

A Comparative Study of Multi-class Classification Based on Imbalanced Data

Rojan Zaki Abdulkareem¹, Adnan Mohsin Abdulazeez^{2,*}

¹*Akre University for Applied Science, Technical College of Informatics – Akre
Department of Information Technology – Kurdistan Region – Iraq*

²*Duhok Polytechnic University – Kurdistan Region – Iraq*

Abstract Class imbalance presents a significant challenge in creating reliable and precise medical diagnostic models, especially in multi-class classification contexts where rare yet clinically important cases are insufficiently represented. This work addresses the imbalance problem across three different medical datasets: HAM10000, Skin Cancer ISIC, and Non-Alcoholic Fatty Liver Disease (NAFLD) by presenting an extensive deep learning framework using Cycle-Consistent GANs (CycleGAN) for data balancing, integrating advanced data augmentation methods, and applying Focal Loss to enhance training. The suggested architecture utilizes EfficientNet-B3 for image classification and a custom-built Multi-Layer Perceptron (MLP) for evaluating tabular clinical data. The CycleGAN model is employed to create realistic images of minority classes and to replicate oversampling in tabular domains, thus generating balanced and semantically varied datasets. To enhance generalization, we implement real-time augmentation techniques, which encompass image data augmentation via flipping, rotation, and color jittering, alongside normalization strategies for tabular features. This study presents a unified deep learning pipeline that implements real CycleGAN-based oversampling for both image and tabular medical datasets, distinguishing it from previous research. The amalgamation of CycleGAN with Focal Loss and EfficientNetB3 yields improved efficacy in minority-class detection, setting a novel benchmark for imbalanced multi-class medical classification. The performance evaluation was performed using stratified 5-fold cross-validation, employing measures like macro F1-score, balanced accuracy, and ROC-AUC. The proposed method attained superior results across all datasets, with an efficient peak accuracy of 99.33%, a macro F1-score of 96.85%, and a ROC-AUC of 0.9852 on the HAM10000 dataset. The comparative analysis with previous studies illustrates the supremacy of our pipeline in overall accuracy and minority-class sensitivity. The comparison analysis with past studies shows the superiority of our pipeline in generic accuracy and minority-class sensitivity.

Keywords Class Imbalance, CycleGAN, EfficientNet-B3, Focal Loss, Medical Image Classification

DOI: 10.19139/soic-2310-5070-2731

1. Introduction

Medical data imbalance presents a significant challenge in the development of intelligent diagnostic systems, especially in dermatology and hepatology. Conditions such as skin cancer and NAFLD frequently demonstrate imbalanced class distributions, with prevalent cases significantly outnumbering rare yet clinically significant instances. The existing imbalance impairs the generalization ability of machine learning models and diminishes their sensitivity to minority classes, resulting in inaccurate or biased predictions in critical healthcare contexts [1][2]. Skin cancer is among the most common types of cancer worldwide, imposing a significant burden on healthcare systems [3]. Dermoscopic datasets, including HAM10000 and ISIC, have expedited progress in automated analysis through deep learning (DL). However, these datasets exhibit significant imbalance, with benign lesions, such as melanocytic nevi (nv), vastly outnumbering malignant types, including vascular lesions (vasc)

*Correspondence to: Rojan Zaki (Email: rojan.zaki@auas.edu.krd). Department of Information Technology, Technical College of Informatics - Akre, Akre University for Applied Sciences, Duhok, Kurdistan Region, Iraq.

and dermatofibromas (df) [4]. The NAFLD dataset, which emphasizes liver fat staging, includes a greater number of early-stage samples compared to advanced cases. These distributions lead to diagnostic models that exhibit bias towards dominant classes and lack the necessary sensitivity for identifying rare but severe conditions [5][6]. Traditional approaches to addressing class imbalance encompass oversampling, undersampling, and the generation of synthetic data through methods such as SMOTE (Synthetic Minority Oversampling Technique) and ADASYN [7][8]. Nonetheless, these methods may fail to maintain semantic integrity or adequately represent complex feature distributions in high-dimensional data [9]. Generative Adversarial Networks (GANs) represent a class of deep learning models that were introduced recently [10]. Comprises two competing neural networks: a generator and a discriminator. Figure 1 illustrates summer. The generator produces synthetic data samples, whereas the discriminator assesses these samples against real data, offering evaluation to enhance the generator's performance. In this adversarial framework, the generator incrementally acquires the ability to generate data that increasingly mirrors the original distribution [11]. This framework has demonstrated significant efficacy in tasks including image synthesis, data augmentation, and domain translation, facilitating the generation of high-dimensional, complex, and realistic data representations without the necessity for explicit labeling or paired datasets [12].




Figure 1. Basic architecture of a GAN.

To overcome these obstacles, we suggest an integrated system that enhances multi-class classification in unbalanced medical datasets by combining state-of-the-art DL classifiers, data augmentation, and CycleGAN-based oversampling. Three different datasets are used to validate our pipeline: the NAFLD tabular dataset, the HAM10000 dataset, and the ISIC skin cancer image dataset. Each dataset suffers from class imbalance despite modality variations, which makes them perfect candidates to evaluate the adaptability of our methodology. The EfficientNet-B3 architecture is used in the suggested system for image-based classification because it strikes a balance between computational efficiency and accuracy. A bespoke MLP is used to represent nonlinear dependencies in structured data for the tabular NAFLD dataset. We use Focal Loss, which dynamically modifies loss contributions according to prediction difficulty, to further improve classification performance on minority classes by concentrating the model's learning on examples that are more difficult to categorize. The following is a summary of our main contributions:

1. We use stratified K-fold cross-validation and a common assessment technique based on macro F1-score, balanced accuracy, and ROC-AUC to evaluate our approach on three datasets: HAM10000, ISIC, and NAFLD. This ensures generalizability across modalities.

2. We suggest a unified architecture that makes use of real-time data augmentation techniques and synthetic samples produced by CycleGAN to lessen multi-class imbalance in medical data. This maintains the semantic characteristics of minority samples while guaranteeing balanced class distributions.

3. We include Focal Loss into both image and tabular classification pipelines to increase model robustness and diagnostic precision, enabling the models to concentrate more on difficult and underrepresented classes.

4. Emphasizing scalability, we use the same training pipeline structure across tabular and image-based datasets with only minor adjustments, demonstrating the adaptability and reusability of our approach in many clinical settings.

This paper’s remaining sections are arranged as follows: Section 2 examines relevant research in medical data classification and GAN-based data augmentation. The preprocessing pipeline, balancing mechanism, and dataset features are described in Section 3. The training processes and experimental setup are described in Section 4. The results and a detailed discussion of the findings are presented in Section 5. Section 6 wraps up the paper and provides guidance for further research.

2. Related Works

Deep learning has considerably progressed medical diagnosis by facilitating automatic and precise virus identification across several modalities, such as dermoscopic pictures and electronic health information. The efficacy of these models is significantly contingent upon the quality and composure of the training datasets. A number of academics have investigated methods to alleviate the detrimental impacts of data imbalance, augment generalization, and increase prediction accuracy for minority classes.

2.1. Class Imbalance in Medical Datasets

Class imbalance is a continual difficulty in medical classification tasks, as datasets generally display an unequal representation of normal and abnormal samples [13][14]. In dermoscopic datasets such as HAM10000 and ISIC, the predominant pictures are classified as benign, specifically (nv), while malignant lesions are inadequately represented [15]. Clinical datasets like NAFLD exhibit a significant proportion of early-stage or mild cases, hence constraining the model’s capacity to acquire distinguishing traits for advanced stages [16]. Conventional methods for tackling class imbalance encompass random oversampling, undersampling, and algorithmic strategies like SMOTE and ADASYN [17]. Although these methods provide slight speed improvements, they frequently result in overfitting or the introduction of noise due to synthetic data production reliant only on feature interpolation, especially in high-dimensional image data [18].

2.2. GAN-based Oversampling Techniques

Recent studies have utilized GANs to generate class-specific, high-fidelity pictures, thereby addressing the constraints of conventional oversampling methods [19][20]. CycleGANs have gained recognition for their capacity to learn inter-class mappings without the necessity of paired training data. Exhibited the efficacy of CycleGAN in producing authentic lesion images to equilibrate dermoscopy datasets [21]. Likewise, GANs to radiological and histological pictures, markedly enhanced minority class recall [22]. Figure 2 depicts the procedure of GAN-based oversampling utilizing CycleGAN. We employ a CycleGAN-based balancing technique for both HAM10000 and NAFLD, training class-to-class generators to replicate underrepresented samples. In contrast to previous

techniques that depend exclusively on image replication or SMOTE, our methodology produces structurally varied and semantically precise samples, leading to balanced and expressive training distributions [23].




Figure 2. Gan-based Overview sampling using CycleGan.

2.3. Deep Learning Models for Skin Lesion and Clinical Data Classification

Convolutional Neural Networks (CNNs) are extensively utilized in dermatological image analysis, with designs including ResNet, VGG, and DenseNet demonstrating robust performance on benchmark datasets. Recently, EfficientNet has emerged as a premier architecture that optimally balances accuracy and computational expense [24]. Introduced a compound scaling algorithm that equally adjusts depth, width, and resolution, allowing EfficientNet to surpass previous CNNs in multiple medical imaging tasks [25]. For tabular medical data, (MLPs) continue to be a favored option because of their straightforwardness and versatility. Research by [26], indicates that when integrated with rigorous preprocessing and class balancing, MLPs may proficiently model intricate feature interactions in clinical datasets, including NAFLD. EfficientNet-B3 is the main image classifier in our method since it is so very effective on small-to-medium medical datasets. We use normalized clinical features [27]. and a custom three-layer MLP to classify steatosis phases in the NAFLD dataset.

2.4. Loss Functions for Imbalanced Classification

Often struggling with imbalanced datasets, standard loss functions, including cross-entropy, allocate uniform weight to all samples, independent of class frequency. Focal Loss was developed to reduce the impact of accurately identified samples and to focus learning on difficult, misclassified instances [28]. Focal Loss has been successfully employed in medical imaging applications such as skin lesion detection, pneumonia classification, and diabetic retinopathy grading. This study integrates Focal Loss into the classification pipelines for both images and tabular data [29]. The loss function enhances convergence by prioritizing errors in the minority class, thereby improving macro-level evaluation metrics, such as F1-score and balanced accuracy [30].

3. Research Methods

The methods and tools applied to reduce data imbalance and enhance classification accuracy over several medical datasets are described in this part. Figure 3 shows, within a multi-modal deep learning architecture, the combination of image-derived spatial features with clinical data. This study utilized three publicly available datasets: the

HAM10000 dataset, which contains 10,000 training photos, the ISIC Skin Cancer dataset, and the NAFLD tabular dataset. The datasets include several imaging and clinical diagnostic modalities, including dermatoscopic images of skin lesions and records of liver conditions. Resampling, normalization, augmentation, and patch-based processing were employed to guarantee consistency and robustness across tests. Resampling, normalization, augmentation, and patch-based processing were employed to guarantee consistency and robustness across tests. The executed measures normalize input data and alleviate class imbalance effects, hence improving model generalization during both the training and evaluation stages.




Figure 3. Complete Workflow of the Proposed Deep Learning Pipeline for Class-Imbalanced Medical Image Classification.

3.1. Datasets

In medical datasets, class imbalance poses a major problem that usually results in biased models and insufficient generalization for underrepresented classes. We used synthetic oversampling techniques—including statistical duplication and CycleGAN-style approaches—along with augmentation strategies unique to the modality of each dataset to solve this problem. The balancing techniques followed for the NAFLD, Skin Cancer ISIC, and HAM10000 datasets are described in this paper.

3.1.1. HAM10000 Dataset

The HAM10000 dataset has considerable imbalance, marked by a prevalence of (nv) samples compared to the rare classes, including (vasc) and (df). We employed a two-phase process that combines CycleGAN-based image generation with data augmentation. Initially, CycleGAN was trained to produce synthetic dermatoscopic images by establishing mappings from the dominant class (nv) to each minority class. The procedure was repeatedly executed to standardize the sample size across all seven lesion types. Figure 4 presents the class distribution prior to and following the balancing operation. During the second stage, real-time data augmentation was implemented to enhance diversity and robustness. The augmentation pipeline comprised horizontal and vertical flipping, random rotations of $\pm 20^\circ$, color jittering with variations of $\pm 20\%$ in brightness, contrast, saturation, and hue, as well as

resizing to dimensions of 224×224 pixels. The transformations enhanced intra-class variability and facilitated effective generalization of the EfficientNet-B3 model. Figure 5 illustrates an instance of augmentation utilized on a dermoscopic lesion.




Figure 4. Class Distribution BEFORE / AFTER balancing.




Figure 5. Data augmentation by image rotation and flipping.

3.1.2. Skin Cancer ISIC Dataset The ISIC Skin Cancer dataset exhibited significant class imbalance across all nine diagnostic categories. To efficiently replicate CycleGAN-style oversampling, minority class samples were duplicated until each class comprised a minimum of 800 images, as demonstrated in Figure 6. This approach, although not utilizing generative modeling, successfully reduced class bias during the training phase. Augmentations were then applied to all training samples to simulate the real-world variability experienced in a clinical setting. incorporated were arbitrary horizontal flips, 20° rotations, moderate adjustments to brightness and contrast, and resizing to 224×224 pixels. In contrast to training, validation, and testing employed standardized transformations to ensure consistency. The implemented balancing strategy yielded a well-distributed and diverse dataset, enhancing the performance of the EfficientNet-B3 classifier for all lesion types.




Figure 6. Class Distribution BEFORE/AFTER Balancing.




Figure 7. Class Distribution BEFORE/AFTER Balancing.

3.1.3. NAFLD Tabular Dataset Clinical tabular features and labels indicating four phases of hepatic steatosis define the NAFLD dataset. Early-stage labels dominated the dataset, and the class distribution showed notable skewness. Using statistical resampling—where minority class samples were oversampled via bootstrapped duplication to match the majority class size—we sought balance. Figure 7 shows the class distribution both before and after balancing. Following balancing, all numerical features underwent normalization through standard scaling. Exclusion of categorical attributes occurred, and the resultant matrix was utilized to train a fully connected neural network MLP. While focal loss was used to solve the remaining imbalance, stratified 5-fold cross-validation was used to guarantee a consistent performance estimation. Performance was assessed using ROC-AUC, balanced accuracy, and macro F1-score.

3.2. Data Preprocessing

To guarantee uniformity and appropriateness for deep learning, all input data underwent a stringent pre-processing pipeline tailored to each dataset type: image-based (HAM10000 and ISIC) and tabular (NAFLD).

3.2.1. Image Preprocessing For both **HAM10000** and **Skin Cancer ISIC** datasets, all **dermoscopic** images were first resized to a uniform resolution of 224×224 pixels using bilinear interpolation. Let $I(x, y)$ be the original image of size $H \times W$. The rescaled image $\Gamma(x', y')$ is computed as:

$$\Gamma(x', y') = I\left(x' \cdot \frac{H}{H'}, y' \cdot \frac{W}{W'}\right), \quad x' \in [0, H'], y' \in [0, W']$$

Where H' and W' are the standardized height and width (i.e., 224 pixels). After resizing, **pixel intensity normalization** was applied using **ImageNet statistics**. The normalized image $I_{norm}(x, y)$ is computed using:

$$I_{norm}(x, y) = \frac{I(x, y) - \mu}{\sigma}$$

where μ and σ are the per-channel mean and standard deviation of ImageNet. This normalization ensures consistency across datasets and accelerates convergence.

3.2.2. Augmentation Strategy The training pipeline included several real-time augmentations to increase model robustness and avoid overfitting:

- **Random Horizontal Flip:**

$$I_{flipped}(x, y) = I(W - x - 1, y)$$

- **Random Rotation:** Given a rotation angle θ , the transformed coordinates are:

$$\begin{bmatrix} x' \\ y' \end{bmatrix} = \begin{bmatrix} \cos \theta & -\sin \theta \\ \sin \theta & \cos \theta \end{bmatrix} \begin{bmatrix} x - x_c \\ y - y_c \end{bmatrix} + \begin{bmatrix} x_c \\ y_c \end{bmatrix}$$

where (x_c, y_c) is the centre of the image.

- **Color Jitter:** was applied to randomly adjust brightness B , contrast C , and saturation S within $\pm 20\%$ range:

$$J_{jitter} = C \cdot (I - \mu) + \mu + B + S$$

These augmentations simulate variability in lighting and orientation commonly encountered in real-world clinical imaging, enabling the model to generalize effectively across different domains.

3.2.3. Tabular Data Preprocessing

Missing values and non-numerical columns were deleted from the NAFLD data. Z-score standardizing helped all features be normalized:

$$Z_i = \frac{x_i - \bar{x}}{\sigma}$$

where $\bar{x}$ and σ are the mean and standard deviation of the feature. The target variable (*Steatosis stage*) was encoded using:

$$y_i = \text{LabelEncoder}(\text{stage}_i)$$

This unified preprocessing framework facilitated direct comparison of models across various modalities—dermoscopic images and clinical features—by employing a consistent classification architecture and validation protocol.

3.3. Proposed Model

This study aims to create a comprehensive deep learning framework for the classification of skin lesions and tabular health data across diverse datasets with different modalities. We established a standardized training pipeline utilizing the EfficientNet-B3 architecture to guarantee consistency and reproducibility. We implemented it uniformly across three separate datasets: HAM10000 (image and metadata), Skin Cancer ISIC (images only), and the NAFLD tabular dataset. The proposed system consists of four primary components: data preprocessing, class balancing, deep model training utilizing stratified k-fold cross-validation, and performance evaluation through advanced metrics such as the macro F1-score, balanced accuracy, and ROC-AUC.

3.3.1. Interpretability of Synthetic Sample Generation




Figure X: Comparison of Real vs. Synthetic Lesion Samples Generated by CycleGAN

Figure 8. Comparison of Real vs. Synthetic Lesion Samples Generated by CycleGAN.

Qualitative visualizations illustrating the semantic integrity of CycleGAN-generated samples are shown in Figure 8. Synthetic dermoscopic images for the HAM10000 and ISIC datasets preserve lesion boundary structures, pigmentation patterns, and morphological characteristics specific to each class. Synthetic samples in the NAFLD tabular dataset maintain the correlations among clinical features, including ALT and lipid levels, thereby preserving authenticity. CycleGAN achieves this by imposing a cycle-consistency loss, which ensures that a sample translated

from the source domain to the target domain and then returned preserves its original structure. This characteristic ensures that the resulting minority-class samples are not random but based on significant changes obtained from authentic medical data.

3.3.2. Justification for CycleGAN and Computational Considerations CycleGAN was chosen as the best generative model for oversampling because it can translate images without needing exact matches between different types of lesions or medical cases. This capability is especially useful in medical imaging, where such direct matches are often missing. In contrast to WGAN, which maintains training stability but lacks consistency in image generation and style transfer, and StyleGAN, which requires extensive datasets and paired examples for optimal performance, CycleGAN offers greater flexibility and produces images that more accurately reflect real-world semantics. CycleGAN successfully generates realistic images and tables that are beneficial for diagnosis, accurately preserving the shape, color, and structure of lesions. Conversely, WGAN generated images of lower clarity, while StyleGAN required significantly more resources and produced fewer diverse outcomes due to its dependence on interpolation within the latent space. The CycleGAN training utilized an NVIDIA RTX 4070 GPU (8 GB VRAM), using roughly 5.2 GPU hours per dataset, with memory consumption varying between 6.5 and 7.1 GB, depending on image resolution. CycleGAN, although more resource-demanding than conventional GANs or SMOTE-based resampling techniques, warrants the incurred expenses owing to its ability to generate high-fidelity samples across unpaired domains. To enhance scalability and facilitate potential implementation, we recommend investigating lightweight alternatives such as FastCUT, TinyGAN, or MobileGAN for the next versions. These architectures minimize model size and inference time while maintaining semantic integrity, rendering them appropriate for resource-limited clinical settings.

3.3.3. Balancing Strategy In medical datasets, class imbalance poses a major obstacle that causes biased learning whereby the model underperforms on minority classes and favors overrepresented classes. Using bootstrapped resampling for tabular data and CycleGAN-based oversampling for image data, our approach generates class-balanced training sets. We employed a CycleGAN-based balancing simulation for the HAM10000 and Skin Cancer ISIC datasets to enhance the representation of underrepresented classes. This was achieved by replicating and transforming samples until each class achieved parity with the majority class. The number of samples per class was standardized to a fixed threshold (e.g., 5000–7000), as illustrated in the class distribution figures prior to and following balancing (refer to Figures 4 and 6). Let $N_{\max}$ represent the number of samples in the majority class, and N_i the number in class i . For each class where $N_i < N_{\max}$, synthetic samples S_i were generated using:

$$S_i = \left\lceil \frac{N_{\max} - N_i}{N_i} \right\rceil \cdot N_i$$

This approach assured that every class engaged in model training equally, so improving fairness and reducing generalization bias. Synthetic oversampling of the NAFLD tabular dataset was executed by random resampling with replacement, specifically aimed at augmenting the minority class count. This guaranteed uniform class representation throughout all five folds in cross-validation. The balancing technique facilitated more consistent training, accelerated convergence, and enhanced recall for minority classes, as indicated by the evaluation measures.

3.3.4. Feature Transformation and Normalization To maintain consistency and comparability among picture and tabular features, all input data were standardized through suitable transformation methods specific to the data modality. Preprocessing for image-based datasets (HAM10000 and ISIC) was conducted at the pixel level to standardize brightness, contrast, and scale. Each dermoscopic image was scaled to a standardized resolution of

224×224 pixels to ensure consistent input dimensions across the deep learning model process. Following resizing, intensity normalization was performed using min-max scaling, which adjusts pixel values into a fixed range $[0, 1]$ as indicated by Equation (1):

$$I_{\text{norm}}(x, y) = \frac{I(x, y) - I_{\min}}{I_{\max} - I_{\min}} \quad (1)$$

In this case, $I(x, y)$ represents the pixel intensity at position (x, y) , and $I_{\min}$, $I_{\max}$ are the minimum and maximum intensities within the image. This normalization guaranteed uniform brightness and eradicated the effects of ambient light or device-specific discrepancies. Furthermore, each image was altered utilizing conventional ImageNet normalization parameters, incorporating the subsequent channel-wise mean and standard deviation:

$$\text{Mean} = [0.485, 0.456, 0.406], \quad \text{Std} = [0.229, 0.224, 0.225]$$

In the NAFLD dataset, non-numeric properties were initially removed or encoded, and the continuous features were subsequently scaled using z-score normalization (zero mean and unit variance), as stated in equation (2):

$$z_i = \frac{x_i - \mu}{\sigma} \quad (2)$$

where x_i is the raw feature value, μ is the mean, and σ signifies the standard deviation of the feature column. This modification promotes convergence in neural networks by centering the distribution and ensuring all characteristics contribute uniformly during optimization.

3.3.5. Model Architecture This study utilized the EfficientNet-B3 convolutional neural network architecture as the foundation for all classification tasks across three datasets: HAM10000, Skin Cancer ISIC, and NAFLD. EfficientNet-B3, which is used for its optimal equilibrium between computational efficiency and accuracy, employs a compound scaling technique to equally adjust depth, width, and resolution with fixed coefficients. The model was initialized using pre-trained ImageNet weights to leverage transfer learning and expedite convergence. The network received inputs comprising photos downsized to 224×224 pixels for both the HAM10000 and ISIC datasets, together with normalized tabular vectors for the NAFLD dataset.

The EfficientNet-B3 model processed image-based datasets through a sequence of convolutional layers with Swish activation functions, followed by global average pooling. A dropout layer was implemented after the final feature extraction layer to alleviate overfitting by randomly deactivating a portion of the neurons during training. We modified the final dense layer to correspond with the number of output classes for each dataset in the classification process.

For the NAFLD tabular dataset, a dedicated MLP (Multi-Layer Perceptron) was utilized to process structured clinical data. The MLP has three fully connected layers utilizing dropout and Leaky ReLU activation functions, adeptly simulating intricate nonlinear relationships in the normalized tabular features. This architecture facilitates the strong classification of hepatic steatosis phases, even in imbalanced situations.

We utilized Focal Loss rather than regular cross-entropy loss to mitigate class imbalance by diminishing the weight of easy examples and directing the model's attention towards more challenging, misclassified instances. The loss function is formally defined as:

$$FL(p_t) = -\alpha(1 - p_t)^\gamma \log(p_t)$$

where p_t is the model's predicted probability for the correct class, α is a balancing factor, and γ is the focusing parameter set to 2.0 in our implementation. This approach enabled the model to acquire more distinctive features while being resilient to imbalanced class distributions and overfitting. The comprehensive architecture employed in our investigation is encapsulated in Figure 9.

- **AUC-ROC:** Evaluates the area under the ROC curve to show how well the model differentiates between classes.
- **Confusion Matrix:** Offers a clear summary of classification outcomes, emphasizing performance across different classes.

All models underwent training and evaluation by 5-fold Stratified Cross-Validation to ensure reliability and mitigate bias from random partitions. Throughout each fold, we monitored training and validation accuracy curves to assess convergence and identify overfitting.

3.3.7. Parameter Selection in Focal Loss

Focal Loss comprises two adjustable hyperparameters— α and γ —that regulate class balance and model emphasis, respectively. The parameter $\alpha \in [0, 1]$ balances the contributions of positive and negative classes to mitigate class imbalance, while $\gamma > 0$ modifies the emphasis on difficult instances by reducing the weight of correctly classified examples. The study empirically established that $\alpha = 0.25$ and $\gamma = 2.0$ produced optimal performance across all datasets. Reduced γ values resulted in underfitting, especially within minority classes, whereas elevated values induced training instability. Table 1 presents a summary of an ablation study performed on different parameters.

Table 1. Ablation study on Focal Loss parameters

α (alpha)	γ (gamma)	Macro F1	Balanced Acc	ROC-AUC
0.25	1.0	0.91	0.89	0.9752
0.25	2.0	0.9685	0.9730	0.9852
0.5	2.0	0.94	0.92	0.9801

4. Experimental Setup

A methodical experimental framework was developed and applied to guarantee strict evaluation of the suggested approach. All experiments were carried out on a local machine: an ASUS TUF Gaming A17 FA707XI laptop equipped with an AMD Ryzen 9 7940HS processor (8 cores, 16 threads, 4.00 GHz), 16 GB RAM, and an NVIDIA GeForce RTX 4070 GPU with 8 GB GDDR6 memory. The operating system used was Windows 11 Home, and all deep learning tasks were executed using Python 3.11 with PyTorch as the primary framework.

4.1. Training Configuration

Each fold underwent training for 20 epochs with a batch size of 32, and weights were preserved after each fold to provide checkpoint-based assessment and continuation. The Adam optimizer was used starting with a 0.001 learning rate.

4.2. Hyperparameter Tuning

Hyperparameters, including learning rate, batch size, dropout rate, and activation functions, were optimized using a manual grid search. Table 2 captures the search range and the ideal values: Comparing the Leaky ReLU function

Table 2. Hyperparameter Tuning

Parameter	Values Tested	Optimal Value
Learning Rate	1×10^{-3} , 1×10^{-4} , 1×10^{-5}	1×10^{-4}
Batch Size	16, 32, 64	32
Activation	ReLU, GELU	Leaky ReLU
Loss Function	CrossEntropy, Focal Loss ($\gamma = 2.0$, $\alpha = 0.25$)	Focal Loss

to ReLU and GELU reveals better convergence. In managing class imbalance, Focal Loss ($\gamma = 2.0$, $\alpha = 0.25$) exceeded conventional CrossEntropy. A dropout rate of 0.1 avoided overfitting without sacrificing model capacity. Choosing a batch size of 32 allowed one to balance training stability with memory efficiency.

4.3. Balancing and Augmentation Strategy

CycleGAN-based synthetic augmentation was used on the HAM10000 and ISIC skin cancer datasets to provide realistic samples for underrepresented classes, hence addressing class imbalance. Real-time data augmentation—using random horizontal and vertical flips, rotation ($\pm 20^\circ$), and color jittering (brightness, contrast)—then improved model robustness and variability in feature space. A CycleGAN-style resampling simulation was utilized for the NAFLD tabular dataset, oversampling minority classes using bootstrapped duplication to align with the majority class count. Additionally, Min-Max normalization was employed on all numeric features to standardize their ranges, while categorical variables were omitted to streamline the feature space and ensure training uniformity.

5. Results and Discussion

This section delineates the empirical results of the proposed deep learning architecture across three separate datasets: HAM10000, Skin Cancer ISIC, and NAFLD. Each model underwent evaluation by stratified five-fold cross-validation, with performance measured by important metrics such as accuracy, macro F1-score, balanced accuracy, ROC-AUC, and confusion matrices. All assessments were performed after the implementation of CycleGAN-based data balance and robust data augmentation techniques.

5.1. HAM10000 Results

The EfficientNet-B3 architecture, trained on the balanced and augmented HAM10000 dataset, attained outstanding classification performance. The model exhibited a high level of consistency across all folds, with an average macro

F1-score of 0.988 and a balanced accuracy of 0.988. The confusion matrix for Fold 1 Figure 10 demonstrates accurate predictions across all seven skin lesion categories, exhibiting negligible confusion even across superficially analogous forms such as melanocytic nevi (nv) and melanoma (mel). Table 3 provides a comprehensive comparison of model performance, illustrating the efficacy of CNN models across all measures when trained with GAN-augmented data.

Table 3. Performance of CNN Models on HAM10000 (ORG vs. GAN)

Dataset	Model	Data	F1-Score	ROC-AUC	Balanced Acc	Training Acc	Validation Acc
HAM10000	EfficientNetB3	ORG	79.28	98.13	78.44	95.90	87.62
		GAN	99.33	99.98	99.30	99.34	99.37
	ResNet50	ORG	46.39	91.38	44.07	73.51	74.94
		GAN	98.39	99.93	98.27	98.06	98.36
	VGG19	ORG	68.07	95.17	66.26	83.93	80.85
		GAN	97.13	99.80	97.09	95.05	96.92
	MobileNet	ORG	77.15	97.52	75.76	92.18	86.02
		GAN	98.83	99.94	98.85	97.41	97.90




Figure 10. (a) The accuracy Training/validation (b) Confusion Matrix for HAM10000 Fold 1.

5.2. ISIC Skin Cancer Results

The model exhibited strong classification capabilities on the ISIC dataset. The confusion matrix in Fold 3 Figure 11 demonstrates robust memory for challenging classes such as actinic keratosis and melanoma, which were previously underrepresented. The model attained an average macro F1-score of 0.982 and a ROC-AUC over 0.99, so corroborating the efficacy of the balancing technique. Table 4 provides a comprehensive comparison of model performance, illustrating the efficacy of CNN models across all measures when trained with GAN-augmented data.

Table 4. Performance of CNN Models on ISIC (ORG vs. GAN)

Dataset	Model	Data	F1-Score	ROC-AUC	Balanced Acc	Training Acc	Validation Acc
Skin Cancer	EfficientNetB3	ORG	62.01	94.76	62.12	91.30	69.27
		GAN	94.05	99.40	94.22	94.57	94.13
	ResNet50	ORG	60.77	94.15	60.59	90.38	68.82
		GAN	92.94	99.26	93.15	92.74	92.99
	VGG19	ORG	59.74	94.65	60.57	77.57	70.92
		GAN	87.52	98.88	88.12	89.87	88.08
	MobileNet	ORG	76.68	95.73	66.85	89.51	77.40
		GAN	93.63	99.32	93.88	93.88	93.71




Figure 11. (a) The accuracy Training/validation (b) Confusion Matrix for ISIC Fold 3.

5.3. NAFLD Tabular Classification Results

A tabular MLP classifier was assessed on the original (ORG) and CycleGAN-augmented (GAN) data for the NAFLD dataset. Table 5 demonstrates that the implementation of CycleGAN-Real oversampling markedly enhanced classification performance across all criteria. The model trained on GAN-augmented data shown a significant enhancement in macro F1-score, with MobileNet-based features attaining 85.96%, in contrast to merely 31.67% on the original data. Comparable patterns were noted in balancing accuracy, ROC-AUC, and validation accuracy. The findings indicate that the incorporation of CycleGAN-Real oversampling and min-max normalization significantly improved the generalization capacity of the MLP classifier, especially for underrepresented steatosis stages in the imbalanced NAFLD dataset.

5.4. Cross-Fold Metric Comparison

Figure 13 presents bar charts that compare the F1-score and balanced accuracy across all folds for each dataset. The findings indicate negligible variation among folds, validating the model's robustness. HAM10000 consistently

Table 5. Performance of MLP Classifier on Original and CycleGAN-Augmented NAFLD Tabular Dataset

Dataset	Model	Data	F1-Score	ROC-AUC	Balanced Acc	Training Acc	Validation Acc
NAFLD	EfficientNetB3	ORG	29.82	68.80	33.72	72.59	57.49
		GAN	86.38	95.68	87.00	91.22	91.67
	ResNet50	ORG	46.44	72.18	49.00	92.29	55.26
		GAN	36.29	81.26	42.84	84.52	55.09
	VGG19	ORG	43.09	70.18	46.89	87.23	59.47
		GAN	91.62	96.71	91.67	93.24	91.67
	MobileNet	ORG	31.67	67.57	31.94	92.11	55.56
		GAN	85.96	94.31	86.26	94.40	86.22




Figure 12. (a) The accuracy Training/validation (b) Confusion Matrix for NAFLD Fold 1.

produced the highest results, whereas NAFLD exhibited dependable performance in categorizing stages of liver steatosis. Insert bar charts for macro F1 and balanced accuracy here.




Figure 13. Model Architecture.

5.5. Discussion

The incorporation of CycleGAN balancing and targeted augmentation proved essential for alleviating class imbalance and improving model performance. Specifically, samples generated by CycleGAN exhibited semantic variety that conventional oversampling methods could not attain. Moreover, the implementation of Focal Loss enhanced sensitivity towards underrepresented classes. EfficientNet-B3 significantly surpassed standard CNNs, such as VGG16 and ResNet18, in performance on the HAM10000 and ISIC datasets, as documented in related studies. The tabular MLP architecture accurately identified clinical biomarkers into distinct steatosis categories for NAFLD, confirming that deep learning techniques may be applied to non-image medical data when appropriately preprocessed and balanced. Beyond technical performance, this study recognizes the ethical and clinical considerations associated with synthetic data generation. Although CycleGAN successfully mitigates class imbalance, the generated samples may still contain subtle artifacts that could influence diagnostic interpretation if applied directly in clinical practice. To address this, future work will include close collaboration with medical specialists, who will visually inspect and validate the realism and diagnostic plausibility of the generated images. In addition, a fairness assessment will be conducted, particularly on the NAFLD dataset, to evaluate potential demographic or sampling bias. These steps aim to ensure that the proposed framework not only performs effectively but also upholds clinical transparency, reliability, and ethical integrity in real-world healthcare applications.

5.6. Comparison Analysis

A comprehensive comparison was conducted to illustrate the efficacy of the proposed framework against numerous current state-of-the-art researches in skin lesion classification. The investigation emphasizes classification performance, data balancing techniques, and model robustness in the context of imbalanced medical imaging datasets. Your proposed approach, which incorporates CycleGAN-based class balancing, EfficientNet-B3, and Focal Loss, demonstrated exceptional performance across all assessment parameters. The model achieved a maximum accuracy of 99.33%, a macro F1-score of 96.85%, and a ROC-AUC of 0.9852, establishing it as one of the most efficient pipelines documented in the literature to date.

5.6.1. Comparison to Prior Works To evaluate the proposed CycleGAN-Focal Loss-EfficientNet pipeline, we compared our results with several recent studies addressing class imbalance in medical classification tasks utilizing GANs or advanced balancing strategies. Table 6 presents a summary of the comparative results. The studies utilized diverse methodologies, including intra-class boundary augmentation, self-attention GANs, and style knowledge transfer.

IB-GAN (2024) presents an intra-class boundary-aware GAN designed to generate samples in proximity to decision margins. The framework attained an accuracy of 95.67% on a retinal fundus image dataset. Although it successfully improved class boundaries, it exhibited insufficient semantic variation in the generated images. In contrast, the samples generated by our CycleGAN exhibited enhanced structural diversity and intra-class realism, leading to improved classification performance in rare classes [31].

Style-KD (2023) introduced a method for style knowledge distillation aimed at addressing the imbalance in dermoscopic datasets. Accuracy was reported as 94.02%. While style transfer facilitated feature alignment, the approach did not generate new samples, thereby constraining enhancements in minority-class representation. Our approach utilizes CycleGAN to directly augment the volume of the minority class, thereby improving generalization and recall [32].

Chen and Du (2025) examined the issue of medical text imbalance through the implementation of a self-attentive adversarial augmentation network (SAAN) utilizing BERT. The results indicated an accuracy of 96.02%. Their

framework, while robust in NLP domains, is inapplicable to image classification. Our method, evaluated through dermoscopic and clinical imaging, demonstrates robustness in both image-based and tabular domains [33].

CycleGAN with Mask Self-Attention (2023) incorporated attention modules into CycleGAN to enhance the quality of generated samples. The reported accuracy is 96.90%. Although this is comparable to our CycleGAN baseline, their method incurs higher computational costs. Our model demonstrates performance that is comparable to or exceeds that of existing methods, while eliminating the overhead associated with self-attention, particularly across various datasets [34].

GANSMIA-CDIA (2025) utilized Generative Adversarial Networks for the augmentation of dermoscopic images. The reported metrics indicated an accuracy of 94.78%. However, the analysis was limited to a single dataset and binary classification. The framework accommodates multi-class scenarios and demonstrates effective generalization across three medical datasets: HAM10000, ISIC, and NAFLD [6].

Table 6. Comparison of the proposed models with existing studies.

Ref.	Study	Method	Acc.	F1	AUC	Balancing Method
[31]	IB-GAN (2024)	Boundary-focused generation	95.67%	93.41%	0.9712	Boundary-aware GAN
[32]	Style-KD (2023)	Style transfer + distillation	94.02%	91.76%	0.9650	Style transfer
[33]	Chen & Du (2025)	GAN + multi-task BERT	96.02%	94.12%	0.9723	SAAN + BERT
[34]	CycleGAN + Attn (2023)	Self-attention CycleGAN	96.90%	95.30%	0.9780	CycleGAN + attention
[6]	GANSMIA (2025)	Dermoscopic GAN augmentation	94.78%	93.00%	0.9600	Standard GAN

The comparison demonstrates that the CycleGAN-augmented EfficientNet-B3 model surpasses current methods in overall accuracy and per-class metrics. The incorporation of GAN-based data balancing, robust augmentation, and Focal Loss markedly improves classification performance on highly imbalanced datasets. This method overcomes significant shortcomings of previous studies and establishes a new standard for multi-class medical image classification.

6. Conclusion

This study introduces a cohesive deep learning pipeline designed to tackle the significant issue of class imbalance in multi-class medical data classification. The integration of CycleGAN-based synthetic oversampling with Focal Loss and EfficientNet-B3 results in significant enhancements in model sensitivity, especially for minority classes, in both image and tabular datasets. This study presents a consistent and scalable CycleGAN-based balancing method applicable to diverse medical domains, specifically dermoscopic images (HAM10000, ISIC) and tabular clinical records (NAFLD). The application of real CycleGAN oversampling across domains is innovative and exhibits performance that surpasses prior achievements documented in the literature. The model attained a peak accuracy of 99.33%, a macro F1-score of 96.85%, and a ROC-AUC of 0.9852, establishing a new standard for imbalanced multi-class medical classification. The proposed architecture enhances accuracy and markedly improves performance across individual classes, especially for rare and clinically significant conditions. The integration of real-time augmentation, adaptive loss functions, and stratified cross-validation enhances the model's robustness and generalizability. This study, while yielding promising results, is constrained by inherent limitations typical of data-driven models, including reliance on dataset quality, potential annotation bias, and the computational expenses associated with GAN training. Nonetheless, these factors do not undermine the innovation or practical contribution of the method. This approach has the potential for extension to additional medical domains, scaling to larger multi-modal datasets, and further optimization for real-time clinical integration via lightweight architectures or model pruning. Future research may investigate self-supervised CycleGAN training or hybrid balancing strategies to improve scalability. This work is distinguished by its innovative cross-domain application of real

CycleGAN-based oversampling, demonstrating minority-class detection performance that, through comparative analysis, exceeds existing methods and establishes a new benchmark in the field.

Acknowledgement

The authors thank Akre University for Applied Science and Duhok Polytechnic University for their academic support. We also acknowledge the creators of the HAM10000, ISIC, and NAFLD datasets for providing the medical data used in this study.

REFERENCES

1. Adnan Mohsin Abdulazeez and Shereen Sadiq Hasan. Classification of heart diseases based on machine learning: A review. *International Journal of Informatics Information System and Computer Engineering*, 6(1):31–49, June 2025. Received 14 Mar 2024; Revised 11 Apr 2024; Accepted 02 Jun 2024; Published online 07 Aug 2024.
2. Paul A. Adedeji, Johnson A. Oyewale, Tunde I. Ogedengbe, Obafemi O. Olatunji, and Nkosinathi Madushele. Comparative assessment of soft computing and svm architectures for multi-class automobile engine fault classification. *International Journal of System Assurance Engineering and Management*, 16(5):1743–1756, March 2025.
3. I Nyoman Mahayasa Adiputra, Pei-Chun Lin, and Paweena Wanchai. The effectiveness of generative adversarial network-based oversampling methods for imbalanced multi-class credit score classification. *Electronics*, 14(4):697, February 2025.
4. Sania Akhtar, Muhammad Hanif, Ahmar Rashid, Khurshed Aurangzeb, Ejaz Ahmad Khan, Hamdi Melih Saraoglu, and Kamran Javed. An optimized data and model centric approach for multi-class automated urine sediment classification. *IEEE Access*, 12:59500–59520, 2024.
5. Madallah Alruwaili and Mahmood Mohamed. An integrated deep learning model with efficientnet and resnet for accurate multi-class skin disease classification. *Diagnostics*, 15(5):551, February 2025.
6. AMAL ALSHARDAN, SAAD ALAHMARI, MOHAMMED ALGHAMDI, MUTASIM AL SADIG, ABDULLAH MOHAMED, and GOUSE PASHA MOHAMMED. Gan-based synthetic medical image augmentation for class imbalanced dermoscopic image analysis. *Fractals*, 33(02), January 2025.
7. Imane Araf, Ali Idri, and Ikram Chairi. Cost-sensitive learning for imbalanced medical data: a review. *Artificial Intelligence Review*, 57(4), March 2024.
8. Samir Brahim Belhaouari, Ashhadul Islam, Khelil Kassoul, Ala Al-Fuqaha, and Abdesselam Bouzerdoun. Oversampling techniques for imbalanced data in regression. *Expert Systems with Applications*, 252:124118, October 2024.
9. Doanh C. Bui, Boram Song, Kyungeun Kim, and Jin Tae Kwak. Dax-net: A dual-branch dual-task adaptive cross-weight feature fusion network for robust multi-class cancer classification in pathology images. *Computer Methods and Programs in Biomedicine*, 248:108112, May 2024.
10. Pengdi Chen, Yong Liu, Yuanrui Ren, Baoan Zhang, and Yuan Zhao. A deep learning-based solution to the class imbalance problem in high-resolution land cover classification. *Remote Sensing*, 17(11):1845, May 2025.
11. Hongwei Ding, Nana Huang, and Xiaohui Cui. Leveraging gans data augmentation for imbalanced medical image classification. *Applied Soft Computing*, 165:112050, November 2024.
12. Veronica Distefano, Monica Palma, and Sandra De Iaco. Multi-class random forest model to classify wastewater treatment imbalanced data. *Socio-Economic Planning Sciences*, 95:102021, October 2024.
13. Kushankur Ghosh, Colin Bellinger, Roberto Corizzo, Paula Branco, Bartosz Krawczyk, and Nathalie Japkowicz. The class imbalance problem in deep learning. *Machine Learning*, 113(7):4845–4901, December 2022.
14. Ross Greer, Akshay Gopalkrishnan, Maitrayee Keskar, and Mohan M. Trivedi. Patterns of vehicle lights: Addressing complexities of camera-based vehicle light datasets and metrics. *Pattern Recognition Letters*, 178:209–215, February 2024.
15. Folasade Olubusola Isinkaye, Michael Olusoji Olusanya, and Ayobami Andronicus Akinyelu. A multi-class hybrid variational autoencoder and vision transformer model for enhanced plant disease identification. *Intelligent Systems with Applications*, 26:200490, June 2025.
16. Najdavan A. Kako, Adnan M. Abdulazeez, and Diler N. Abdulqader. Multi-label deep learning for comprehensive optic nerve head segmentation through data of fundus images. *Heliyon*, 10(18):e36996, September 2024.
17. Hesham Kamal and Maggie Mashaly. Enhanced hybrid deep learning models-based anomaly detection method for two-stage binary and multi-class classification of attacks in intrusion detection systems. *Algorithms*, 18(2):69, January 2025.

18. Dewant Katare, David Solans Noguero, Souneil Park, Nicolas Kourtellis, Marijn Janssen, and Aaron Yi Ding. Analyzing and mitigating bias for vulnerable road users by addressing class imbalance in datasets. *IEEE Open Journal of Intelligent Transportation Systems*, 6:590–604, 2025.
19. Kurnianingsih, Sou Nobukawa, Melyana Nurul Widyawati, Cipta Pramana, Nurseno Bayu Aji, Afandi Nur Aziz Thohari, Dwiana Hendrawati, Eri Sato-Shimokawara, and Naoyuki Kubota. A novel ensemble xgboost and deep q-network for pregnancy risk prediction on multi-class imbalanced datasets. *ICT Express*, 11(4):648–656, August 2025.
20. Shuxian Li, Liyan Song, Xiaoyu Wu, Zheng Hu, Yiu-ming Cheung, and Xin Yao. Multi-class imbalance classification based on data distribution and adaptive weights. *IEEE Transactions on Knowledge and Data Engineering*, 36(10):5265–5279, October 2024.
21. Junzhuo Liu, Zhixiang Wang, Ye Zhang, Alberto Traverso, Andre Dekker, Zhen Zhang, and Qiaosong Chen. Cyclegan clinical image augmentation based on mask self-attention mechanism. *IEEE Access*, 10:105942–105953, 2022.
22. Yajun Liu, Hong Fan, Jianguang Zhao, Jianfang Zhang, and Xinxin Yin. Efficient and generalized image-based cnn algorithm for multi-class malware detection. *IEEE Access*, 12:104317–104332, 2024.
23. Inhyuk Park, Won Hwa Kim, and Jongbin Ryu. Style-kd: Class-imbalanced medical image classification via style knowledge distillation. *Biomedical Signal Processing and Control*, 91:105928, May 2024.
24. Amirreza Salehi and Majid Khedmati. Hybrid clustering strategies for effective oversampling and undersampling in multiclass classification. *Scientific Reports*, 15(1), January 2025.
25. Merdin Shamal Salih and Adnan Mohsin Abdulazeez. A fusion-based deep approach for enhanced brain tumor classification. *Journal of Soft Computing and Data Mining*, 5(1), June 2024.
26. Prabu Selvam, P. Karthikeyan, S. Manochitra, A. V. L. N. Sujith, T. Ganesan, Rajaram Ayyasamy, Mohammed Shuaib, Shadab Alam, and A. Rajendran. Federated learning-based hybrid convolutional recurrent neural network for multi-class intrusion detection in iot networks. *Discover Internet of Things*, 5(1), April 2025.
27. Anees Tariq, Muhammad Munwar Iqbal, Muhammad Javed Iqbal, and and Iftikhar Ahmad. Transforming brain tumor detection empowering multi-class classification with vision transformers and efficientnetv2. *IEEE Access*, 13:63857–63876, 2025.
28. J.S. Thanga Purni and R. Vedhapriyavadhana. Eosa-net: A deep learning framework for enhanced multi-class skin cancer classification using optimized convolutional neural networks. *Journal of King Saud University - Computer and Information Sciences*, 36(3):102007, March 2024.
29. Lei Tong, Adam Corrigan, Navin Rathna Kumar, Kerry Hallbrook, Jonathan Orme, Yinhai Wang, and Huiyu Zhou. Clanet: A comprehensive framework for cross-batch cell line identification using brightfield images. *Medical Image Analysis*, 94:103123, May 2024.
30. Alex X. Wang, Viet-Tuan Le, Hau Nguyen Trung, and Binh P. Nguyen. Addressing imbalance in health data: Synthetic minority oversampling using deep learning. *Computers in Biology and Medicine*, 188:109830, April 2025.
31. Akdeas Oktanae Widodo, Bambang Setiawan, and Rarasmaya Indraswari. Machine learning-based intrusion detection on multi-class imbalanced dataset using smote. *Procedia Computer Science*, 234:578–583, 2024.
32. Lang Wu. A meta-learning network method for few-shot multi-class classification problems with numerical data. *Complex & Intelligent Systems*, 10(2):2639–2652, December 2023.
33. Jenny Yang, Rasheed El-Bouri, Odhran O'Donoghue, Alexander S. Lachapelle, Andrew A. S. Soltan, David W. Eyre, Lei Lu, and David A. Clifton. Deep reinforcement learning for multi-class imbalanced training: applications in healthcare. *Machine Learning*, 113(5):2655–2674, November 2023.
34. Gheyath Mustafa Zebari, Diloan Asaad Zebari, Diyar Qader Zeebaree, Habibollah Haron, Adnan Mohsin Abdulazeez, and Kamil Yurtkan. Efficient cnn approach for facial expression recognition. *Journal of Physics: Conference Series*, 2129(1):012083, December 2021.