



HAL
open science

Diabetic Retinopathy Screening within Unlabeled Dataset based on Least Squares Cycle-GAN Domain Transfer

Zineb Sadok, Mohamed Akil, Rostom Kachouri, Ali Ahaitouf

► **To cite this version:**

Zineb Sadok, Mohamed Akil, Rostom Kachouri, Ali Ahaitouf. Diabetic Retinopathy Screening within Unlabeled Dataset based on Least Squares Cycle-GAN Domain Transfer. The 13th International Conference on Image Processing Theory, Tools and Applications IPTA 2024, Oct 2024, Rabat, Morocco. hal-04813357

HAL Id: hal-04813357

<https://hal.u-pec.fr/hal-04813357v1>

Submitted on 1 Dec 2024

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

Diabetic Retinopathy Screening within Unlabeled Dataset based on Least Squares Cycle-GAN Domain Transfer.

SADOK Zineb
Laboratory of Intelligent Systems, Georesources and Renewable Energies,
Faculty of Sciences and Technology,
Sidi Mohamed Ben Abdellah
University
Fez, Morocco.
Gaspard Monge Computer Science
Laboratory, ESIEE Paris, Gustave
Eiffel University
F-77454 Marne-la-Vallée, France.
zineb.sadok@usmba.ac.ma
zineb.sadok@esiee.fr

AKIL Mohamed
Gaspard Monge Computer Science
Laboratory, ESIEE Paris, Gustave
Eiffel University
F-77454 Marne-la-Vallée, France.
Mohammed.akil@esiee.fr

KACHOURI Rostom
Gaspard Monge Computer Science
Laboratory, ESIEE Paris, Gustave
Eiffel University
F-77454 Marne-la-Vallée, France.
Rostom.kachouri@esiee.fr

AHAITOUF Ali
Laboratory of Intelligent Systems, Georesources and Renewable Energies,
Faculty of Sciences and Technology,
Sidi Mohamed Ben Abdellah
University
Fez, Morocco.
Ali.ahaitouf@usmba.ac.ma

Abstract

Diabetic retinopathy (DR) causes blindness in young adults worldwide. Thanks to early detection, patients with diabetic retinopathy can be properly treated in time, and the deterioration of diabetic retinopathy can be prevented. Early detection is therefore essential for screening the DR disease. To perform a supervised DR classification, we need labeled images, otherwise we are required to conduct a traditional manual diagnosis, with the help of an ophthalmologist, which is time-consuming and costly expensive. The question here is what we can do if we have unlabeled images? We can benefit from the knowledge of a pre-trained classifier using a labeled base, but the problem is that the images in the datasets may come from different domains. To solve all these problems, we propose a three-step method: first, train a classifier using a labeled dataset; second, adapt the unlabeled dataset (source domain) to the labeled dataset (target domain) using Least Squares Cycle-GAN; and finally, classify these adapted images using the pre-trained classifier on the target domain. Following this procedure, the results show that we succeeded in classifying 77% of unlabeled images, meaning that the ophthalmologist can concentrate on the remaining 23% of cases for diagnosis, which is clear time saving. This approach is contribution to the existing supervised learning and transfer learning methods that are requiring a lot of labeled data, which is always not available in DR screening due to the time consuming of image fundus annotating.

Keywords

Diabetic Retinopathy Grading, Generative Adversarial Network, Domain Transfer, Cycle-GAN, Unlabeled Dataset, Retinal image.

I. INTRODUCTION

A supervised classification of DR necessitates a labeled dataset, which is not always readily available. Domain adaptation consists of the transfer a trained model in a source

domain, where dataset may be limited and unlabeled, to a target domain with labeled dataset.

In the context of this study, a private database will be considered for the automatic diagnosis of the RD. However, this private dataset, though interpretable by expert ophthalmologists, is limited in size and exhibits an imbalance in diabetic retinopathy stages, along with a lack of labeling for different signs of DR. Consequently, the application of supervised machine learning approaches for classifying or segmenting these signs is difficult and even practically impossible. Domain adaptation will then be helpful to classify this dataset and more generally non-labelled datasets.

Then, we use our private retinal fundus image dataset [1] as the source domain and the IDRiD (Indian Diabetic Retinopathy Images Dataset) [2] as the target domain.

To address this, we propose a Least Squares (LS) Cycle-GAN to generate fundus images. This adaptation enhances the original Cycle-GAN framework by adopting a least squares loss function, which mitigates issues like the Vanishing Gradient problem [3] during training and preserves essential retinal fundus image characteristics.




Figure 1 : Retinal Unpaired images from different domain, IDRiD dataset is the Source Domain (X) and Private dataset is the Target Domain (Y).

Before initiating domain adaptation, we first prepared a classifier using the IDRiD dataset (Target Domain) to classify diabetic retinopathy stages. The IDRiD dataset supports both diabetic retinopathy grade classification and segmentation of various associated signs.

To verify the quality of the generated images, the Image Spatial Quality Evaluator (BRISQUE) and Shannon Entropy were employed to assess each generated image and its original. Finally, to validate the effectiveness of our domain adaptation approach, we use t-Distributed Stochastic Neighbor Embedding (t-SNE) Visualization [4] for domain alignment. This visualization will show how closely the generated images align with the surface of the target domain images. The remainder of this paper is organized as follows: Section II reviews related work, including applications of GANs in fundus imaging and transfer learning approaches for DR grading. Section III describes our proposed method in three steps: DR Grading, LS Cycle-GAN, and classification of adapted images. Section IV details our experiments, including the framework, datasets, data processing, and evaluation metrics. Finally, we present the results and conclude the paper.

II. RELATED WORK

The definition of a domain for images in the context of machine learning and computer vision refers to a specific distribution of images along with their corresponding labels or characteristics such as a set of images and the associated features or labels that related to those images. This can include various attributes such as pixel values, color histograms, texture features, object classes, or any other relevant information that characterizes the images within that domain. Formally, a domain D is represented by a pair $(X, P(X))$ where X is the feature space or input space where the data points reside, it defines the possible values that the data instances can take, and $P(X)$ is a probability distribution over the feature space X , it specifies the likelihood or probability of observing different values or configurations of features within the domain [5].

The differences between domains can manifest themselves in variations in the feature space or in the data distribution [6]. The Generative Adversarial Networks (GANs) [7] are sophisticated deep learning techniques widely used for realistic image generation across several applications.

The architecture of GANs consists of two key components: a Generator Network that learns to produce data indistinguishable from real data, and a Discriminator Network that learns to distinguish between generated data and real data. This adversarial setup ensures that the generator continually improves its ability to produce realistic outputs, guided by the feedback from the discriminator.

The Cycle-GAN [8], a type of generative adversarial networks (GANs), is among the recent advancements in domain transfer methods. It comprises two paired generator-discriminator modules, tasked with learning two mappings: from domain X to domain Y $\{G, D_y\}$ and the inverse Y to X $\{F, D_x\}$. The generators (G, F) translate images between the source and target domains, while the discriminators (D_x, D_y) aim to distinguish the original data from the translated ones.

The Cycle-GAN is guided by two losses, the adversarial loss L_{adv} and the cycle-consistency loss L_{cycle} .

Adversarial loss enhances the local realism of the translated data. If we take the adversarial loss of the translation from domain X to domain Y as an example, it can be written as follows:

$$L_{GAN}(G, D_Y, X, Y) = E_{x \sim p_{data}(x)} [\log(1 - D_Y(G(x)))] + E_{y \sim p_{data}(y)} [\log D_Y(y)] \quad (1)$$

The expectation terms $E_{x \sim p_{data}(x)}$ and $E_{y \sim p_{data}(y)}$ denote averages over the data distributions $p_{data}(x)$ and $p_{data}(y)$. The loss of cycle consistency overcomes the requirement for matched training data. The concept of cycle consistency loss is that the converted data from the target domain can be transcribed back into the source domain, $x \rightarrow G(x) \rightarrow F(G(x)) \approx x$, which can be noted as follows:

$$L_{cycle}(G, F) = E_{x \sim p_{data}(x)} [\|F(G(x)) - x\|_1] + E_{y \sim p_{data}(y)} [\|G(F(y)) - y\|_1] \quad (2)$$

The $\|\cdot\|_1$ denotes the L_1 norm, measuring the absolute difference between two vectors, employed in the cycle-consistency loss $L_{cycle}(G, F)$ to ensure consistent mappings between X and Y .

The application of GANs and their various versions, such as Cycle-GAN, in ophthalmology imaging has garnered significant attention for their effectiveness across various facets of medical image analysis as we will explore in the next section.

A. Applications of Generative Adversarial Networks in fundus image

Various studies have leveraged Generative Adversarial Networks (GANs) for diverse applications in medical imaging. In reference [9], researchers applied GANs to denoise retinal OCT (Optical Coherence Tomography) images, aiming to improve image quality and diagnostic accuracy. Another study [10] focused on enhancing the resolution of fundus photography using GANs, which is crucial for detailed examination of the retina. Research in reference [11] explored GAN-based data augmentation techniques specifically for OCT chorio-retinal segmentation tasks, enhancing the robustness and generalizability of segmentation algorithms. Additionally, reference [12] investigated GANs for optic disc and cup segmentation in fundus photography, demonstrating advancements in automated analysis and diagnosis of glaucoma. These studies illustrate the versatility of GANs in medical imaging, addressing various challenges from denoising and resolution enhancement to segmentation and diagnostic support.

These advancements highlight the potential of GANs, opening new avenues for precise and personalized diagnostic applications in ophthalmology. For example, researchers in Ref. [13] utilized Cycle-GAN to successfully translate images from the ultra-widefield fundus photography (UWFP)

domain to the traditional fundus photography (TFP) domain, preserving the essential structural details of the retina and optic nerve. Another study [14].

Recent advancements in GANs have also been tailored for domain adaptation in medical image segmentation. One notable example is the Edge-Cycle-GAN [15], designed to retain edge details during image translation between domains and more recently, the Structure-Preserving Cycle-GAN (SP Cycle-GAN) have been introduced [16] to enhance segmentation accuracy in unsupervised medical image adaptation.

B. Transfer learning approaches for Diabetic Retinopathy Grading (RD):

Machine learning, deep learning and transfer learning technologies have made remarkable progress in the automated detection of DR. Several research studies have explored to integrate domain adaptation into the convolutional neural network [17][18]. To address the problem of insufficient labeled data, authors in [17] sought to refine the pre-trained InceptionNet-V3 model for classification by subsampling a smaller version of the EyePACS [19] dataset and testing on a previously unseen subset of the data, with an accuracy of 90.9%. In the author study [18] the VGG19 model [20] was trained using the IDRiD dataset and extracted features from retinal images, which are then fed into different classifiers such as logistic regression (LR), SVM (Support Vector Machine) and KNN (K-nearest neighbors). This work has demonstrated that transfer learning is potentially effective for retinal image classification by distilling useful knowledge from source images.

classification using a target generation model. This method aims to enhance diabetic retinopathy detection accuracy through the benefits of transfer learning for automated retinal image analysis. In this paper, we are considering further this transfer learning and we will use it to evaluate fundus images for the diabetic retinopathy (DR) screening. In fact, we aim to solve the problem of images classification from unlabeled databases (Source Domain), by classifying images generated by the LS Cycle-GAN from the pre-trained source model with images from labeled databases (Target Domain).

III. THE PROPOSED METHOD

To classify an unlabeled dataset using supervised learning, our proposed method involves three sequential steps. Figure 2 summarizes these steps, with the first step involving the preparation of a pre-trained classifier using the IDRiD dataset [2]. The second step employs the LS Cycle-GAN to adapt the unlabeled datasets (Source Domain) to the domain of the labeled dataset (Target Domain). Finally, the third step involves classifying these adapted images using the classifier.

A. Deep Neural Network and RD Grading "STEP 1"

In the first step we utilizing a pre-trained DenseNet121 model [22] (Figure 2, Step 1) on the IDRiD dataset to classify generated images. Fine-tuning begins by loading the pre-trained DenseNet121 model and adapting it for target classification. Initially, we remove the last fully-connected layer, originally designed for 1000 classes in ImageNet, and replace it with three dense layers: the first and second with 1024 and 512 neurons respectively activated by ReLU, and finally an output layer with 2 neurons activated by softmax for a binary classification (No RD, RD).

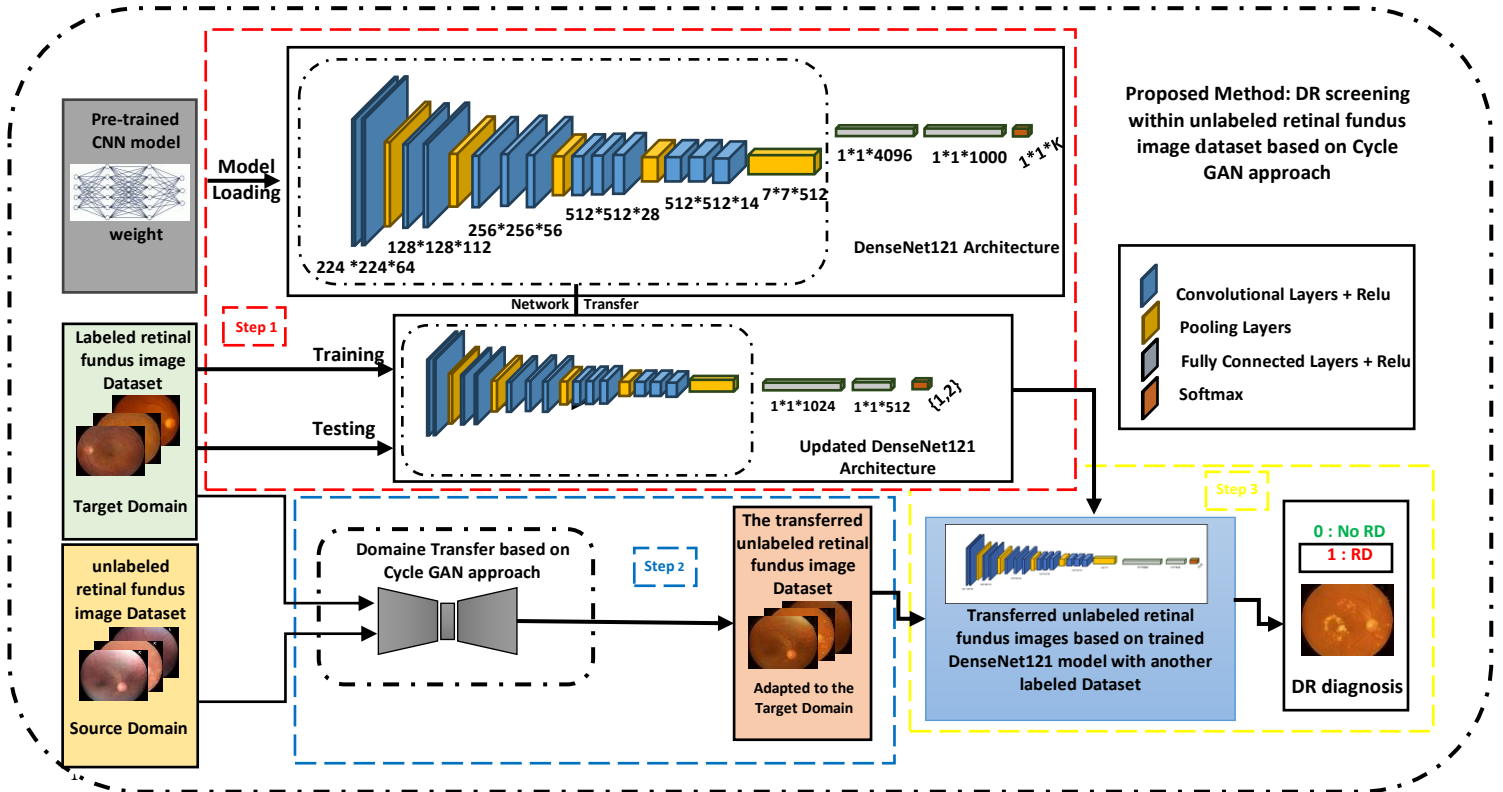


Figure 2 : Framework of our Proposed Approach and its Different Steps in Domain Adaptation for Diabetic Retinopathy Classification of the Unlabeled Retinal Fundus Images

free transfer learning approach for diabetic retinopathy « warm-up » and "deep tuning" phases of training-validation.

The model's weights are initially frozen to retain previously learned features. Then, training proceeds using the Adam optimizer with a learning rate (LR) set to $\{10^{-5}, 10^{-4}\}$, employing a binary cross-entropy loss function. To prevent overfitting, we implement a 10-fold cross-validation strategy. The warm-up phase spans 120 epochs, followed by 80 epochs for the fine-tuning phase. Early stopping criteria are applied throughout both stages, halting training if validation accuracy fails to improve over 20 consecutive epochs. A comprehensive summary of hyperparameters is provided in Table 1.

Table 1: Specified Hyperparameters of Classification Model.

Learning rate	Optimizer	Batch size	Epochs	Early stop
$\{10^{-5}, 10^{-4}\}$	Adam	{12 ;32}	{120,80}	20

After the initial training, the model enters a secondary training phase to further enhance its performance. This is demonstrated in recent studies [23], which have shown the effectiveness of "deep tuning," involving the fine-tuning of all model layers under appropriate conditions. The second step, based on domain adaptation, involves the use of the Least-Square Cycle-GAN framework to adapt the unlabeled dataset from the source domain (Private dataset) to the labeled dataset in the target domain (IDRiD dataset). That is the object of the next section.

This modification aims to mitigate issues such as vanishing gradients, which are particularly critical in medical image generation tasks like producing retina images.

The LSGAN-specific loss function replaces the adversarial loss between the generators G and F , ensuring better preservation of essential structural details in the generated images. This update is crucial for improving the fidelity and reliability of generated images, thereby supporting accurate medical diagnosis and analysis, it can be written as follows:

$$L_{LSGAN}(G, D_Y, X, Y) = \frac{1}{2} E_{x \sim p_{data}(x)} [(D(x) - 1)^2] + \frac{1}{2} E_{y \sim p_{data}(y)} [G(D(y))^2] \quad (3)$$

LS Cycle-GAN loss is defined as follows:

$$L_{LSCycleGAN}(G, F, D_x, D_y) = L_{LSGAN}(G, D_y, X, Y) + L_{LSGAN}(F, D_x, Y, X) + L_{Cycle}(G, F) \quad (4)$$

Our method enables retinal images to be generated by adapting unannotated images from a source domain X to an annotated target domain Y , overcoming the problem of unannotated retinal databases. Figure 3 simply illustrates the LS Cycle-GAN-based domain transfer method, which utilizes two types of inputs: images from the source domain (unlabeled database) and images from the target domain (labeled database).

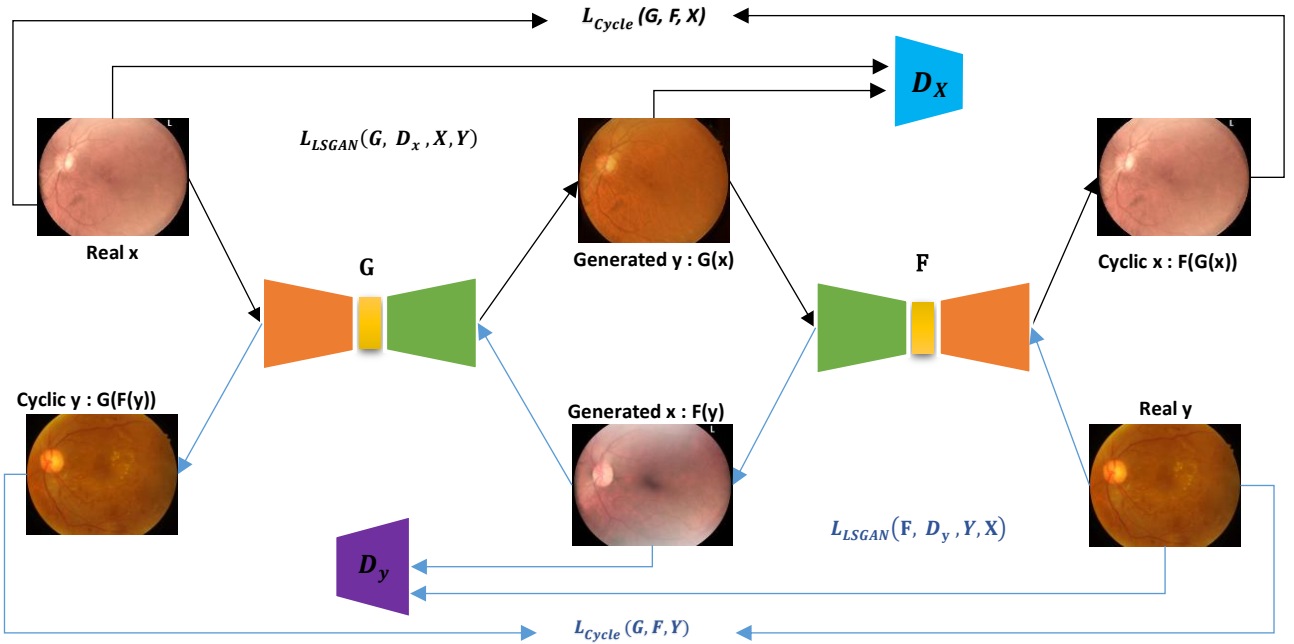


Figure 3 : Domaine Transfer based on Least Square Cycle-GAN approach

B. Least Square Cycle-GAN "STEP 2"

Least Squares Cycle-GAN (LS) Cycle-GAN was implemented by updating the original Cycle-GAN loss in Equation (2), focusing specifically on GAN losses (1). We moved from the standard GAN loss function to LSGAN in Equation (3) [3].

After generation, generators G and F produce two adapted images for the target and source domains, respectively. To assess the quality of the generated images, we employ the discriminator and the LSGAN loss function, updating the weights of the generators and discriminator to enhance generation.

Subsequently, both cycles are completed to compute the cycle loss. This involves reversing the roles of the two

generators and regenerating the input images (real images), a process known as "Cyclic image". The transferred unlabeled retinal fundus images dataset serves as input for the classifier, which has been pretrained on the target domain. This allows us to leverage its expertise in this domain.

C. Classification of Adapted Images "STEP 3"

The third step is the final stage, such as we validate our method and achieve our goal of Diabetic Retinopathy screening of an unlabeled dataset using LS Cycle-GAN. The transferred unlabeled retinal fundus images dataset serves as input for the classifier, which has been pretrained on the target domain. This allows us to leverage its expertise in this domain.

IV. EXPERIMENTS

A. Experimental Framework and implementation details

1) Datasets

The private dataset (1) includes 1,310 fundus images that can be interpreted by an ophthalmologist. For each image, the ophthalmologist indicated its stage, either "No RD" (no diabetic retinopathy) or "RD" (presence of diabetic retinopathy). Images were collected and captured using a portable, mobile retinal camera called oDocs nun IR. It provides high-resolution imaging (2880 x 2160 pixels) for precise retinal scanning, featuring a wide field of view (45 to 55 degrees) and manual focus (-20D to +20D). Compatible with Android smartphones, it utilizes a private dataset collected by El Mehdi Chakour from Omar Drissi Hospital in Fez, Morocco, under the supervision of Pr. Idriss Benatiya Andaloussi [1].

The Indian Diabetic Retinopathy Image Dataset (IDRiD) [2] includes information on the severity of diabetic retinopathy. Medical specialists evaluated all 516 photos, which included a variety of clinical situations related to diabetic retinopathy. Diabetic retina photos were divided into four groups, from 0 (no obvious DR) to 4 (severe DR). This database will provide us with the annotations that we need to improve the quality and diversity of annotations in our private database. We take group 0 for the class without RD and the other 4 groups for the RD class.

2) Data Progressing & Augmentation

Preprocessing of retinal images is essential to improve the outcomes of any analysis using these images. The preprocessing chain [1] consists of three main parts. Firstly, noise reduction is performed using a denoising filter such as the Gaussian filter (GF). Secondly, contrast enhancement is achieved using Contrast-Limited Adaptive Histogram Equalization (CLAHE). Lastly, the intensity values of pixels in an image are adjusted based on a gamma correction curve (GCC). Additionally, the images are cropped around the retina to eliminate dark aspects and spatially for IDRiD images. The images are then resized to (224×224). Finally, image normalization is applied to scale pixel intensities from [0,255] to [-1,1].

3) Evaluation Metrics

a) Image Quality Measurements

It is crucial to ensure that the domain adaptation has not contributed to degrading the image quality during transfer. To this end, we evaluate the quality of images in the private database before and after the domain transfer in order to measure the influence of the transfer on image quality.

We use unreferenced image quality measures, such as Image Spatial Quality Evaluator (BRISQUE) [24], [25] and Entropy $H(p)$ [25], to objectively assess these changes. BRISQUE is based on a regression model that relates features extracted from the image to predicted human quality scores. A score between 0 (best) and 100 (worst) is obtained in each case. Entropy quantifies the amount of information in an image, ranging from 0 to 8, with higher entropy values indicating a greater level of detail. The mathematical equation of Shannon Entropy $H(p)$ is given by:

$$H(p) = - \sum_{j=0}^{L-1} p(j) \log p(j) \quad (5)$$

Where p is the number of normalized histograms and j is the number of intensity levels.

b) Performance Metrics for Classification

The classification metrics offer valuable insights into various facets of the model's performance. These metrics assess the accuracy of data for the two classes, as well as the sensitivity of the model.

- **Sensitivity:** The proportion of correctly predicted positive observations in the class to actual positive observations.

$$\text{Sensitivity} = \frac{\text{True Positive}}{\text{True Negative} + \text{False Negative}} \quad (6)$$

- **Accuracy:** It's the ratio of correctly predicted observations to all observations.

$$\text{Accuracy} = \frac{\text{True Negative} + \text{True Positive}}{\text{Total}} \quad (7)$$

B. Results of our proposed method

Firstly, we verify the quality of images before and after domain transfer adaption, using the Image quality measurements. Table 2 shows that the application of domain transfer preserves image quality. The results show some improvement in image quality after the application of domain transfer, with a significant reduction in the mean BRISQUE scores and a slight increase in mean Entropy for No RD and RD images. This suggests that domain transfer is effective in improving the clarity and detail of fundus images.

Table 2: Impact of Domain Transfer on Image Quality

Images	Mean Brisque		Mean Entropy	
	No RD	RD	No RD	RD
Original Images	53,61162	52,07758	7,05221	7,04966
Adapted Images	18,35729	20,92824	7,18680	7,17876

Secondly, the classification employs a pre-trained classifier trained on the IDRiD dataset (Table 3). The DenseNet121 model was chosen for its superior performance during testing, achieving a test accuracy of 0.9128 and sensitivity of 0.9199, surpassing VGG16 (Test accuracy: 0.8313, Sensitivity: 0.8325) and Inception-V3 [26] (Test Accuracy: 0.8623, Sensitivity: 0.8701).

Table 3 : Obtained results on model evaluation, according to accuracy and sensitivity values.

Model	IDRiD database	
	Accuracy	Sensitivity
DenseNet121	0.9128	0.9199
Inception-V3	0.8623	0.8701
VGG16	0.8313	0.8325

The results presented in Table 4 illustrate the effectiveness of our domain adaptation approach using the DenseNet121 pre-trained model on the IDRiD dataset. Initially, the model applied to the unannotated private database achieved an accuracy of 70.15% and a sensitivity of 71.43%.

However, after domain adaptation, the DenseNet121 model's performance increased considerably, reaching an accuracy of 77.41% and a sensitivity of 78.28%. This significant improvement confirms the effectiveness of using the pre-trained model to classify the generated images.

Table 4: Performance of Classification of Images of DR Using the DenseNet121 Pre-trained Model on the IDRiD Dataset

Dataset	DenseNet121	
	Accuracy	Sensitivity
Original private	0,7015	0,7143
Adapted private	0,7741	0,7828

The classifier has never seen or trained on the private database, so these are directly the test results. We were able to classify 77% of the unlabeled images, meaning that the ophthalmologist can only focus on the remaining 23% of cases for diagnoses. To have a clear idea on the generated images, we call the t-Distributed Stochastic Neighbor Embedding (t-SNE) [4] which is a powerful nonlinear technique to reduce the dimensionality of high-dimensional feature representations to a low-dimensional space (typically two or three dimensions) for a clear visualization.




Figure 4 : t-SNE visualization of image characteristics of source, target and generated domain.

In our context, when the classifier performs effectively, features extracted from images belonging to the same category tend to cluster closely together in the t-SNE plot. The t-SNE visualization is illustrated in Figure 4. The blue dots (source) are clearly separated from the red dots (target) and green dots (generated), indicating significant differences between the source and the target domains. The green dots are mainly located between the red and the blue dots, showing that the generated images have intermediate characteristics and are still closer to the red dots than the blue ones.

In addition, we can see a visual comparison of the generated images in Figure 5. We note that an excellent adaptation of these images (B) is ensured and a good similarity to the target domain images (C) is obtained.




Figure 5 : (A) the real images of the private database (source domain), (B) the generated images of the private database adapted to the IDRiD database domain, and (C) the real images of the IDRiD database (target domain).

V. CONCLUSION

In this work, we introduce LS-Cycle-GAN, a domain adaptation method designed for diabetic retinopathy classification using unlabeled images from the source domain and a pre-trained classifier on images from the target domain. Our approach focuses on maintaining high image quality throughout the adaptation process. The obtained results illustrate successful adaptation of generated images in the target domain. Our proposed method reduces significantly the manual diagnosis workload for ophthalmologists. Specifically, it reduces the need for manual diagnosis to only 23% of unconfirmed images, while automatically diagnosing the remaining 77%. This streamlined approach not only optimizes the diagnostic process but also saves considerable time, demonstrating the potential of LS-Cycle-GAN in enhancing efficiency and accuracy in medical image analysis.

KNOWLEDGE

This work is carried out in collaboration with the “Intelligent Systems, Geo-resources and Renewable Energies” laboratory, USMBA-FST Fez Morocco and the Gaspard-Monge Computer Laboratory, UGE-ESIEE Paris France, as part of the MAGHREB 2023 PHC project (Project Code: 48672RG).

REFERENCES

- [1] EM Chakour, Z Sadok, M Akil, R Kachouri, A Mansouri, I Benatiya, A Ahaitouf ‘Color fundus image enhancement - A deep learning based desktop app for earlier screening of diabetic retinopathy using real-time handheld fundus camera’, in *International Conference on Optimization and Data Science in Industrial Engineering (ODSIE 2023)*, Istanbul, Turkey, Nov. 2023. Accessed: Jun. 14, 2024.
- [2] ‘IDRI - Grand Challenge’, grand-challenge.org. Accessed: May 16, 2024. [Online].
- [3] X. Mao, Q. Li, H. Xie, R. Y. K. Lau, Z. Wang, and S. P. Smolley, ‘Least Squares Generative Adversarial Networks’. arXiv, Apr. 05, 2017. doi: 10.48550/arXiv.1611.04076.
- [4] L. van der Maaten and G. Hinton, ‘Visualizing data using t-SNE’, *Journal of Machine Learning Research*, vol. 9, pp. 2579–2605, Nov. 2008.
- [5] S. Ben-David, J. Blitzer, K. Crammer, A. Kulesza, F. Pereira, and J. W. Vaughan, ‘A theory of learning from different domains’, *Mach Learn*, vol. 79, no. 1, pp. 151–175, May 2010, doi: 10.1007/s10994-009-5152-4.
- [6] ‘A Brief Review of Domain Adaptation’, arXiv. Accessed: Jun. 12, 2024.
- [7] I. Goodfellow *et al.*, ‘Generative Adversarial Nets’, in *Advances in Neural Information Processing Systems*, Curran Associates, Inc., 2014. Accessed: Dec. 02, 2023.
- [8] J.-Y. Zhu, T. Park, P. Isola, and A. A. Efros, ‘Unpaired Image-to-Image Translation using Cycle-Consistent Adversarial Networks’. arXiv, Aug. 24, 2020. doi: 10.48550/arXiv.1703.10593.
- [9] K. J. Halupka *et al.*, ‘Retinal optical coherence tomography image enhancement via deep learning’, *Biomed Opt Express*, vol. 9, no. 12, pp. 6205–6221, Dec. 2018, doi: 10.1364/BOE.9.006205.
- [10] D. Mahapatra, B. Antony, S. Sedai, and R. Garnavi, ‘Deformable medical image registration using generative adversarial networks’, in *2018 IEEE 15th International Symposium on Biomedical Imaging (ISBI 2018)*, Apr. 2018, pp. 1449–1453. doi: 10.1109/ISBI.2018.8363845.
- [11] J. Kugelman, D. Alonso-Caneiro, S. A. Read, S. J. Vincent, F. K. Chen, and M. J. Collins, ‘Data augmentation for patch-based OCT chorio-retinal segmentation using generative adversarial networks’, *Neural Computing and*
- [12] S. Wang, L. Yu, X. Yang, C.-W. Fu, and P.-A. Heng, ‘Patch-Based Output Space Adversarial Learning for Joint Optic Disc and Cup Segmentation’, *IEEE Transactions on Medical Imaging*, vol. 38, no. 11, pp. 2485–2495, Nov. 2019, doi: 10.1109/TMI.2019.2899910.
- [13] T. K. Yoo *et al.*, ‘Deep learning can generate traditional retinal fundus photographs using ultra-widefield images via generative adversarial networks’, *Comput Methods Programs Biomed*, vol. 197, p. 105761, Dec. 2020, doi: 10.1016/j.cmpb.2020.105761.
- [14] S. K. S. S and K. G. Srinivasa, ‘Artifact Reduction in Fundus Imaging using Cycle Consistent Adversarial Neural Networks’. arXiv, Dec. 25, 2021. doi: 10.48550/arXiv.2112.13264.
- [15] T. Vo and N. Khan, ‘Edge-preserving Domain Adaptation for semantic segmentation of Medical Images’. arXiv, Nov. 18, 2021. doi: 10.48550/arXiv.2111.09847.
- [16] P. Iacono and N. Khan, ‘Structure Preserving Cycle-GAN for Unsupervised Medical Image Domain Adaptation’. arXiv, Apr. 18, 2023. doi: 10.48550/arXiv.2304.09164.
- [17] M. T. Hagos and S. Kant, ‘Transfer Learning based Detection of Diabetic Retinopathy from Small Dataset’. arXiv, May 22, 2019. doi: 10.48550/arXiv.1905.07203.
- [18] S. Gupta, A. Panwar, S. Goel, A. Mittal, R. Nijhawan, and A. K. Singh, ‘Classification of Lesions in Retinal Fundus Images for Diabetic Retinopathy Using Transfer Learning’, in *2019 International Conference on Information Technology (ICIT)*, Dec. 2019, pp. 342–347. doi: 10.1109/ICIT48102.2019.00067.
- [19] ‘Data Analysis’, EyePACS. Accessed: Jun. 15, 2024.
- [20] K. Simonyan and A. Zisserman, ‘Very Deep Convolutional Networks for Large-Scale Image Recognition’. arXiv, Apr. 10, 2015. doi: 10.48550/arXiv.1409.1556.
- [21] ‘Diabetic Retinopathy Grading by a Source-Free Transfer Learning Approach - ScienceDirect’. Accessed: Jun. 14, 2024.
- [22] G. Huang, Z. Liu, L. Van Der Maaten, and K. Q. Weinberger, ‘Densely Connected Convolutional Networks’, in *2017 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, Jul. 2017, pp. 2261–2269. doi: 10.1109/CVPR.2017.243.
- [23] N. Tajbakhsh *et al.*, ‘Convolutional Neural Networks for Medical Image Analysis: Full Training or Fine Tuning?’, *IEEE transactions on medical imaging*, vol. 35, no. 5, pp. 1299–1312, May 2016, doi: 10.1109/TMI.2016.2535302.
- [24] A. Mittal, A. K. Moorthy, and A. C. Bovik, ‘No-Reference Image Quality Assessment in the Spatial Domain’, *IEEE Transactions on Image Processing*, vol. 21, no. 12, pp. 4695–4708, Dec. 2012, doi: 10.1109/TIP.2012.2214050.
- [25] H. Zhao, B. Yang, L. Cao, and H. Li, ‘Data-Driven Enhancement of Blurry Retinal Images via Generative Adversarial Networks’, in *Medical Image Computing and Computer Assisted Intervention – MICCAI 2019*, pp. 75–83. doi: 10.1007/978-3-030-32239-7_9.
- [26] C. Szegedy, V. Vanhoucke, S. Ioffe, J. Shlens, and Z. Wojna, ‘Rethinking the Inception Architecture for Computer Vision’, in *2016 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, Jun. 2016, pp. 2818–2826. doi: 10.1109/CVPR.2016.308.