpropagation update (feature loss weight, classification loss weight) of feature loss + classification loss is performed on LightGAN, so that the features of the generated images approach the mean value of the target category, and the category consistency of the generated images is constrained by the classification loss; for DDIM, hyperparameter adjustment (noise loss weight, feature loss weight, classification loss weight) of noise loss + feature loss + classification loss is performed, and the generation pertinence is enhanced by optimizing the classification guidance weight and sampling coefficient. This fusion update strategy completely avoids the problem of gradient dependence on wrong samples, and only drives the generator iteration based on the mean value of category features and classification results, which can accurately capture the feature differences between the two types of images in the HCC MVI binary classification task. The generated images are used to augment the training dataset after effectiveness screening, so as to improve the generalization ability of the subsequent classification model.

2.3. Experimental Setup

To verify the stability and repeatability of the model during training, the dataset was randomly divided into training set and validation set at a ratio of 1:1, as shown in Table. 1. For the supplementary images generated by the generator, they were merged with the original data to construct an augmented dataset after size calibration, range cropping, and normalization. Meanwhile, a fixed random seed was set to further control the randomness of data division and training process, ensuring the reproducibility of experimental results. The core evaluation metrics selected for experimental evaluation included accuracy, precision, recall, F1-score, confusion matrix, ROC curve, and AUC value, to comprehensively measure the model performance.

Table 1 Dataset Distribution

Category	Training Set	Validation Set	Total
MVI-negative	55	60	115
MVI-positive	51	47	98
Total	106	107	213

The experiment was implemented based on the PyTorch framework, with CUDA acceleration prioritized for training. The classification networks included ResNet18, AlexNet, VGG16, etc. The generation modules included three types: LightGAN generation module, DDIM generation module, and hybrid generation module, all of which were pre-trained for one epoch for subsequent module initialization. In terms of training parameters, the batch size was set to 4, the maximum number of closed-loop iterations was 15, the number of model training epochs within each closed-loop was 100, and the gradient accumulation step was 4 to alleviate the memory limitation; the optimizer adopted SGD with an initial learning rate of 0.001, momentum of 0.9, and weight decay of 0.0001, and the learning rate was decayed to 0.9 times the original every 20 epochs through the StepLR scheduler; the loss function adopted cross-entropy loss, and the generator training integrated feature loss, classification loss, etc., and the contribution of different losses was balanced through weight configuration (LightGAN feature loss weight = 0.8, LightGAN classification loss weight = 0.2, DDIM noise loss weight = 0.7, DDIM feature loss weight = 0.2, DDIM classification loss weight = 0.1); among other parameters, the effectiveness screening threshold for generated images was (0.1, 0.9), the number of generated images per category was 50, mixed-precision training was adopted to accelerate the process, and the maximum gradient clipping norm was set to prevent gradient explosion. The experimental results were named in folders by timestamp, and training logs, model weights, evaluation charts, and other files were automatically saved to ensure traceability. All experiments were conducted on the PyTorch framework under the Windows 11 system and accelerated training on RTX4050 16GB.

3. Result

3.1. Result Evaluation and Ablation Experiments

To compare the performance differences among different generator modules, ResNet18 was adopted as the baseline network, and the classification results of the model after adding training images generated by different generator modules to the training set were evaluated. The results are shown in Table. 2 and Fig. 3. First, to evaluate the impact of training set division on model training, the results of dataset division ratios of 4:1 and 1:1 were compared. It can be seen from Table. 2 that after using the 1:1 dataset division, all evaluation indicators decreased significantly: the accuracy dropped by 12.69%, and the F1-score dropped by 13.21%. This indicates that the insufficient sample size of the training set has a significant negative impact on the fitting ability of the model.

Table 2 Ablation Study and Comparison Results of Different Attention Modules (%)

Model	Iteration	Accuracy	Precision	Recall	F1-score
ResNet18(Dataset 4:1)	-	83.72	83.77	83.48	83.58
ResNet18(Dataset 1:1)	-	71.03	71.10	70.27	70.37
ResNet18_LightGAN	1	76.64	76.73	76.07	76.23
	5	78.50	78.94	77.80	78.01
	10	**82.24**	**82.52**	**81.72**	**81.93**
	15	77.57	78.12	76.78	76.99
ResNet18_DDIM	1	78.50	81.10	77.16	77.37
	5	80.37	81.83	79.36	79.66
	10	**83.18**	**85.23**	**82.11**	**82.50**
	15	80.37	80.90	79.68	79.93
ResNet18_Hybrid	1	79.44	79.52	78.98	79.13
	5	**86.92**	**86.95**	**86.66**	**86.78**
	10	83.18	83.35	82.74	82.93
	15	84.11	84.45	83.60	83.84

Table. 2 also compares the classification performance of ResNet18 after adding the LightGAN generation module, DDIM generation module, and Hybrid generation module respectively. Among them, the LightGAN generation module is a lightweight modified version of the traditional generative adversarial network, the DDIM generation module is a lightweight variant of the diffusion model, and the Hybrid generation module is a hybrid generative module that combines the advantages of the two. As shown in Table. 2, the addition of generator modules has a positive effect on improving the classification performance of the model.

When only the LightGAN generation module was used, the overall evaluation indicators of the model, such as accuracy and F1-score, were relatively low and decreased significantly in the later stage. The reason is that LightGAN relies on the adversarial training between the generator and the discriminator: although it can learn the detailed and real features of medical images, it is prone to mode collapse. After excessive iteration, the generated samples are highly homogeneous, which cannot continuously provide diverse and effective supplementary data, resulting in limited improvement in the generalization ability of the model.

When only the DDIM generation module was used, the model achieved outstanding precision (reaching 85.23% at 10 epochs), but the overall accuracy and recall were lower than those of the Hybrid generation module. Although DDIM ensures the diversity and stability of generated samples through gradual denoising and avoids mode collapse, the generated images are slightly inferior to LightGAN in the fit of medical feature details and cannot accurately capture the core features of categories, making it difficult to achieve a breakthrough in comprehensive performance.

In contrast, after introducing the Hybrid generation module, all evaluation indicators of the model were significantly improved. At 5 iterations, the accuracy reached 86.92%, which was 4.68% higher than the peak of LightGAN and 3.74% higher than the peak of DDIM; the F1-score and recall also reached 86.78% and 86.66%, respectively. The core reason is that the Hybrid generation module integrates the advantages of LightGAN in accurately fitting the detailed features of medical images and DDIM in sample diversity and training stability. By dynamically adjusting the weights of the two, it not only avoids the shortcomings of a single module but also fully covers the feature complexity of medical data, providing augmented data with both realism and diversity for model training, and ultimately achieving a leapfrog improvement in comprehensive performance.

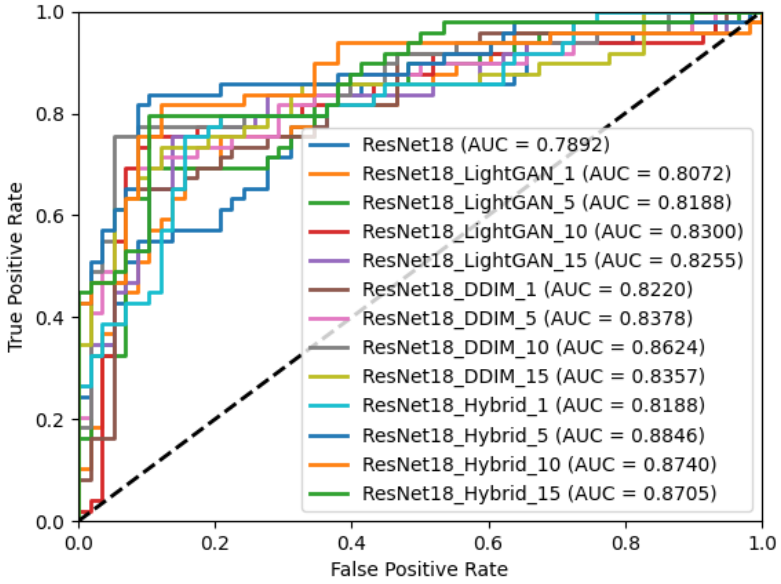

Fig. 3 Multi-epoch ROC Curves of Ablation Study

It can be found from Fig. 3 that the baseline ResNet18 without a generation module (AUC = 0.7892) had the worst discrimination ability, confirming the necessity of data augmentation under small-sample scenarios. Among the single generation modules, the AUC of LightGAN fluctuated significantly with the number of iterations (reaching 0.8300 at 10 epochs and then declining), reflecting the unstable sample quality caused by GAN mode collapse; while the AUC of DDIM improved more smoothly (reaching 0.8624 at 10 epochs), showing the generation stability of the diffusion model. The Hybrid generation module performed the best: ResNet18_Hybrid_5 with 5 iterations achieved an AUC of 0.8846, which was the closest to the perfect classification among all curves, integrating the feature realism of LightGAN and the sample diversity of DDIM. Even when iterated to 15 epochs, its AUC (0.8705) was still higher than the peak of DDIM. All generation modules reached the AUC peak at 5-10 iterations and declined at 15 iterations, further verifying the optimal iteration interval of closed-loop iteration. The high AUC of Hybrid_5 means that it has the strongest ability to balance missed diagnosis and misdiagnosis under different thresholds, making it the most suitable solution for clinical needs in terms of comprehensive performance.

3.2. Comparative Experiments

To explore the generalization of the Hybrid generation module, we also compared its performance on five models: AlexNet[22], VGG16[23], GoogleNet[24], ResNet50[25], and ConvTransResNet. All models adopted the transfer learning strategy, with parameter settings consistent the same, and the number of closed-loop training epochs was set to 5, which was the optimal epoch in the ablation experiments.

As shown in the results of Table. 3, without the Hybrid plug-in, the basic performance of each backbone network was different but generally low: AlexNet had the best basic accuracy of 76.64%, while ResNet50 was the lowest at 71.03%. Moreover, most models had the problem of "imbalance between precision and recall" (e.g., VGG16 had a precision of 76.08% but a recall of only 71.20%, and GoogleNet had a precision of 80.15% but a recall of 71.75%). This reflects that under small-sample medical data, a single backbone network is difficult to balance the needs of "accurate classification" and "comprehensive identification".

Table 3 Comparison Results of 5 Models (%)

Model	Plugin	Accuracy	Precision	Recall	F1-score
AlexNet	-	76.64	76.99	75.91	76.10
	Hybrid	**79.44**	**79.52**	**78.98**	**79.13**
VGG16	-	72.90	76.08	71.20	70.90
	Hybrid	**76.34**	**78.29**	**75.44**	**75.60**
GoogleNet	-	73.83	80.15	71.75	71.08
	Hybrid	**78.50**	**78.66**	**77.96**	**78.13**
ResNet50	-	71.03	73.77	69.32	68.90
	Hybrid	**77.57**	**79.66**	**76.30**	**76.48**
ConvTransResNet	-	73.83	73.67	73.49	73.55
	Hybrid	**81.31**	**82.08**	**80.54**	**80.82**

After introducing the Hybrid plug-in, all evaluation indicators of all models were improved: AlexNet's accuracy increased by 2.80% and F1-score increased by 3.03%; VGG16's accuracy increased by 3.44% and F1-score increased by 4.70%; GoogleNet's accuracy increased by 4.67% and F1-score increased by 7.05%; ResNet50's accuracy increased by 6.54% and F1-score increased by 7.58%; ConvTransResNet achieved the most significant improvement, with an accuracy increase of 7.48% and an F1-score increase of 7.27%. In addition, the gap between its precision (82.08%) and recall (80.54%) was greatly reduced, achieving a balance between "low misdiagnosis" and "low missed diagnosis".

This result confirms the generalization and adaptability of the Hybrid generation module: the "real + diverse" medical image data generated by it can effectively make up for the shortcomings of feature learning of different backbone networks—it has a greater improvement effect on models with weaker basic performance (such as ResNet50), and can give full play to the feature extraction potential of more complex models (such as ConvTransResNet). Finally, the comprehensive performance of various backbone networks in medical image classification tasks is optimized.

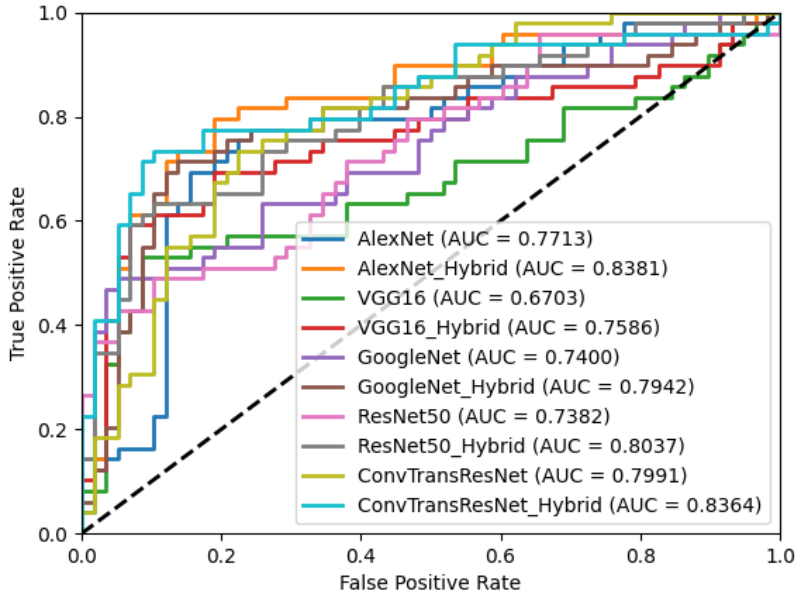

Fig. 4 ROC Curves of Comparison Experiment

It can also be clearly seen from Fig. 4 that the Hybrid generation module improved the ability of different backbone networks to distinguish positive and negative samples: without the Hybrid plug-in, the AUC of each basic model was generally low and the gap was significant, reflecting the limited

discrimination ability of a single backbone network under small-sample medical data; after introducing the Hybrid plug-in, the AUC of all models was significantly improved, and the performance gap between networks was greatly reduced. This not only confirms that the Hybrid generation module can help different backbone networks learn more comprehensive sample features through "real + diverse" augmented data but also reflects that it can make up for the shortcomings of weak-performance networks and release the potential of strong-performance networks, ultimately making the discrimination ability of various models more balanced and improving the possibility of balancing missed diagnosis and misdiagnosis in medical scenarios.

4. Discussion

This study conducted multi-dimensional experiments on the impact of generative data augmentation modules on medical image classification models, interpreting the optimization logic of small-sample tasks from data distribution, generative module characteristics, and clinical adaptability, and discussing the model's innovation, limitations, and implications.

The core advantage of the hybrid module is its dynamic fusion weight mechanism: LightGAN's weight increases for low accuracy to align target features, while DDIM's weight rises for high accuracy to enhance quality. A closed-loop optimization process updates LightGAN via feature + classification loss gradient backpropagation and DDIM via noise + feature + classification loss hyperparameter adjustment, avoiding wrong sample gradient dependence, capturing image feature differences, and using screened generated images to augment the training set and improve generalization.

In terms of data distribution, the 4:1 vs 1:1 dataset comparison (accuracy 83.72% vs 71.03%) confirms inherent clinical data scarcity, caused by high image acquisition and annotation costs, which limits model performance. The generative module addresses this by generating feature-consistent supplementary data, alleviating overfitting and outperforming traditional data augmentation in feature authenticity.

Regarding generative module characteristics, LightGAN generates detailed images but suffers from mode collapse and late performance decline, while DDIM ensures stability and diversity but lacks precision in fitting medical core features. The proposed hybrid module balances realism and diversity via dynamic weight adjustment: on ResNet18 at 5 iterations, it achieved 86.92% accuracy and 0.8846 AUC (the highest), and it significantly improved accuracy, F1-score, and AUC across various backbone networks, narrowing performance gaps and enabling weak networks (e.g., ResNet50) to meet clinical needs.

Clinically, the model alleviates ICG image scarcity, improves MVI classification accuracy, and provides reliable diagnostic support for HCC treatment and prognosis. Academically, it offers a new approach to small-sample medical image classification by integrating generative model advantages. Limitations include single-institution data and focus on HCC MVI classification; future research will expand datasets and explore broader applications.

5. Conclusion

This study systematically verifies the value of generative data augmentation in small-sample medical image classification tasks through multiple sets of experiments, confirming that a dataset conforming to real clinical distribution is the foundation of model performance; the proposed Hybrid hybrid generative module, which integrates the detailed fitting ability of LightGAN and the generation stability of DDIM, not only improves the peak performance of a single backbone network but also has good generalization to different backbone networks, narrowing the performance gap between networks, and the optimal interval of closed-loop iteration is 5-10 rounds, with excessive

iteration leading to performance degradation; importantly, the generative module can improve model accuracy, F1-score and AUC without additional clinical data collection, has a greater improvement effect on weak-performance networks, reduces the dependence on high-end computing resources in medical scenarios, and ultimately provides a feasible technical framework for small-sample medical image classification tasks and experimental support for the development of subsequent medical image auxiliary diagnosis systems.

References

*[1] Bray, F. et al. Global cancer statistics 2022: GLOBOCAN estimates of incidence and mortality worldwide for 36 cancers in 185 countries. CA A Cancer J Clinicians **74**, 229–263 (2024).*

*[2] Cheng, Z. et al. Risk factors and management for early and late intrahepatic recurrence of solitary hepatocellular carcinoma after curative resection. HPB **17**, 422–427 (2015).*

*[3] Raoul, J.-L. & Edeline, J. Systemic treatment of hepatocellular carcinoma: standard of care in China and elsewhere. The Lancet Oncology **21**, 479–481 (2020).*

*[4] Tang, Y. et al. Identification and Validation of a Prognostic Model Based on Three MVI-Related Genes in Hepatocellular Carcinoma. Int. J. Biol. Sci. **18**, (2022).*

*[5] Wang, F. et al. Predicting microvascular invasion in small (≤ 5 cm) hepatocellular carcinomas using radiomics-based peritumoral analysis. Insights Imaging **15**, 90 (2024).*

*[6] Zheng, R. et al. Comparison of non-radiomics imaging features and radiomics models based on contrast-enhanced ultrasound and Gd-EOB-DTPA-enhanced MRI for predicting microvascular invasion in hepatocellular carcinoma within 5 cm. Eur Radiol **33**, 6462–6472 (2023).*

*[7] Xu, C., Jiang, D., Tan, B., Shen, C. & Guo, J. Preoperative diagnosis and prediction of microvascular invasion in hepatocellularcarcinoma by ultrasound elastography. BMC Med Imaging **22**, 88 (2022).*

[8] Wang, G. Prediction of Microvascular Invasion of Hepatocellular Carcinoma Based on Preoperative Diffusion-Weighted MR Using Deep Learning. Academic Radiology (2020).

*[9] Zhou, W. Prediction of Microvascular Invasion of Hepatocellular Carcinoma Based on Contrast-Enhanced MR and 3D Convolutional Neural Networks. Frontiers in Oncology **11**, (2021).*

[10] Cen Y-Y, et al. Computed tomography-based deep learning and multi-instance learning for predicting microvascular invasion and prognosis in hepatocellular carcinoma. World Journal of Gastroenterology, 2025, 31(30): 109186. doi:10.3748/wjg.v31.i30.109186.

[11] Cao, L. et al. MVI-TR: A Transformer-Based Deep Learning Model with Contrast-Enhanced CT for Preoperative Prediction of Microvascular Invasion in Hepatocellular Carcinoma. https://doi.org/10.3390/cancers15051538 (2023) doi:10.3390/cancers15051538.

[12] Zhou, M. A rock joint roughness coefficient determination method incorporating the SE-Net attention mechanism and CNN. Environmental Earth Sciences https://doi.org/10.1007/s12665-025-12414-x (2025) doi:10.1007/s12665-025-12414-x.

[13] Wang, Q. et al. ECA-Net: Efficient Channel Attention for Deep Convolutional Neural Networks. https://doi.org/10.1109/CVPR42600.2020.01155 doi:10.1109/CVPR42600.2020.01155.

[14] Anshelevich, D. Synthetic tabular data generation using a VAE-GAN architecture. https://doi.org/10.1016/j.knosys.2025.113997 (2025) doi:10.1016/j.knosys.2025.113997.

[15] Radford, A., Metz, L. & Chintala, S. Unsupervised Representation Learning with Deep Convolutional Generative Adversarial Networks. Preprint at http://arxiv.org/abs/1511.06434 (2016).

[16] Isola, P., Zhu, J.-Y., Zhou, T. & Efros, A. A. Image-to-Image Translation with Conditional Adversarial Networks. https://doi.org/10.1109/CVPR.2017.632 doi:10.1109/CVPR.2017.632.

[17] Ho, J., Jain, A. & Abbeel, P. Denoising Diffusion Probabilistic Models. Preprint at http://arxiv.org/abs/2006.11239 (2020).

[18] Song, J., Meng, C. & Ermon, S. DENOISING DIFFUSION IMPLICIT MODELS. https://doi.org/10.48550/arXiv.2010.02502 (2021).

*[19] Kose, E. et al. A comparison of indocyanine green fluorescence and laparoscopic ultrasound for detection of liver tumors. HPB **22**, 764–769 (2020).*

[20] Szostek, A. K., Grimes, J. A. & Wallace, M. L. Average fluorescence intensity is affected by distance and penetration depth with near-infrared fluorescence imaging with indocyanine green in an ex vivo model. https://doi.org/10.2460/ajvr.25.01.0001 doi:10.2460/ajvr.25.01.0001.

[21] Cai, X. Does Using Indocyanine Green Fluorescence Imaging for Tumors Help in Determining the Safe Surgical Margin in Real-Time Navigation of Laparoscopic Hepatectomy? A Retrospective Study. https://doi.org/10.1245/s10434-

022-12893-3 doi:10.1245/s10434-022-12893-3.
*[22] Krizhevsky, A., Sutskever, I. & Hinton, G. E. ImageNet classification with deep convolutional neural networks. Commun. ACM **60**, 84–90 (2017).*
[23] Simonyan, K. & Zisserman, A. VERY DEEP CONVOLUTIONAL NETWORKS FOR LARGE-SCALE IMAGE RECOGNITION. https://doi.org/10.48550/arXiv.1409.1556 (2015).
[24] Szegedy, C. et al. Going Deeper with Convolutions. https://doi.org/10.1109/CVPR.2015.7298594 doi:10.1109/CVPR.2015.7298594.
[25] He, K., Zhang, X., Ren, S. & Sun, J. Deep Residual Learning for Image Recognition. Preprint at http://arxiv.org/abs/1512.03385 (2015).