

VTGAN: Semi-supervised Retinal Image Synthesis and Disease Prediction using Vision Transformers

Sharif Amit Kamran
University of Nevada, Reno
skamran@nevada.unr.edu

Khondker Fariha Hossain
University of Nevada, Reno
khondkerfariyah@nevada.unr.edu

Alireza Tavakkoli
University of Nevada, Reno
tavakkol@unr.edu

Stewart Lee Zuckerbrod
Houston Eye Associates
szuckerbrod@houstoneye.com

Salah A. Baker
University of Nevada, Reno
sabubaker@med.unr.edu

Abstract

In Fluorescein Angiography (FA), an exogenous dye is injected in the bloodstream to image the vascular structure of the retina. The injected dye can cause adverse reactions such as nausea, vomiting, anaphylactic shock, and even death. In contrast, color fundus imaging is a non-invasive technique used for photographing the retina but does not have sufficient fidelity for capturing its vascular structure. The only non-invasive method for capturing retinal vasculature is optical coherence tomography-angiography (OCTA). However, OCTA equipment is quite expensive, and stable imaging is limited to small areas on the retina. In this paper, we propose a novel conditional generative adversarial network (GAN) capable of simultaneously synthesizing FA images from fundus photographs while predicting retinal degeneration. The proposed system has the benefit of addressing the problem of imaging retinal vasculature in a non-invasive manner as well as predicting the existence of retinal abnormalities. We use a semi-supervised approach to train our GAN using multiple weighted losses on different modalities of data. Our experiments validate that the proposed architecture exceeds recent state-of-the-art generative networks for fundus-to-angiography synthesis. Moreover, our vision transformer-based discriminators generalize quite well on out-of-distribution data sets for retinal disease prediction.

1. Introduction

Fluorescein Angiography (FA) combined with retinal funduscopy is a standard tool for diagnosing various retinal vascular abnormalities and degenerative conditions [40]. In Fluorescein Angiography, a fluorescent fluid is injected into the blood stream and becomes visible 8 to 10 minutes after

insertion. The dosage depends on the viscosity of the dye, age of the patient, and the retina vascular structure [39]. In general, the procedure is safe, but there have been reported cases of allergic reactions, ranging from nausea and vomiting to anaphylactic shock and death [38, 33, 34]. The risk associated with FA signifies the need for non-invasive mechanisms to assess retinal vascular structure.

Various automated systems incorporating image processing and machine learning techniques have been proposed for diagnosing retinal abnormalities and degenerative diseases from fundus photographs [16, 15, 45, 38]. Currently, there are no inexpensive imaging alternative for imaging the vascular structure of the retina. Although OCTA [22] is capable of mapping the retinal subspace in 3D, the equipment is expensive and its highest resolutions are only achievable on very small areas. Consequently, it is imperative to develop non-invasive and inexpensive techniques for measuring retinal vasculature to circumvent the risks associated with the existing invasive procedures. The first contribution of this paper is a novel deep learning architecture for producing retinal vascular images (i.e., FA images) from non-invasive fundus photographs.

A major obstacle in using machine learning architectures for ophthalmic applications is the lack of publicly available data – in our work Fluorescein Angiography data [28, 29]. Only a handful of machine learning systems have been proposed to predict disease using FA images. However, these systems are trained and tested on privately-held data set. Pan *et al.* [43] utilized three such pre-trained architectures on 4067 privately curated images for Abnormal and Normal FA prediction. Similarly, Li *et al.* [36] incorporated the encoder of U-Net [46] for classifying different levels of degradation for 3935 privately held Ultra-widefield FA images. To our knowledge, the only two available public data sets have only 59, and 70 FA images in total [17, 18]. The second contribution of this paper is a semi-supervised

approach in conjunction with a vision transformer architecture to address the challenges in training models that learn pathologies from such small amounts of data.

In particular, we introduce VTGAN, a semi-supervised conditional GAN that can simultaneously produce the retinal vascular structure (i.e. FA images) from fundus photographs, while differentiating between healthy and abnormal retina. The proposed architecture incorporates a novel discriminator based on vision transformers for both patch-level adversarial image detection and image-level disease classification. Specifically, we utilize a novel embedding feature loss incorporating patch and positional information to find similarities between angiography and fundus images for better disease prediction. For qualitative assessment, we compare our proposed architecture with recent state-of-the-art fundus-to-angiogram synthesis architectures [27, 26, 54, 31, 9]. For quantitative evaluation, we use Frechet inception Distance (FID) [21] and Kernel Inception Distance (KID) [1] for quantifying image features and measuring structural similarity. To validate the discriminator’s robustness, we use standard metrics for out-of-distribution classification.

2. Related Work

Generative adversarial networks (GAN) have become a staple for image-to-image synthesis [6, 47], inpainting [58, 11] and style transfer [54, 55]. By a combination of multi-scale architectures, these networks can detect and learn fine and coarse features from images [3, 2]. This approach, in turn, can be employed for both conditional [23, 12] and unconditional variants of GAN architectures [5, 57]. Generative networks have also seen success in computer tomography (CT), magnetic resonance imaging (MRI), and X-ray for image segmentation, augmentation and cross-domain information fusion tasks [41, 42, 49, 53]. Recently, GAN models are utilized for synthesizing FA images from fundus photographs [51, 26]. These models employ two generators for coarse and fine image generation. The generators were trained on randomly cropped patches of different scales, while multi-scale discriminators are used to discern local and global features from images. Usually, a discriminator incorporates a linear output for classifying among real or fake (generated) images. In addition, these models utilize the PatchGAN architecture [35] as a discriminator for adversarial and real image classification, by producing the output as a feature map of size $N \times N$. Incorporating these ideas, the state-of-the-art models employ an architecture that can retain global information like the shape of optic-disc, contrast, and local features like venular structures, arteries, and microaneurysm [51, 26]. However, the problem with this approach is that the discriminator works on patch-level information. As a result, the cohesive relationship between global and local features is lost while

generating FA from crops of fundus images.

The advent of Vision Transformers (ViT) have improved the performance of state-of-the-art architectures in image classification tasks. The premise of ViT is that image pixels contain inherent spatial coherencies. Therefore, utilizing an image as a sequence of non-overlapping patches and incorporating them into a transformer can better extract intrinsic and spatial features [52]. Unlike image generative pre-trained models such as iGPT [4], which apply transformers to obtain image-level features, ViT works on patch-level features and incorporates their positions by utilizing embedding layers [30, 4].

Vision Transformers retain the cohesiveness of coarse and fine features by utilizing the position of each patch. By mapping this information for $N \times N$ patches into an $N \times N$ feature-map, we propose a novel architecture based on vision transformers, called VTGAN and illustrated in Fig. 1. Moreover, we extend VTGAN’s capability by adding a multi layer perceptron (MLP) head for classifying Abnormal and Normal FA images in a semi-supervised manner. This architecture is a significant contribution that addresses training models for FA image classification task based on limited amounts of publicly available data. Qualitative and quantitative evaluations demonstrate that VTGAN surpasses other state-of-the-architectures both in terms of synthesis and classification tasks. Additionally, we demonstrate that the produce FA images by VTGAN are such high quality that expert ophthalmologists cannot reliably identify fake image from real ones.

3. Proposed Methodology

This paper proposes a vision-transformer-based generative adversarial network (GAN) consisting of residual, spatial feature fusion, upsampling and downsampling blocks for generators, and transformer encoder blocks for discriminators. We incorporate multiple losses for generating vivid fluorescein angiography images from normal and abnormal fundus photographs for training. In the first section, we elaborate coarse and fine generators in section 3.1. We detail each distinct blocks in sections 3.2, 3.3, and 3.4. We then elaborate on our newly proposed vision-transformer-based discriminators and the interrelationship between the generators and discriminators to establish the overall generative network in section 3.5. Finally, in section 3.6 we discuss the associated loss functions and their weight multipliers for each distinct architecture that forms the proposed model.

3.1. Multi-scale Generators

In order to capture large and fine-scale features to produces realistic vascular images, we combine multi-scale coarse and fine generators. We adopt two generators (G_f and G_c), as visualized in Fig. 1 in our architecture. G_f syn-

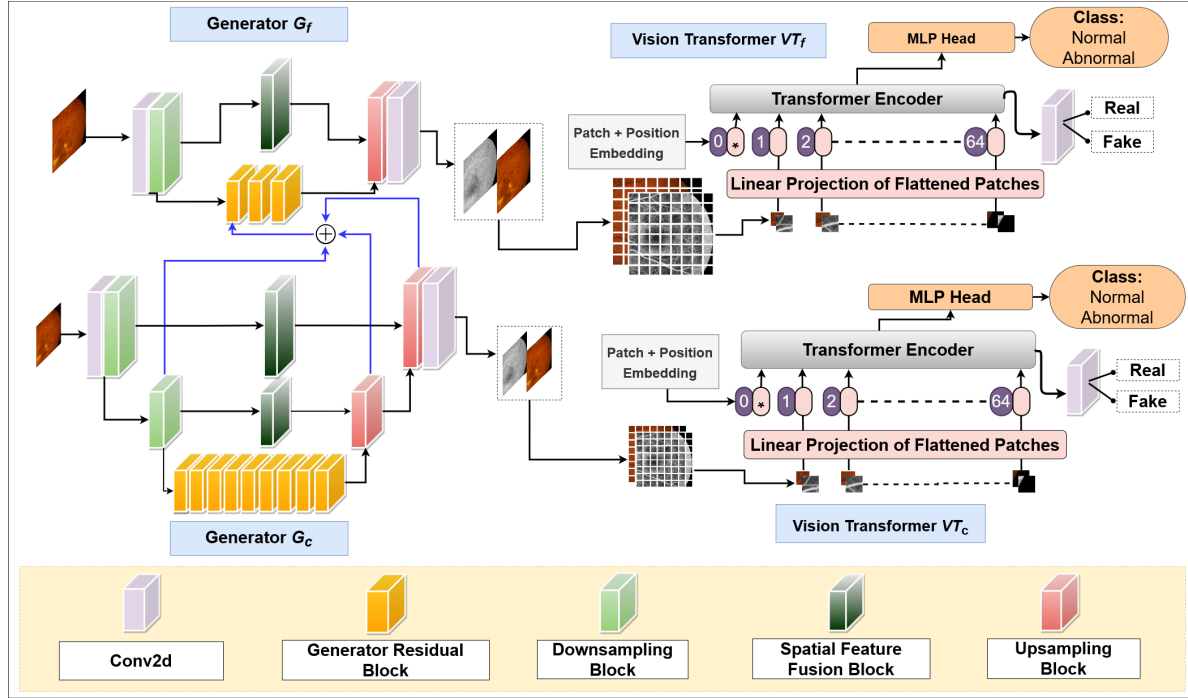


Figure 1. VTGAN consists of Coarse and Fine generators G_f, G_c and Vision Transformers as discriminators VT_f, VT_c . The generators take fundus as input and synthesize angiograms. Whereas the vision transformers take patches of the concatenated fundus and angiograms as input and outputs, i) feature-map for adversarial example detection, and ii) image-level classification for Fluorescein angiograms. The generators consist of Downsampling, Upsampling, Spatial Feature Fusion, Residual blocks and multi-scale feature summation (Blue arrow). The Vision transformers consist of Transformer Encoder, Multi-layer perceptron (MLP) head, and adversarial feature output block.

thesizes local features such as arteries and venules. Conversely, G_c translates global features such as large blood vessels, optic disc, and overall contrast and illumination. The generators consist of multiple downsampling, upsampling, spatial feature fusion, residual blocks, and a multi-scale feature summation block (Blue arrow) between the G_f and G_c . The input and output dimension of G_f are 512×512 , while G_c has half the dimension of G_f (256×256). Furthermore, the G_c generates $256 \times 256 \times 64$ size feature vector, which is element-wise added with one of the intermediate layers of the G_f generator. Both G_c and G_f take fundus images and synthesize FA images. The detailed arrangement is visualized in Fig. 1.

3.2. Downsampling and Upsampling Blocks

We use, as generators, auto-encoders comprising of multiple downsampling and upsampling blocks for feature extraction. A single downsampling block contains a convolution layer, a batch-norm layer [24] and a Leaky-ReLU activation function successively and is given in Fig. 2(i). In contrast, an upsampling block consists of a transposed convolution layer, batch-norm [24], and Leaky-ReLU activation layer consecutively and is illustrated in Fig. 2(ii). We use the downsampling block twice in G_c , followed by nine successive residual identity blocks. Finally, the upsam-

pling blocks are used $2 \times$ again to make the spatial output the same as the input. For G_f , we utilize the downsampling once, and after three consecutive residual blocks, a single upsampling block is employed to get the same spatial output as the input. We use kernel size, $k = 3$, and stride, $s = 2$ for convolution 2D and transposed convolution 2D layers.

3.3. Generator's Residual Blocks

For spatial and depth feature extraction, residual blocks have become the fundamental building blocks for image-to-image translation, inpainting, and style transfer tasks [48, 54, 9, 8, 44]. The original residual block consisted of successive convolutions and a skip connection linking the pre-residual layer with the post-residual layer [19]. Regular convolution operations are computationally inefficient and fail to retain accurate spatial and depth information, unlike separable convolution [10]. Separable convolution incorporates a depth-wise convolution layer followed by a point-wise convolution. Consequently, it obtains and retains depth and spatial features better. We use residual blocks for our generators, as illustrated in Fig. 2(iv). The residual block consists of Reflection padding, Separable Convolution, Batch-Norm, and Leaky-ReLU layers followed by two branches of the same repetitive layers. The main difference

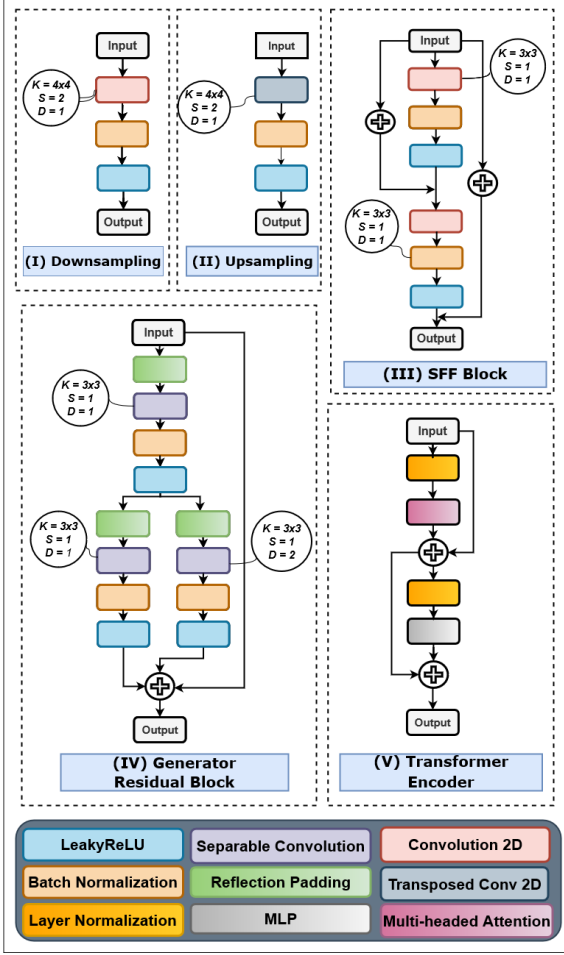


Figure 2. Distinct building blocks of our proposed Generative network comprising of (i) Downsampling, (ii) Upsampling, (c) Spatial Feature Fusion (SFF) block, (d) Generator Residual Block and (e) Transformer Encoder. Here, K = for kernel size, S = stride and D = Dilation rate.

is, one branch consists of a dilation rate of $d = 1$ and the other with a dilation rate, $d = 2$ in the separable convolution layers. We use, kernel size, $k = 3$ and stride, $s = 1$. The output from the two branches and the skip-connection successively element-wise summed in the final output.

3.4. Spatial Feature Fusion Block

In this section, we elaborate on the spatial feature fusion (SFF) block, as illustrated in Fig. 2(iii). The block consists of two residual units with Convolution, Batch-Norm, Leaky-ReLU layer successively. The convolution layer have kernel, $k = 3$, and stride, $s = 1$. Additionally, there are two skip connections, one going from the input and element-wise summed to the first residual unit's output. The next one comes out of the input layer and is added with the last residual unit's output. We use spatial feature fusion block for combining spatial features from the bottom lay-

ers with the topmost layers of the architecture, as visualized in Fig. 1. G_c comprises two SFF blocks that connect each of the two downsampling blocks with the two upsampling blocks successively. In contrast, G_f has only one SFF block between the single downsampling and upsampling block. The reason behind incorporating the SFF block is to extract and retain spatial information that is otherwise lost due to consecutive downsampling and upsampling. As a result, we can combine these features with the learned features of the later layers of the network to get an accurate approximation, as seen in similar GAN architectures [56, 7].

3.5. Vision-transformer as Discriminators

GAN discriminators require adapting to local and global information changes for differentiating real and fake images. To alleviate this inherent problem, we need a heavy architecture with a large number of parameters. In contrast, convolution with a large receptive field can be employed for obtaining multi-scale features but can cause over-fitting on training data. To resolve this problem, we propose a new Vision Transformer-based Markovian discriminator, in the same vein as PatchGAN [35]. We use eight vision transformer encoders consisting of a multi-headed attention layer and multi-layer perceptron (MLP) block. The Layer Normalization layer precedes each block, and residual skip connection is added to the output from the input. The block is visualized in Fig. 2(v). To handle 2D images of 512×512 , we reshape the images into a sequence of flattened 2D patches with resolution 64×64 . By doing so, we end up having 64 patches in total. The Transformer uses a constant latent vector size of $D = 64$ through all of its layers, so we flatten the patches and map to D dimensions with a trainable linear projection. The output of this projection is called the patch embeddings as mentioned in [13]. Position embeddings are added to the patch embeddings to preserve positional information. We use regular learnable 1D position embeddings in a similar manner to Dosovitskiy *et al.* [13]. For multi-headed attention, we use $n = 4$ heads. For MLP blocks, we use two dense layers with features size $h = [128, 64]$, each succeeded by a GeLU activation [20] and a dropout of 0.1. Contrarily, our vision transformer has two outputs, an MLP head, and a Convolutional layer, as illustrated in Fig. 1. The MLP head has two output hidden units for FA image classification (Abnormal and Normal). In contrast, the convolution layer outputs a feature map of 64×64 for classifying each patch in the original image as Real or Fake.

We use two vision transformer-based discriminators that incorporate identical structures but operate at two different scales. We term the two discriminators as, VT_f and VT_c as visualized in Fig. 1. The coarse angiograms and fundus are resized to 256×256 by a factor of 2 using the Lanczos filter [14]. Both discriminators have identical transformer

encoder and output layers (in Fig. 2(v)).

VT_c dictates G_c to extract more global features such as the fovea, brightness, optic disc, and color sensitivity. Contrarily, the VT_f steers the fine generator, G_f to synthesize more accurate fine features such as small vasculature, venules, exudates, arteries. Consequently, we fuse learnable elements from both generators while training them with their paired vision transformer-based discriminators.

3.6. Adversarial Cost Functions

Our whole network's objective function can be contrived as Eq. 1. So it's a problem spanning multiple distinct objectives for the generator and discriminator networks, where our target is to maximize VT_f and VT_c 's loss while minimizing the G_f and G_c 's loss.

$$\min_{G_f, G_c} \max_{VT_f, VT_c} \mathcal{L}_{adv}(G_f, G_c, VT_f, VT_c) \quad (1)$$

We employ Hinge-Loss [56, 37] to train our model in an adversarial setting, as illustrated in Eq. 2 and Eq. 3. All the fundus images and their corresponding angiogram pairs are normalized to $[-1, 1]$, and we employ \tanh as the output activation for the adversarial feature-map. We multiple $\lambda_{adv}(G)$ with a weight multiplier $\mathcal{L}_{adv}(G)$ and add $\lambda_{adv}(D)$ to the output as given in Eq. 4. For classification we use *softmax* activation after the MLP head output. In Eq. 5 the categorical cross-entropy loss is given where y is the real class and \hat{y} is the predicted class.

$$\mathcal{L}_{adv}(VT) = -\mathbb{E}_{x,y} [\min(0, -1 + VT(x, y))] - \mathbb{E}_x [\min(0, -1 - VT(x, G(x)))] \quad (2)$$

$$\mathcal{L}_{adv}(G) = -\mathbb{E}_{x,y} [(VT(G(x), y))] \quad (3)$$

$$\mathcal{L}_{adv}(G, VT) = \mathcal{L}_{adv}(VT) + \lambda_{adv}(\mathcal{L}_{adv}(G)) \quad (4)$$

$$\mathcal{L}_{cce}(VT) = \mathbb{E}_{y, \hat{y}} \left[- \sum_{i=1}^k y_i \log \hat{y}_i \right] \quad (5)$$

Eq. 2 and Eq. 3 demonstrate VT_f and VT_c being trained on the real fundus (x) and FA (y) images and then real fundus (x) and fake FA ($G(x)$) images in succession. We batch-wise train the vision transformers VT_f and VT_c for a couple of repetitions on images sampled randomly. Following that, G_c and G_f are trained on a batch of random images while freezing the weights the vision transformers, VT_c and VT_f .

We utilize MSE loss and VGG-19 based perceptual loss [25] for G_c and G_f as given in Eq. 6 and Eq. 7. We ensure the synthesized images retain more realistic color, contrast, and vascular structure by employing these losses. For our discriminators, VT_c and VT_f we incorporate a novel embedding feature loss as provided in Eq. 8.

$$\mathcal{L}_{MSE}(G) = \mathbb{E}_{x,y} \|G(x) - y\|^2 \quad (6)$$

$$\mathcal{L}_{perc}(G) = \mathbb{E}_{x,y} \sum_{i=1}^k \frac{1}{P} \|F_{VGG19}^i(y) - F_{VGG19}^i(G(x))\| \quad (7)$$

$$\mathcal{L}_{ef}(G, VT) = \mathbb{E}_{x,y} \sum_{i=1}^k \frac{1}{Q} \|VT_{em}^i(x, y) - VT_{em}^i(x, G(x))\| \quad (8)$$

For synthesizing FA at different scales, we utilize the mean-squared-error loss \mathcal{L}_{MSE} as shown in Eq. 6. The difference between the real FA, y , and the synthesized FA, $G(x)$ is calculated using this loss. In contrast, the perceptual loss, \mathcal{L}_{perc} is employed for findings similarities between real and fake FA as given in Eq. 7. Both of these images are pushed successively in VGG19 [50] architecture, and then the resultant feature is subtracted. Finally, the novel embedding feature loss is calculated by obtaining positional and patch features from the transformer encoder layers of VT by inserting the real and synthesized FA successively, as shown in Eq. 8. Here, P and Q stand for the number of features extracted from VGG19 layers and the embedding layers of the transformer-encoder consecutively.

We combine Eq. 4, 5, 6, 7 and 8 to configure our ultimate cost function as provided in Eq. 9.

$$\min_{G_f, G_c} \left(\max_{VT_f, VT_c} (\mathcal{L}_{adv}(G_f, G_c, VT_f, VT_c)) + \lambda_{MSE} [\mathcal{L}_{MSE}(G_f, G_c)] + \lambda_{ef} [\mathcal{L}_{ef}(G_f, G_c, VT_f, VT_c)] + \lambda_{perc} [\mathcal{L}_{perc}(G_f, G_c)] + \lambda_{cce} [\mathcal{L}_{cce}(VT_f, VT_c)] \right) \quad (9)$$

λ_{adv} , λ_{ef} , λ_{perc} , λ_{MSE} , and λ_{cce} implies weights of the losses, and prioritizes a specific loss during training the networks. In our generative network, we emphasize more on $\mathcal{L}_{adv}(G)$, \mathcal{L}_{MSE} , \mathcal{L}_{perc} , \mathcal{L}_{cce} , and thus we pick larger λ values for these.

4. Experiments

The following section outlines our proposed architectures experimentations and evaluations on quantitative & qualitative metrics. First, we discuss the pre-processing pipeline for our dataset in Sec. 4.1. Next, we detail our hyper-parameter selection and tuning in Sec. 4.2. Moreover, we compare our model with other fundus-to-angio synthesis models with various qualitative metrics in Sec. 4.3. Lastly, in Sec. 4.4 and Sec. 4.5, we examine the expert quantifications and the performance on the out-of-distribution dataset.

4.1. Dataset

For our experimentations, we utilize the dataset provided by [17]. The dataset comprises thirty healthy and twenty-nine unhealthy sets of fundus and fluorescein angiograms (FA). Each image pair is collected from individual patients. After close observation, we found seventeen sets that are

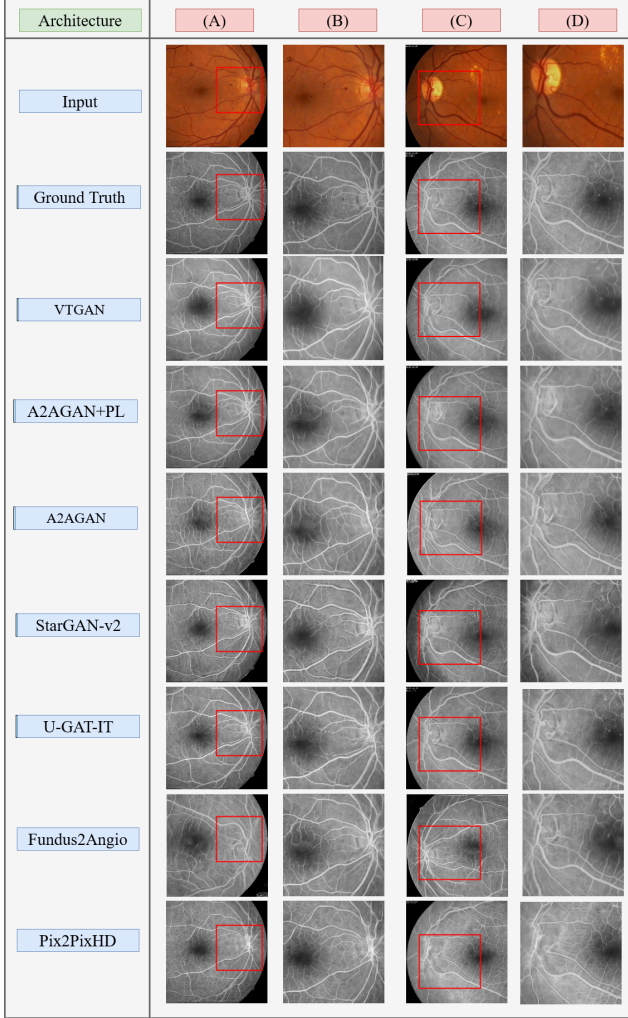


Figure 3. Relative differences in Fluroscein Angiograms (FA) synthesized by various generative networks. Columns (A) and (C) depict real fundus, real FA, and synthesized FA images. In contrast, column (B) and (D) displays zoomed-in red rectangle boxes from Column (A) & (C), to show similar local vessels and venular arrangement.

either perfectly or closely aligned. From the original images of resolution 576×720 , we extract 50 from each set with an overlapping crop size of 512×512 . For the image synthesis training task, we extract 850 images. The fundus photographs are in 3 channel RGB form, and FA images are in 1 channel Gray-scale form. As the dataset is categorized into Abnormal and Normal classes, we use this annotation for our supervised classification training. Out of seventeen images, ten are Abnormal, and seven are Normal patients. Due to cropping, we end up having 500 for abnormal and 350 for normal images. We use data augmentation to increase the sample size the same as the abnormal one. To test the models, we crop four overlapping quadrants of each image pair and generate fifty-six images in total along with

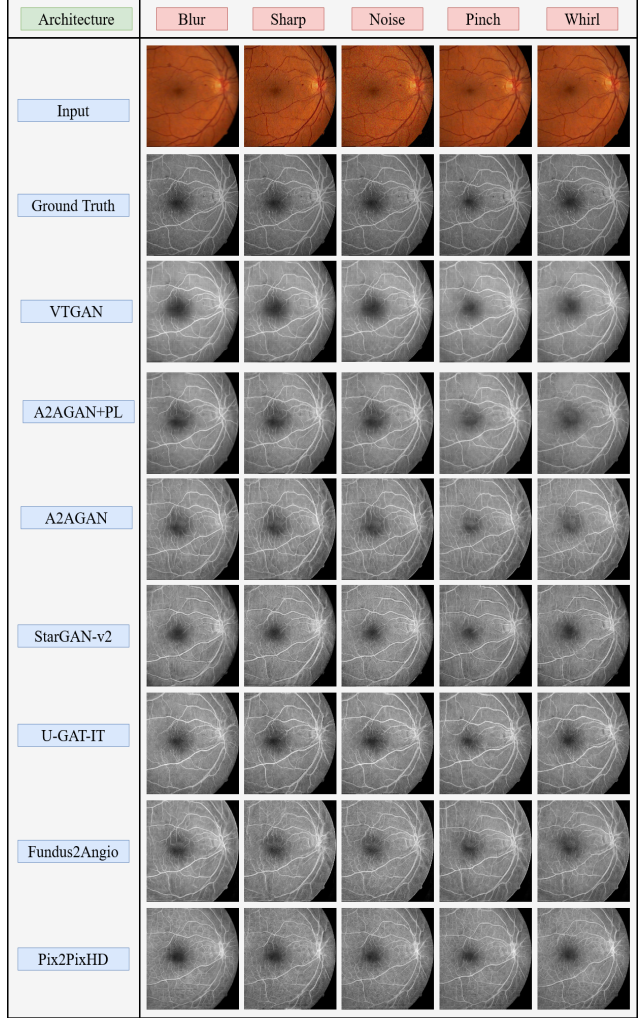


Figure 4. **1st & 2nd rows:** Fundus and Angiogram pairs altered using spatial and radial distortions. **3rd-8th rows:** Angiograms generated from distorted fundus imitating imaging errors and distinct biological markers.

their associated labels.

4.2. Hyper-parameter tuning

We utilized hinge loss [56, 37] for training our models in an adversarial setting. We chose $\lambda_{adv} = 10$ (Eq. 4) and $\lambda_{perc} = 10$, $\lambda_{ef} = 1$, $\lambda_{MSE} = 10$, $\lambda_{cce} = 10$ (Eq. 9). For optimizer we used Adam [32], with learning rate $\alpha = 0.0002$, $\beta_1 = 0.5$ and $\beta_2 = 0.999$. The batch size was, $b = 2$ and we trained for 200 epochs for 50 hours with NVIDIA GPUs. Code Link : github.com/SharifAmit/VTGAN.

4.3. Qualitative Assessments

For evaluating our architecture's performance, we used 14 test samples and cropped four quadrants of the image with a resolution of 512×512 . We conducted two experiments for estimating the accurate visual representation i)

Table 1. Test results for different architectures

Fr�chet Inception Distance (FID)						
Architecture	Orig.	Noise	Blur	Sharp	Whirl	Pinch
VTGAN	17.3	18.7 (1.4�)	24.1 (6.8�)	24.5 (7.2�)	28.4 (11.1�)	22.3 (5.0�)
A2AGAN w/ PL ¹ [26]	24.6	21.6 (3.0�)	30.0 (5.4�)	25.6 (1.0�)	40.0 (15.4�)	24.9 (0.3�)
A2AGAN [26]	20.7	20.8 (0.1�)	23.5 (2.8�)	24.9 (4.2�)	27.8 (7.1�)	19.5 (1.2�)
StarGAN-v2 [9]	27.7	35.1 (7.4�)	32.6 (4.9�)	27.4 (0.3�)	32.7 (5.0�)	26.7 (1.0�)
U-GAT-IT [31]	24.5	26.0 (1.5�)	30.4 (5.9�)	26.8 (2.3�)	33.0 (9.5�)	29.1 (4.6�)
Fundus2Angio [27]	30.3	41.5 (11.2�)	32.3 (2.0�)	34.3 (4.0�)	38.2 (7.9�)	33.1 (2.8�)
Pix2PixHD [54]	42.8	53.0 (10.2�)	43.7 (1.1�)	47.5 (4.7�)	45.9 (3.1�)	39.2 (3.6�)
Kernel Inception Distance (KID)						
Architecture	Orig.	Noise	Blur	Sharp	Whirl	Pinch
VTGAN	0.00053	0.04953	0.00205	0.04927	0.04815	0.04542
A2AGAN w/ PL [26]	0.00087	0.05045	0.00235	0.05162	0.05390	0.04575
A2AGAN [26]	0.00392	0.05390	0.00505	0.05301	0.05657	0.05341
StarGAN-v2 [9]	0.00118	0.05274	0.00235	0.05331	0.05539	0.05271
U-GAT-IT [31]	0.00131	0.05610	0.00278	0.05533	0.05815	0.05719
Fundus2Angio [27]	0.00184	0.05328	0.00272	0.05267	0.05278	0.04985
Pix2PixHD [54]	0.00258	0.05613	0.00254	0.05788	0.06029	0.05838

¹ PL = Perceptual Loss² FID: Lower is better; KID: Lower is better

without transformation and ii) with spatial and radial transformations. By doing so, we measured the network’s ability to adjust to structural changes due to imaging error and patient eye movements.

For our first experiment, we compare our network with other state-of-the-art image-to-image translation architectures. A side-by-side comparison of the results is visualized in Fig. 3. Column, A & C in Fig. 3, display the global feature differences while column B & D are zoomed-in local vasculature changes. By its looks, our model produces vivid and convincing results compared to other architectures. Attention2Angio (A2AGAN) and U-GAT-IT also yield impressive results. However, if observed for the closed-up versions in columns B & D, we can witness that the optic disc contains fewer blood vessels. StarGAN-v2 and Pix2PixHD also fail to generate rich arteries, exudates, and vasculature.

In the second set of experiments, we applied three transformations and two distortions on the fundus images: 1) Blur to represent severe cataracts, 2) Sharpening to represent dilated pupils, 3) Noise to sensor impedance during fundoscopy, 4) Pinch to visualize squinting as a vascular change, and, 5) Whirl, for distortions caused by increased intraocular pressure (IOP). We can see in Fig. 4 a side-by-side comparison of different architecture predictions on these transformed images. As observed from the results, VTGAN synthesizes images very similar to the ground-truth (GT) and performs robustly to preserve the vascular structures.

For **Blurred** fundus images, VTGAN is less impacted by the transformation compared to other state-of-the-art architecture, as seen in (row 4 to 9 of column 1) of Fig. 4.

The vascular structures are better retained as opposed to U-GAT-IT and Fundus2Angio. For **Sharpened** fundus, the angiogram generated by StarGAN-v2 and A2AGAN (row 4 to 6 of column 2) exhibits small artifacts around the blood vessels not present in our case. For **Noisy** images, our result is unaffected by this pixel-level alteration. However, all other state-of-the-art models (row 4 to 9 of column 3) fail to synthesize tiny and slim venular branches.

For **Pinch** and **Whirl**, our experimental result shows the versatility and reproducibility of VTGAN for retaining vascular structure and is illustrated in Fig. 4 (row 3 of column 4 and 5). Compared to ours, only A2AGAN and U-GAT-IT reserves the flattening condition and manifestation of vascular changes but loses the overall smoothness in the process (row 5 to 7 of columns 4 and 5). In Fig. 4 VTGAN encodes the vascular feature information and is much less affected by both kinds of warping. The other architectures failed to generate microvessels due to IOP or vitreous variations, as can be seen in Fig. 4. Consequently, For all kinds of transformation and contortion, VTGAN surpasses existing state-of-the-art image-to-image translation models.

4.4. Quantitative Assessments

For quantitative evaluation, we performed two experiments. In the first experiment, we use the Fr chet inception distance (FID) [21], and Kernel Inception distance (KID) [1] which has been previously employed for measuring similar image-to-image translation GANs [9, 31]. We computed the FID and KID scores for different architectures on the generated FA image and original angiogram, including the five spatial and radial transformations. The results are

Table 2. Results of Qualitative with Undisclosed Portion of Fake/Real Experiment

Architecture		Results		Average		
		Correct	Incorrect	Missed ¹	Found ¹	Precision ²
VTGAN	Fake	6%	94%	57%	43%	45.9%
	Real	80%	20%			

¹ Missed higher is better; Found lower is better

² Precision Lower is better

given in Table. 1. A lower FID and KID score means the synthesized images are more close to the real angiogram.

From Table. 1, out of our three networks, the best KID is achieved for VTGAN. And it reaches the lowest scores out of all other architecture, for both with and without distortions. For FID, our model achieves the lowest score for three out of five types of distortions. Attention2Angio scores lower FID for Blur and Whirl transformation.

In the second experiment, we assess the generated angiogram’s quality by asking two expert ophthalmologists to identify them. We use a balanced set of 50 images, 25 real and 25 fake. We then shuffle the data before the expert evaluation. For this experiment, experts did not know the exact number of fake and real images. By not disclosing this, we tried to estimate the following criterion: 1) Correct fake and real angiograms found by the experts, where lower equates to better, 2) Incorrect fake and real angiograms missed by the experts, where higher equates to better, and 3) The average precision of the experts for identifying fake angiograms, where lower equates to better. Table 2 illustrates the detailed result.

As it can be seen from Table 2, experts assigned 94% of the fake angiograms as real, for images synthesized by our models. The result also shows that experts had trouble identifying fake images, while they easily identified real angiograms with 80% certainty. On average, the experts misclassified 57% of all images produced by VTGAN. The average precision diagnosis of the experts are 45.9%. Consequently, our model successfully fools the experts to identify fake angios as real.

4.5. Disease Classification

For our subsequent experimentation, we test our visual transformer’s accuracy on in-distribution and out-of-distribution classification tasks. We use data provided by Hajeb *et al.* [17]. We crop four quadrants from the 14 Fundus and Angiogram pairs because of the shortage of data. Out of 56 images, 20 are for Abnormal, and 36 are for Normal classes. We name this test set as in-distribution and

Table 3. Test Accuracy on in-distribution Abnormal/Normal Angiograms

Accuracy	Sensitivity	Specificity
85.7	83.3	90.0

Table 4. Test Accuracy on out-of-distribution Abnormal/Normal Angiograms

Distortion	Accuracy	Sensitivity	Specificity
Blur	78.6 (7.1 ↓)	72.2 (11.1 ↓)	90.0 (-)
Sharp	76.7 (8.0 ↓)	69.4 (13.9 ↓)	90.0 (-)
Noise	78.6 (7.1 ↓)	72.2 (11.1 ↓)	90.0 (-)
Pinch	78.6 (7.1 ↓)	72.2 (11.1 ↓)	90.0 (-)
Whirl	78.6 (7.1 ↓)	72.2 (11.1 ↓)	90.0 (-)

measure the performance using three standard metrics: Accuracy, Sensitivity, and Specificity. The result is provided in Table. 3, and our vision transformer-based model scores 85.7%, 83.3%, 90% for accuracy, sensitivity, and specificity successively.

We use spatial and radial transformation on test images for out-of-distribution evaluation. The model’s performance is illustrated in Table. 4. As can be seen from the table, for Blur, Noise, Pinch, and Whirl transformations, the Accuracy and Sensitivity decreased to 78.6%, 72.2%, consecutively. Compared to the in-distribution data, the decrease is 7.1%, 11.1% for Accuracy and Sensitivity. For Sharp transformation, the accuracy and sensitivity are worse than the other distortions. They are 76.7%, 69.4% successively. An interesting revelation is the specificity has no effect due to these distortions. The specificity is firm 90.0% for all five distortions. A higher specificity signifies our model accurately predicts Abnormal classes better compared to Normal ones. This is significant, as we want to identify patients with degenerative conditions compared to predicting false positives for healthy patients.

5. Conclusion

In this paper, we proposed a new fundus-to-angiogram translation architecture called VTGAN. The architecture generates realistic angiograms from fundus images without any expert intervention. Additionally, we provided results for its robustness and adaptability conditioned upon radial and spatial transformations, which imitate biological markers seen in real fundus images. We believe the proposed network can be incorporated in the wild to generate precise FA images of patients developing disease overtime. It can be a complimentary disease progression monitoring system for predicting the development of diseases in vivo. We hope to extend this work to other areas of ophthalmological image modalities.

Acknowledgement

This material is based upon work supported by the National Aeronautics and Space Administration under Grant No. 80NSSC20K1831 issued through the Human Research Program (Human Exploration and Operations Mission Directorate).

References

- [1] Mikołaj Bińkowski, Dougal J Sutherland, Michael Arbel, and Arthur Gretton. Demystifying mmd gans. *arXiv preprint arXiv:1801.01401*, 2018.
- [2] Matthew Brown, David G Lowe, et al. Recognising panoramas. In *ICCV*, volume 3, page 1218, 2003.
- [3] Peter Burt and Edward Adelson. The laplacian pyramid as a compact image code. *IEEE Transactions on communications*, 31(4):532–540, 1983.
- [4] Mark Chen, Alec Radford, Rewon Child, Jeffrey Wu, Heewoo Jun, David Luan, and Ilya Sutskever. Generative pre-training from pixels. In *International Conference on Machine Learning*, pages 1691–1703. PMLR, 2020.
- [5] Qifeng Chen and Vladlen Koltun. Photographic image synthesis with cascaded refinement networks. In *Proceedings of the IEEE international conference on computer vision*, pages 1511–1520, 2017.
- [6] Wengling Chen and James Hays. Sketchygan: Towards diverse and realistic sketch to image synthesis. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 9416–9425, 2018.
- [7] Xinyuan Chen, Chang Xu, Xiaokang Yang, and Dacheng Tao. Attention-gan for object transfiguration in wild images. In *Proceedings of the European Conference on Computer Vision (ECCV)*, pages 164–180, 2018.
- [8] Yunjey Choi, Minje Choi, Munyoung Kim, Jung-Woo Ha, Sunghun Kim, and Jaegul Choo. Stargan: Unified generative adversarial networks for multi-domain image-to-image translation. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 8789–8797, 2018.
- [9] Yunjey Choi, Youngjung Uh, Jaejun Yoo, and Jung-Woo Ha. Stargan v2: Diverse image synthesis for multiple domains. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 8188–8197, 2020.
- [10] François Chollet. Xception: Deep learning with depthwise separable convolutions. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 1251–1258, 2017.
- [11] Tali Dekel, Chuhan Gan, Dilip Krishnan, Ce Liu, and William T Freeman. Sparse, smart contours to represent and edit images. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 3511–3520, 2018.
- [12] Emily L Denton, Soumith Chintala, Rob Fergus, et al. Deep generative image models using a laplacian pyramid of adversarial networks. In *Advances in neural information processing systems*, pages 1486–1494, 2015.
- [13] Alexey Dosovitskiy, Lucas Beyer, Alexander Kolesnikov, Dirk Weissenborn, Xiaohua Zhai, Thomas Unterthiner, Mostafa Dehghani, Matthias Minderer, Georg Heigold, Sylvain Gelly, et al. An image is worth 16x16 words: Transformers for image recognition at scale. *arXiv preprint arXiv:2010.11929*, 2020.
- [14] Claude E Duchon. Lanczos filtering in one and two dimensions. *Journal of applied meteorology*, 18(8):1016–1022, 1979.
- [15] Huazhu Fu, Jun Cheng, Yanwu Xu, Changqing Zhang, Damon Wing Kee Wong, Jiang Liu, and Xiaochun Cao. Disc-aware ensemble network for glaucoma screening from fundus image. *IEEE transactions on medical imaging*, 37(11):2493–2501, 2018.
- [16] Nikita Gurudath, Mehmet Celenk, and H Bryan Riley. Machine learning identification of diabetic retinopathy from fundus images. In *2014 IEEE Signal Processing in Medicine and Biology Symposium (SPMB)*, pages 1–7. IEEE, 2014.
- [17] Shirin Hajeb Mohammad Alipour, Hossein Rabbani, and Mohammad Reza Akhlaghi. Diabetic retinopathy grading by digital curvelet transform. *Computational and mathematical methods in medicine*, 2012, 2012.
- [18] Shirin Hajeb Mohammad Alipour, Hossein Rabbani, and Mohammadreza Akhlaghi. A new combined method based on curvelet transform and morphological operators for automatic detection of foveal avascular zone. *Signal, Image and Video Processing*, 8(2):205–222, 2014.
- [19] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 770–778, 2016.
- [20] Dan Hendrycks and Kevin Gimpel. Gaussian error linear units (gelus). *arXiv preprint arXiv:1606.08415*, 2016.
- [21] Martin Heusel, Hubert Ramsauer, Thomas Unterthiner, Bernhard Nessler, and Sepp Hochreiter. Gans trained by a two time-scale update rule converge to a local nash equilibrium. In *Advances in neural information processing systems*, pages 6626–6637, 2017.
- [22] David Huang, Eric A Swanson, Charles P Lin, Joel S Schuman, William G Stinson, Warren Chang, Michael R Hee, Thomas Flotte, Kenton Gregory, Carmen A Puliafito, et al. Optical coherence tomography. *science*, 254(5035):1178–1181, 1991.
- [23] Xun Huang, Yixuan Li, Omid Poursaeed, John Hopcroft, and Serge Belongie. Stacked generative adversarial networks. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 5077–5086, 2017.
- [24] Sergey Ioffe and Christian Szegedy. Batch normalization: Accelerating deep network training by reducing internal covariate shift. *arXiv preprint arXiv:1502.03167*, 2015.
- [25] Justin Johnson, Alexandre Alahi, and Li Fei-Fei. Perceptual losses for real-time style transfer and super-resolution. In *European conference on computer vision*, pages 694–711. Springer, 2016.
- [26] Sharif Amit Kamran, Khondker Fariha Hossain, Alireza Tavakkoli, and Stewart Lee Zuckerbrod. Attention2angiogan: Synthesizing fluorescein angiography from retinal fundus images using generative adversarial networks. *arXiv preprint arXiv:2007.09191*, 2020.
- [27] Sharif Amit Kamran, Khondker Fariha Hossain, Alireza Tavakkoli, and Stewart Lee Zuckerbrod. Fundus2angio: A novel conditional gan architecture for generating fluorescein angiography images from retinal fundus photography. *arXiv preprint arXiv:2005.05267*, 2020.
- [28] Sharif Amit Kamran, Sourajit Saha, Ali Shihab Sabbir, and Alireza Tavakkoli. Optic-net: A novel convolutional neu-

- ral network for diagnosis of retinal diseases from optical tomography images. In *2019 18th IEEE International Conference On Machine Learning And Applications (ICMLA)*, pages 964–971. IEEE, 2019.
- [29] Sharif Amit Kamran, Alireza Tavakkoli, and Stewart Lee Zuckerbrod. Improving robustness using joint attention network for detecting retinal degeneration from optical coherence tomography images. In *2020 IEEE International Conference On Image Processing (ICIP)*, pages 2476–2480. IEEE, 2020.
- [30] Salman Khan, Muzammal Naseer, Munawar Hayat, Syed Waqas Zamir, Fahad Shahbaz Khan, and Mubarak Shah. Transformers in vision: A survey. *arXiv preprint arXiv:2101.01169*, 2021.
- [31] Junho Kim, Minjae Kim, Hyeonwoo Kang, and Kwanghee Lee. U-gat-it: unsupervised generative attentional networks with adaptive layer-instance normalization for image-to-image translation. *arXiv preprint arXiv:1907.10830*, 2019.
- [32] Diederik P Kingma and Jimmy Ba. Adam: A method for stochastic optimization. *arXiv preprint arXiv:1412.6980*, 2014.
- [33] Anthony SL Kwan, Chris Barry, Ian L McAllister, and Ian Constable. Fluorescein angiography and adverse drug reactions revisited: the lions eye experience. *Clinical & experimental ophthalmology*, 34(1):33–38, 2006.
- [34] Kris A Kwiterovich, Maureen G Maguire, Robert P Murphy, Andrew P Schachar, Neil M Bressler, Susan B Bressler, and Stuart L Fine. Frequency of adverse systemic reactions after fluorescein angiography: results of a prospective study. *Ophthalmology*, 98(7):1139–1142, 1991.
- [35] Chuan Li and Michael Wand. Precomputed real-time texture synthesis with markovian generative adversarial networks. In *European conference on computer vision*, pages 702–716. Springer, 2016.
- [36] Henry H Li, Joseph R Abraham, Duriye Damla Sevgi, Sunil K Srivastava, Jenna M Hach, Jon Whitney, Amit Vasanthi, Jamie L Reese, and Justis P Ehlers. Automated quality assessment and image selection of ultra-widefield fluorescein angiography images through deep learning. *Translational Vision Science & Technology*, 9(2):52–52, 2020.
- [37] Jae Hyun Lim and Jong Chul Ye. Geometric gan. *arXiv preprint arXiv:1705.02894*, 2017.
- [38] Rodrigo Pessoa Cavalcanti Lira, Cleriston Lucena de Andrade Oliveira, Marta Virgínia Ribeiro Brito Marques, Elaine Rocha Silva, and Cristiano de Carvalho Pessoa. Adverse reactions of fluorescein angiography: a prospective study. *Arquivos brasileiros de oftalmologia*, 70(4):615–618, 2007.
- [39] NARESH Mandava, ELIAS Reichel, D Guyer, et al. Fluorescein and icg angiography. *St Louis: Mosby*, 106:800–808, 2004.
- [40] Viola Stella Mary, Elijah Blessing Rajsingh, and Ganesh R Naik. Retinal fundus image analysis for diagnosis of glaucoma: a comprehensive survey. *IEEE Access*, 2016.
- [41] Dong Nie, Roger Trullo, Jun Lian, Caroline Petitjean, Su Ruan, Qian Wang, and Dinggang Shen. Medical image synthesis with context-aware generative adversarial networks. In *International conference on medical image computing and computer-assisted intervention*, pages 417–425. Springer, 2017.
- [42] Dong Nie, Roger Trullo, Jun Lian, Li Wang, Caroline Petitjean, Su Ruan, Qian Wang, and Dinggang Shen. Medical image synthesis with deep convolutional adversarial networks. *IEEE Transactions on Biomedical Engineering*, 65(12):2720–2730, 2018.
- [43] Xiangji Pan, Kai Jin, Jing Cao, Zhifang Liu, Jian Wu, Kun You, Yifei Lu, Yufeng Xu, Zhaoan Su, Jiekai Jiang, et al. Multi-label classification of retinal lesions in diabetic retinopathy for automatic analysis of fundus fluorescein angiography based on deep learning. *Graefes’s Archive for Clinical and Experimental Ophthalmology*, 258(4):779–785, 2020.
- [44] Taesung Park, Ming-Yu Liu, Ting-Chun Wang, and Jun-Yan Zhu. Semantic image synthesis with spatially-adaptive normalization. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 2337–2346, 2019.
- [45] Ryan Poplin, Avinash V Varadarajan, Katy Blumer, Yun Liu, Michael V McConnell, Greg S Corrado, Lily Peng, and Dale R Webster. Prediction of cardiovascular risk factors from retinal fundus photographs via deep learning. *Nature Biomedical Engineering*, 2(3):158, 2018.
- [46] Olaf Ronneberger, Philipp Fischer, and Thomas Brox. U-net: Convolutional networks for biomedical image segmentation. In *International Conference on Medical image computing and computer-assisted intervention*, pages 234–241. Springer, 2015.
- [47] Patson Sangkloy, Jingwan Lu, Chen Fang, Fisher Yu, and James Hays. Scribbler: Controlling deep image synthesis with sketch and color. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 5400–5409, 2017.
- [48] Tamar Rott Shaham, Tali Dekel, and Tomer Michaeli. Singan: Learning a generative model from a single natural image. In *Proceedings of the IEEE International Conference on Computer Vision*, pages 4570–4580, 2019.
- [49] Hoo-Chang Shin, Neil A Tenenholz, Jameson K Rogers, Christopher G Schwarz, Matthew L Senjem, Jeffrey L Gunter, Katherine P Andriole, and Mark Michalski. Medical image synthesis for data augmentation and anonymization using generative adversarial networks. In *International workshop on simulation and synthesis in medical imaging*, pages 1–11. Springer, 2018.
- [50] Karen Simonyan and Andrew Zisserman. Very deep convolutional networks for large-scale image recognition. *arXiv preprint arXiv:1409.1556*, 2014.
- [51] Alireza Tavakkoli, Sharif Amit Kamran, Khondker Fariha Hossain, and Stewart Lee Zuckerbrod. A novel deep learning conditional generative adversarial network for producing angiography images from retinal fundus photographs. *Scientific Reports*, 10(1):1–15, 2020.
- [52] Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Lukasz Kaiser, and Illia Polosukhin. Attention is all you need. In *NIPS*, 2017.

- [53] Abdul Waheed, Muskan Goyal, Deepak Gupta, Ashish Khanna, Fadi Al-Turjman, and Plácido Rogerio Pinheiro. Covidgan: data augmentation using auxiliary classifier gan for improved covid-19 detection. *Ieee Access*, 8:91916–91923, 2020.
- [54] Ting-Chun Wang, Ming-Yu Liu, Jun-Yan Zhu, Andrew Tao, Jan Kautz, and Bryan Catanzaro. High-resolution image synthesis and semantic manipulation with conditional gans. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 8798–8807, 2018.
- [55] Wenqi Xian, Patsorn Sangkloy, Varun Agrawal, Amit Raj, Jingwan Lu, Chen Fang, Fisher Yu, and James Hays. Texturegan: Controlling deep image synthesis with texture patches. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 8456–8465, 2018.
- [56] Han Zhang, Ian Goodfellow, Dimitris Metaxas, and Augustus Odena. Self-attention generative adversarial networks. In *International Conference on Machine Learning*, pages 7354–7363, 2019.
- [57] Han Zhang, Tao Xu, Hongsheng Li, Shaoting Zhang, Xiaogang Wang, Xiaolei Huang, and Dimitris N Metaxas. Stackgan: Text to photo-realistic image synthesis with stacked generative adversarial networks. In *Proceedings of the IEEE international conference on computer vision*, pages 5907–5915, 2017.
- [58] Jun-Yan Zhu, Philipp Krähenbühl, Eli Shechtman, and Alexei A Efros. Generative visual manipulation on the natural image manifold. In *European Conference on Computer Vision*, pages 597–613. Springer, 2016.