Protecting Privacy while Improving Choroid Layer Segmentation in OCT Images: A GAN-based Image Synthesis Approach

Kiran Kumar Vupparaboina [1], Mohammed Nasar Ibrahim [2], Amrish Selvam [1], Shiva Vaishnavi Kurakula [1], Abdul Rasheed Mohammed [1], Natasha Mayer [1], José Alain Sahel [1], Jay Chhablani [1], and Sandeep Chandra Bollepalli [1]

[1]Affiliation not available
[2]University of Pittsburgh School of Medicine

December 7, 2023

Abstract

THIS WORK HAS BEEN SUBMITTED TO THE MACHINE LEARNING: SCIENCE AND TECHNOLOGY FOR POSSIBLE PUBLICATION. COPYRIGHT MAY BE TRANSFERRED WITHOUT NOTICE, AFTER WHICH THIS VERSION MAY NO LONGER BE ACCESSIBLE.

The choroid, positioned behind the retina, nourishes the retina by supplying oxygen and nutrients. Choroidal structural changes are associated with severe vision-threatening conditions including age-related macular degeneration (AMD) and central serous chorioretinopathy (CSCR). Optical Coherence Tomography (OCT) imaging enables the visualization of choroidal changes, and clinicians rely on quantifying choroidal biomarkers through segmentation of the choroid layer in OCT scans for precise diagnosis and disease management. Accordingly, various attempts are made at automated choroid layer segmentation, however, their practicality is constrained by the limited and biased nature of training data. Privacy regulations hinder data aggregation, and supervised machine learning requires substantial annotated data. To tackle this, we propose an innovative image synthesis approach using generative adversarial networks (GANs). It involves a three-step process: generation of choroid-labeled B-scans using a standard GAN architecture, the transformation of these scans to unlabeled B-scans via Pix2Pix-GAN, and the training of a Pix2Pix-GAN choroid segmentation model using the synthesized data. To demonstrate the generalizability and efficacy, we evaluated the proposed choroid segmentation algorithm on the real B-scans from two different OCT imaging devices: enhanced depth imaging (EDI) and swept-source (SS) OCT, yielding Dice coefficient values of84.84% and 85.15%, respectively, buttressing its effective-ness. Further, qualitative performance analysis, including manual grading, confirms that the synthesized choroid-labeled images are distinct from real images, thus ensuring data privacy. The proposed methodology marks an initial step towards developing a comprehensive choroid layer quantification tool using synthetic images, and its adaptability makes it versatile for various medical image segmentation challenges.Â

Posted on 14 Nov 2023 — CC-BY-NC-SA 4 — https://doi.org/10.36227/techrxiv.24065691.v3 — e-Prints posted on TechRxiv are preliminary reports that are not peer reviewed. They should not b...

Protecting Privacy while Improving Choroid Layer Segmentation in OCT Images: A GAN-based Image Synthesis Approach

Kiran Kumar Vupparaboina[a,*], Mohammed Nasar Ibrahim[a], Amrish Selvam[a], Shiva Vaishnavi Kurakula[b], Abdul Rasheed Mohammed[c], Natasha Mayer[a], José-Alain Sahel[a], Jay Chhablani[a], Sandeep Chandra Bollepalli[a]

***Abstract*— The choroid, positioned behind the retina, nourishes the retina by supplying oxygen and nutrients. Choroidal structural changes are associated with severe vision-threatening conditions including age-related macular degeneration (AMD) and central serous chorioretinopathy (CSCR). Optical Coherence Tomography (OCT) imaging enables the visualization of choroidal changes, and clinicians rely on quantifying choroidal biomarkers through segmentation of the choroid layer in OCT scans for precise diagnosis and disease management. Accordingly, various attempts are made at automated choroid layer segmentation, however, their practicality is constrained by the limited and biased nature of training data. Privacy regulations hinder data aggregation, and supervised machine learning requires substantial annotated data. To tackle this, we propose an innovative image synthesis approach using generative adversarial networks (GANs). It involves a three-step process: generation of choroid-labeled B-scans using a standard GAN architecture, the transformation of these scans to unlabeled B-scans via Pix2Pix-GAN, and the training of a Pix2Pix-GAN choroid segmentation model using the synthesized data. To demonstrate the generalizability and efficacy, we evaluated the proposed choroid segmentation algorithm on the real B-scans from two different OCT imaging devices: enhanced depth imaging (EDI) and swept-source (SS) OCT, yielding Dice coefficient values of 84.84% and 85.15%, respectively, buttressing its effectiveness. Further, qualitative performance analysis, including manual grading, confirms that the synthesized choroid-labeled images are distinct from real images, thus ensuring data privacy. The proposed methodology marks an initial step towards developing a comprehensive choroid layer quantification tool using synthetic images, and its adaptability makes it versatile for various medical image segmentation challenges.**

I. Introduction

The choroid, positioned between the sclera and the retina of the eye, is a vascular tissue layer performing vital functions of nourishing the retina with oxygen, and nutrients. Choroidal structural changes can be indicative of various eye diseases and conditions [1]. Timely identification of these changes through screening can allow for early diagnosis and treatment. This is particularly crucial in conditions such as choroidal neovascularization, choroidal melanoma, age-related macular degeneration (AMD), and central serous chorioretinopathy (CSCR), where prompt intervention can help prevent vision loss and preserve visual function [2], [3]. In the current clinical practice, optical coherence tomography (OCT) cross-sectional images (B-scans), which captures distinct layers of the posterior eye segment, assumes a pivotal role in the assessment of choroidal structural alterations (see Fig. 1). For quantitative disease screening and management, it is imperative to perform choroid layer segmentation using OCT images. However, detecting the boundaries of the choroid layer pose a significant challenge. Especially the choroid outer boundary (COB) commonly known as choroid sclera interface (CSI), is difficult to detect due to the absence of a strong intensity gradient in OCT images (see Fig. 1) [4].

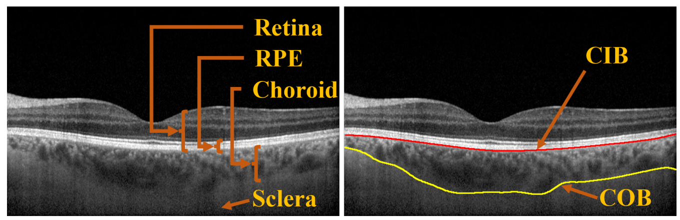

Fig. 1. Sample OCT B-scan depicting various posterior segment layers (left) and choroid layer boundaries (right). Notation: RPE – Retinal pigment epithelium, CIB – Choroid inner boundary, and COB – Choroid outer boundary.

Over the last decade, various image processing methods including gradient-based techniques, statistical models, and machine learning approaches have been explored for choroid segmentation in OCT images [5]–[18]. Despite the promising results shown by certain methods, such as those based on encoder-decoder architectures like UNet, their generalizability is often hindered by the limited and biased nature of the available training data, such as foveal B-scans from healthy individuals. As a result, segmentation models trained on such

The work was supported by the NIH CORE Grant P30 EY08098; the Eye and Ear Foundation of Pittsburgh; the Shear Family Foundation Grant; and an unrestricted grant from Research to Prevent Blindness, New York, USA to the Dept. of Ophthalmology, University of Pittsburgh.

[a]University of Pittsburgh School of Medicine, Pennsylvania, USA.
[b]Nova Scotia health authority, Halifax, Canada.
[c]School of Optometry and Vision Science, University of Waterloo, Ontario, Canada.
*Corresponding author (e-mail: kiran1559@gmail.com)

Fig. 2. Data Preparation: Generating choroid labeled image pairs for training GAN model using ground truth segmentation from ResUNet model [19]. Note that the B-scan depicted corresponds to a CSCR subject.

data may not be suitable for analyzing scans obtained from the non-foveal regions of retina and the individuals with different pathologies or conditions. Therefore, there is a pressing need to train the segmentation models with sufficient variability in training data. However, obtaining a diverse and comprehensive dataset that captures the structural complexity of the choroid layer across different ages, diseases, and disease-severity levels is challenging. Further, inherent data imbalance within healthcare organizations complicate the acquisition of heterogeneous and generalizable training data. Although collating data from multiple sources could address this issue, it is often infeasible due to privacy regulations protecting patient information.

Against this backdrop, we propose an image-synthesis based approach for choroid segmentation that eliminates the reliance on real data. In particular, we generate synthetic B-scans with choroid boundary delineation and corresponding raw B-scans to train a choroid segmentation model. A reliable synthetic data generation scheme not only facilitates the creation of diverse and balanced datasets, but also enables the data sharing while preserving patient privacy. In particular, by synthesizing the data, privacy concerns associated with sharing real patient information are eliminated. This enables collaboration and the pooling of data from different sources, resulting in a more diverse and representative dataset for training segmentation models. Thus, our approach addresses the challenge of limited data availability while still preserving patient privacy. Finally, we plan to evaluate the segmentation model trained on synthesized images using real OCT B-scans to validate its generalizability to real-world data.

GAN-based networks have proven to be effective in addressing various challenges in natural and medical image applications, including image denoising, image quality enhancement, synthetic data augmentation for segmentation and classification tasks [20]–[23]. In the context of retinal image analysis, GANs have gained significant attention, particularly in the field of OCT image analysis [24], [25]. Specifically, GANs have been applied to remove artifacts caused by motion, speckle noise, or scanner imperfections, enhancing the interpretability and reliability of OCT images [24], [26], [27]. GANs have also been utilized for OCT image super-resolution, enhancing image details and resolution through different architectures such as deep convolutional GANs (DCGANs), conditional GANs and Wasserstein GANs (WGANs) [24], [28]. Moreover, GANs have also shown promise in OCT image segmentation tasks, leveraging their generative and discriminative capabilities to accurately delineate structures of interest, such as retinal layers and pathological regions [24].

However, the specific application of GAN-based methods for synthesizing the choroid layer or segmenting the choroid layer has not been previously reported. The complexity of the choroid, characterized by its intricate vascular network and diverse tissue composition, presents unique challenges for GAN-based synthesis approaches. To this end, we propose a novel approach that involves generating choroid boundary-labeled images. Such a scheme not only captures the detailed structure of the choroid layer and its boundaries, but also eliminates the need for manual labeling of the choroid boundaries. Using such synthetic choroid boundary-labeled images, we aim to facilitate the development of subsequent choroid boundary detection algorithm for accurate segmentation of the choroid layer in real OCT images. In summary, the proposed method for choroid segmentation consists of three steps. Firstly, we generate choroid boundary-labeled images, which provide the necessary delineation of the choroid layer for accurate segmentation. Secondly, we synthesize raw OCT B-scans using these boundary-labeled images as input. This involves training a Pix2Pix GAN to generate realistic OCT B-scans that capture the specific characteristics of the choroid layer. Lastly, we train a segmentation model using the synthesized raw OCT B-scans and their corresponding segmented B-scans pairs, employing Pix2Pix GAN. This segmentation model learns to accurately identify and delineate the choroid layer in the real OCT images.

To assess the efficacy of the proposed methodology, we conducted a proof-of-concept study to demonstrate both feasibility and generalizability across different OCT devices. Specifically, we utilized images obtained from two different OCT imaging devices: enhanced depth imaging (EDI) OCT and swept-source OCT B-scans. When evaluated qualitatively, a grader can differentiate between real and synthetic images with an accuracy exceeding 96%, underscoring the capability of maintaining data privacy. Finally, we objectively validated the accuracy of the choroid segmentation model obtained using synthetic images. To this end, we employed multiple healthy and diseased OCT datasets acquired from the respective OCT devices used in our study. The segmentation accuracy was measured using the Dice coefficient (DC), which compares the algorithmic choroid layer segmentation with the ground-truth segmentations provided by experts. The mean Dice coefficient for segmentation accuracy on real images was observed to be 85.15% and 84.84% on SS-OCT and EDI OCT, respectively. This indicates that the segmentation model performed well in accurately identifying and delineating the choroid layer in real OCT images. These results provide evidence for the robustness and effectiveness of our approach in synthesizing high-quality images and achieving accurate choroid segmentation.

The main contributions of the proposed work can be summarized as follows:

1) Proposed a privacy-protected and generalizable method for choroid segmentation using OCT B-scans via GAN-based image synthesis.
2) Synthesized ground truth images with accurately delineated choroid boundaries, enabling the development of accurate choroid layer segmentation models, while eliminating the need for manual annotation.
3) Demonstrated the feasibility and generalizability of the

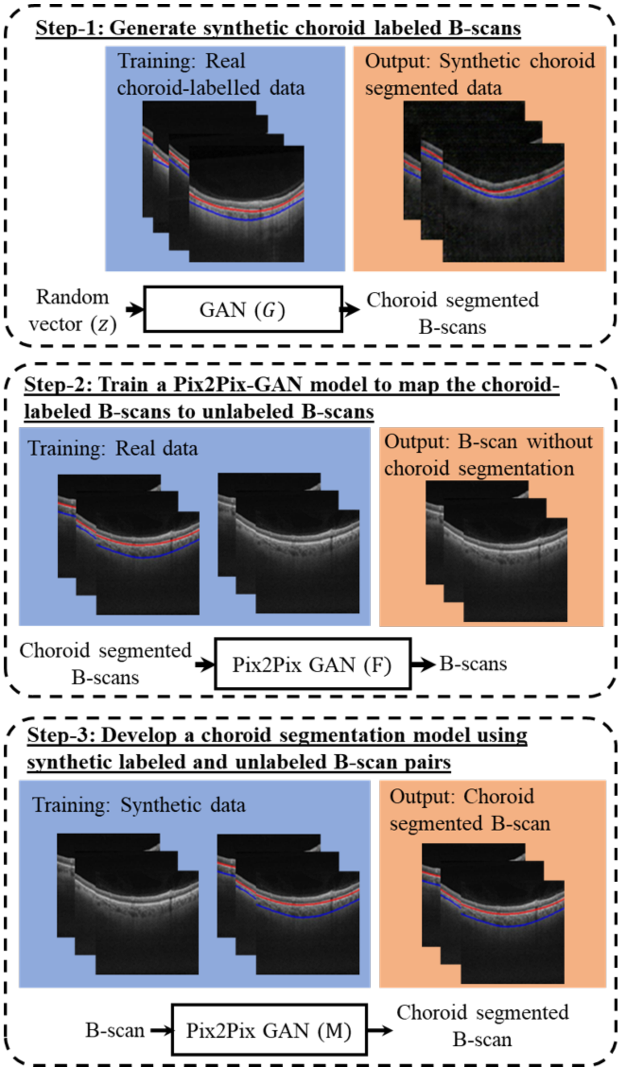

Fig. 3. Schematic of the proposed methodology.

method on healthy and diseased OCTs from two different imaging modalities (EDI-OCT and SS-OCT B-scans).

4) Evaluated the method qualitatively through subjective grading and quantitatively using the Dice coefficient, achieving high scores for synthesis quality and segmentation accuracy.

The rest of the paper is organized as follows. Section II elaborates on the proposed methodology. Experiments and results are presented in section III and finally, in section IV, the paper is concluded with a discussion and future extensions.

II. Material and Methods

We first provide the details of the datasets used in this work and then proceed to describe the proposed OCT image synthesis based choroid layer segmentation method and the evaluation strategy.

A. Dataset

This is a retrospective study performed following the ethical standards of the 1964 Helsinki Declaration with the approval of institutional review board (IRB) of the University of Pittsburgh Medical Center. Informed consent was obtained from the subjects involved. The datasets used in this work encompass a variety of OCT images obtained from two different sources: EDI-OCT (Spectralis, Heidelberg Engineering Inc.) and SS-OCT (Plex Elite 9000, Carl Zeiss) imaging devices. These datasets consist of scans acquired from diverse cohorts, including both healthy individuals and patients diagnosed with various ocular conditions such as age-related macular degeneration (AMD) and central serous chorioretinopathy (CSCR). The Spectralis EDI-OCT device was used to generate volume scans consisting of 97 B-scans each having a $9mm$ lateral and $2mm$ depth resolution and the corresponding pixel resolution is 768×472. On the other hand, the SS-OCT device produces volumetric scans consisting of 1024 B-scans each having a $12mm$ lateral and $3mm$ depth resolution, and the corresponding pixel resolution is 1024×1536.

B. Data preparation

To facilitate the training of GAN models for various tasks involving OCT images, we utilized our previously trained residual encoder-decoder (ResUNet) model to generate choroid masks [29]. Only scans with accurate choroid segmentation, as assessed by an expert grader, were included in the analysis. In particular, 1000 EDI-OCT B-scans from 70 volumes of healthy and AMD subjects as well as 1750 SS-OCT B-scans from 90 volumes of heathy, AMD and CSCR subjects are graded as good segmentations. Subsequently, the inner (CIB) and outer (COB) boundaries of the choroid, determined based on the choroidal masks, were overlaid onto the B-scans using red (CIB) and blue (COB) color markings. Figure 2 illustrates the choroid mask obtained using the ResUNet model, along with the B-scan displaying choroid boundaries marked in red (CIB) and blue (COB). We utilized OCT B-scans and their corresponding choroid boundary-labeled image pairs (CIB-red and COB-blue) as the ground truth. For training the deep learning models, noting the computational complexity, we considered randomly chosen small subsets of 100 EDI-OCT B-scans as well as 100 SS-OCT B-scans. Further, for testing, we considered another 500 EDI and 500 SS-OCT B-scans from the same set of volume scans.

C. Proposed Solution

The proposed privacy preserving choroid segmentation method consists of three steps as depicted in Figure 3 and described in detail in the following.

1) Step-1: Generate synthetic choroid-labeled B-scans: We adopted a standard GAN architecture to train the model to synthesize the choroid-labeled B-scans. Specifically, the generator network (G) takes random noise vector z as input and produces the synthetic choroid-labeled B-scan (x_s) as output, represented as $G(z) = x_s$. The discriminator network D aims to distinguish between the synthetic B-scans $x_s \in X_s$ generated by G and real B-scans $x_r \in X_r$ from a dataset of real choroid labeled B-scans. The GAN training involves iteratively updating the generator and discriminator to improve

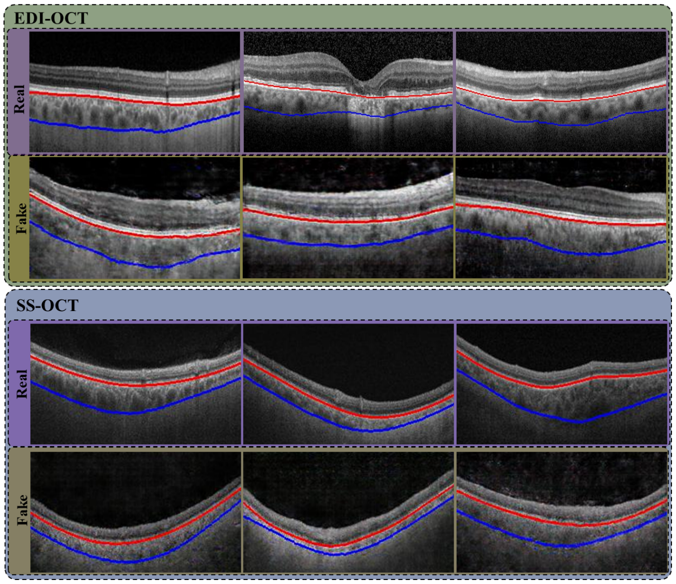

Fig. 4. Depiction of real and synthesized (GAN generated) choroid labeled B-scans.

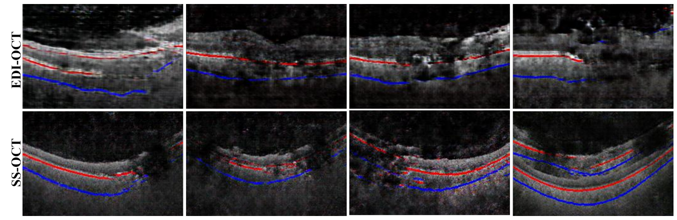

Fig. 5. Spuriously synthesized choroid labeled B-Scans.

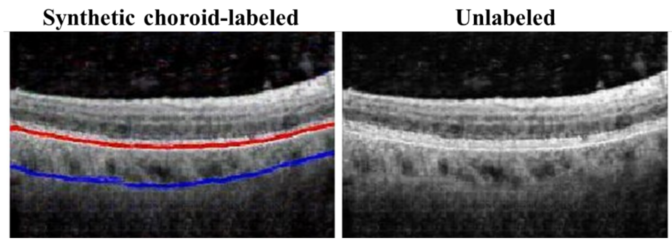

Fig. 6. Representative synthesized choroid-labeled and Pix2Pix-GAN generated unlabeled B-scan pair used for subjective grading.

the quality of generated B-scans, represented by the following min-max objective function:

$$\min_{G}\max_{D}[E(log(D(x_r))) + E(log(1 - D(G(z))))] \quad (1)$$

where E denotes the expectation over the real and synthetic data distributions. In particular, the generator architecture is a traditional feed-forward convolutional neural network (CNN) with the latent vector fully connected to 8×8×128 dimensional vector, subsequently upsampled to 256×256 dimension using 5 convolutional layers with a kernel size of 4×4 using 128 kernels and leakyReLU (α = 0.2) activation. The final layer is a convolutional layer with three kernels of size 2×2 and tanh activation. The discriminator is also a CNN network with 5 convolution layers with 128 kernels of size 3×3 in each layer, with a leakyReLU (α = 0.2) activation followed by a fully connected layer to output. For both EDI-OCT and SS-OCT modalities, we trained the GAN model for 100 epochs and generated synthetic images from both EDI-OCT and SS-OCT models (see Fig. 4). Although the synthesized images appear similar to the OCT B-scans, not all the synthesized images resemble the good quality real OCT B-scans and possess multiple artifacts, as illustrated in Figure 5. Accordingly, only the good synthesized images segregated by the expert are chosen for further steps.

2) Step-2: Train a Pix2Pix-GAN model to map the choroid-labeled B-scans to unlabeled B-scans: Note that our end goal is to perform choroid layer segmentation based on synthetic images. To this end, to train a deep learning segmentation model, we require synthesized B-scans as well as corresponding ground-truth choroid layer segmentations. In step-1, we synthesized choroid-labeled B-scans that serve as ground-truth images. However, we do not yet have the corresponding raw B-scans i.e, scans without choroid markings. In view of this, we now proceed to obtain raw B-scans from choroid-labeled B-scans. Specifically, we adopted Pix2Pix-GAN architecture, a type of conditional GAN that learns a mapping between input and output images [30]. Pix2Pix demonstrated a great ability to translate natural images to corresponding target images. Here, we plan to leverage Pix2Pix-GAN's image translation ability to translate choroid-labeled images to corresponding raw B-scans. In this context, for training the Pix2Pix-GAN model, we employed real pairs of choroid-labeled B-scans (input) and corresponding unlabeled B-scans (output) (see Figure 3) as those pairs from synthetic data are unavailable.

Let F denote the pix2pix model, which learns the mapping between the synthetic choroid labelled B-scans x_l to the corresponding unlabelled B-scans x_u. The model F can be represented as $F(x_l) = \hat{x}_u$, where $\hat{x}_u$ denotes the generated unlabelled B-scans. The pix2pix model training involves minimizing the following objective function:

$$L_{pix2pix} = E(||F(x_l) - x_u||_1) \quad (2)$$

where $||.||_1$ represents the l_1 norm, and E denotes the expectation over the synthetic choroid-labeled B-scans x_l and corresponding unlabeled B-scans x_u. The model learns to generate plausible unlabelled B-scans from the input labeled B-scans. During the training process, the Pix2Pix model learns to capture the relationship between the choroid-labeled data and the unlabeled B-scans. This mapping enables the model to generate accurate synthetic choroid-labeled and choroid-unlabeled B-scan pairs. We employed a residual encoder-decoder (ResUNet) network in the generator network and a PatchGAN in the discriminator network. The image size at the input layer of Pix2Pix model is 256×256 pixels, and we trained the model for 100 epochs.

3) Step-3: Develop a choroid segmentation model using synthetic labeled and unlabeled B-scan pairs: Now we proceed to train a model for choroid layer segmentation based on the synthetic B-scan and the corresponding choroid-labeled image pairs that were obtained in Step-2. To this end, we propose to adopt the same Pix2Pix-GAN architecture described earlier in Step-2 but to perform a reverse mapping i.e., to map synthetic raw B-scans to corresponding choroid-labeled B-scans. Indeed,

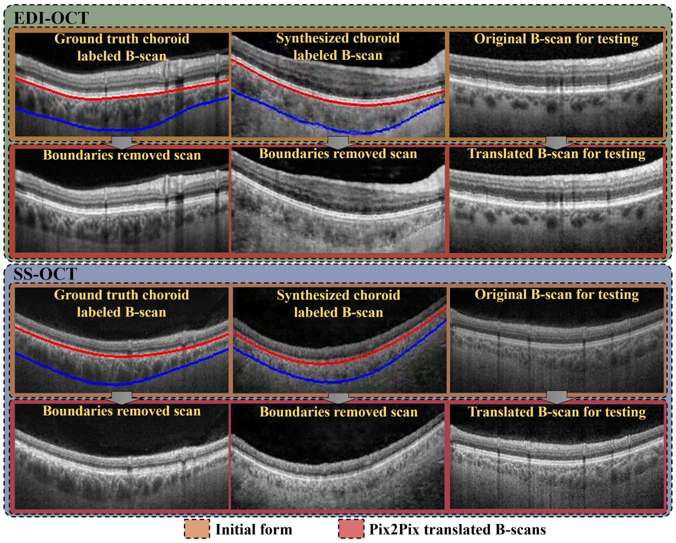

Fig. 7. Pix2Pix-GAN output for removing choroid boundary marking on the ground truth choroid labelled data, synthesized choroid labelled data and the original B-scan without choroid labels.

TABLE I

CONFUSION MATRIX OBTAINED FROM SUBJECTIVE GRADING TO DISTINGUISH REAL AND SYNTHESIZED OCT B-SCANS. NOTATION: GT – GROUND-TRUTH.

EDI-OCT		Expert Annotation	
		Original	Synthesized
GT	Original	48	2
	Synthesized	4	46

SS-OCT		Expert Annotation	
		Original	Synthesized
GT	Original	48	2
	Synthesized	0	50

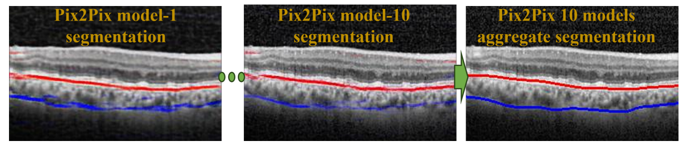

Fig. 8. Pix2Pix multi-instance aggregate segmentation: Graphical depiction of obtaining average segmentation based on 10 intermediate Pix2Pix models.

this is a novel approach for choroid layer segmentation and is reported to have achieved high accuracy of 97% on EDI-OCT B-scans [31]. In particular, let M denote the Pix2Pix model, which learns the mapping between the synthetic unlabelled B-scans x_{su} and its corresponding choroid-labeled B-scan x_{sl}. The Pix2Pix model can be represented as $M(x_{su}) = \hat{x}_{sl}$, where $\hat{x}_{sl}$ represents the generated labeled B-scan with choroid segmentations. Note that the input image size is 256×256 pixels and the number of epochs trained is 100. By leveraging the synthetic labeled and unlabelled pairs, the choroid segmentation model can learn to accurately segment the choroid region in real, unseen B-scans.

Aggregate boundaries of 10 models: We observed that, in some of the test images, the segmentation output of the model after the 100th epoch may have some discontinuities or spurious markings. Further, we observed the same in the output of the intermediate but the discontinuities or the noisy regions may be localized differently (see Figure 8). To mitigate this, we generated boundaries using the models from 10 different iterations, we then extracted the boundary by considering a pixel as the boundary even if one of the models marks that pixel as the boundary. From the aggregated boundary marking we considered the topmost pixel of every A-scan as CIB and bottommost pixel as the COB point. We then perform spline interpolation to extrapolate any discontinuities, and finally the boundaries are smoothed with RLOESS (robust locally estimated scatterplot smoothing) regression [32] with a window size of 0.1. Figure 8 depicts the aggregation of the boundaries generated by individual models and subsequent boundary smoothing.

D. Evaluation Strategy

The performance evaluation strategy for each step of the proposed methodology is described in the following:

1) Step-1: To evaluate the quality of the generated choroid-labeled images, we performed subjective grading by an expert. Specifically, we curated a set of 50 real choroid-labeled B-scans and 50 synthesized choroid-labeled B-scans, randomly mixed together. Subsequently, the expert was asked to review each scan individually and assign a grade indicating whether they perceived the scan as real or synthesized. This evaluation allowed us to gauge the degree to which the generated scans resemble the real OCT B-scans. The aforementioned grading process was carried out separately for two distinct datasets: EDI-OCT and SS-OCT, each comprising 100 scans.

2) Step-2: To evaluate the accuracy of translating choroid-labeled B-scan to raw B-scan, we employed two separate approaches for real and synthetic B-scans. For real B-scans, we test the model on the 50 real unseen choroid-labeled B-scans and compare the resultant unlabeled B-scans vis-à-vis respective original B-scans which are available for real OCT data. To this end, we obtained the mean absolute difference (MAD) of pixel intensities of both real and synthesized raw images. Next, for synthesized images, we adopted subjective grading approach as they do not have ground-truth. In particular, 50 EDI-OCT and 50-SS-OCT choroid-labeled and the corresponding unlabeled B-scan obtained from step-2 are graded by an expert grader. The grader looks at the image pair (see Figure 6 and scores, on a scale of 0 to 100, how accurately the colored choroid boundaries are replaced by actual scan information.

3) Step-3: The evaluate the choroid layer segmentation accuracy, we adopted Dice coefficient (DC) which measures the overlap between segmentations obtained from two different methods. In particular, DC between ground-truth segmentation C_{GT} and corresponding Pix2Pix-GAN based segmentation $C_{Pix2Pix}$ is given by

$$DC = \frac{2Area(C_{GT} \cap C_{Pix2Pix})}{Area(C_{GT}) + Area(C_{Pix2Pix})}. \quad (3)$$

Here, C_{GT} and $C_{Pix2Pix}$ correspond to the region between the detected CIB and COB. To this end, we obtained binary masks of the choroid layer between the boundaries CIB and COB for GT and Pix2Pix choroid segmentations. A DC value of '1' indicates the algorithmic segmentation is at par with ground truth segmentation and a value close to 1 (or 100%) is desirable.

Fig. 9. Pix2Pix-GAN output for choroid Segmentation on real EDI-OCT B-scans. Note that segmentation is performed using the model trained only on synthetic OCT B-scans.

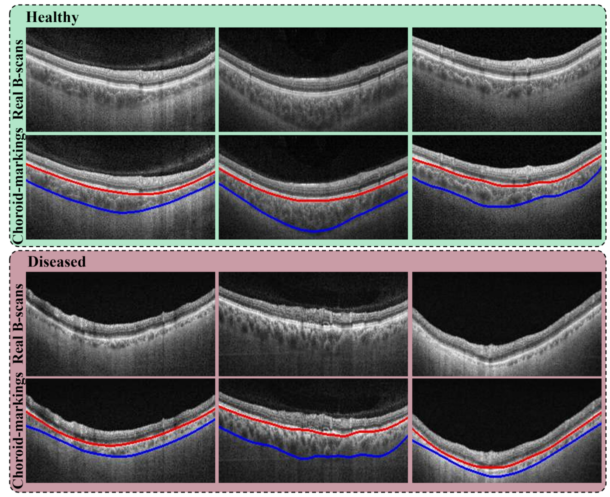

Fig. 10. Pix2Pix-GAN output for choroid Segmentation on real SS-OCT B-scans. Note that segmentation is performed using the model trained only on synthetic OCT B-scans.

E. Implementation Details

The models are trained on a workstation equipped with an Intel i9 processor, 64 GB of RAM, an NVIDIA RTX A4000 GPU, and running the Windows 10 Pro operating system. Programming was carried out within the Spyder IDE using a Python 3.6 environment coupled with TensorFlow version 2.10.0. Training, involving 100 B-scans for 100 epochs, was completed in around 24 hours. Additionally, generating images through GAN-based synthesis typically demanded around 0.2 seconds per image.

III. Results

The outcomes of performance analysis for each step are described in the following.

1) Step-1: We trained a separate GAN for EDI-OCT and SS-OCT data, and synthesized 2000 B-scans for each modality. We curated 500 high-quality images from each category, employing the process outlined in Section II. Representative images from this chosen set are presented in Fig. 4.

Subsequently, we assessed the quality of synthesis using subjective grading. Recall that grading process involved distinguishing real and synthetic OCTs from randomly presented 50 real and 50 synthesized choroid-labeled B-scans from both the EDI-OCT and SS-OCT datasets. Table I provides the confusion matrix based on the subjective grading scores. Notably, the observer achieved an accuracy of 94% for EDI-OCT and 98% for SS-OCT images. In particular, a mere 2 real EDI-OCT and 2 real SS-OCT images were incorrectly identified as synthesized. Conversely, 4 synthesized EDI-OCT images were erroneously categorized as real, while no synthesized SS-OCT images were marked as real choroid-labeled scans. Interestingly, model has not yet reached a stage of potentially confusing the grader, indicating that the generated images do not replicate the information of the real images. Consequently, there is substantial promise in leveraging these generated images to uphold patient privacy.

2) Step-2: To evaluate Pix2Pix model for generating unlabeled scans from the real choroid-labeled scans, we computed the Mean Absolute Deviation (MAD) values by comparing pixel intensities between the synthesized unlabeled scans and the corresponding ground truth unlabelled real OCTs. We achieved a MAD of 1.39 for EDI-OCT and 0.22 for SS-OCT. These results indicate the efficacy of the Pix2Pix-GAN synthesis. Further, to assess the synthesis of unlabeled scans from synthesized choroid-labeled scans, we employed subjective grading approach describe in Section II-D. Specifically, after evaluating 50 pairs of synthesized EDI-OCT images, an average grading score of 85.61% was observed. Similarly, for 50 pairs of SS-OCT images, the average grading score was recorded at 90.10%. These outcomes collectively emphasize the robustness of generating synthetic unlabeled images, affirming their quality and authenticity.

3) Step-3: Finally, we evaluated the performance of the Pix2Pix-GAN choroid segmentation model, which was trained using synthesized image pairs and tested on real images. This evaluation encompassed both the initial training dataset consisting of 100 EDI-OCT and SS-OCT B-scans from Step-1, as well as additional unseen data. For the initial dataset, we compared the GT segmentation with the Pix2Pix-GAN segmentations and obtained a Dice Coefficient (DC) values of 84.84% for EDI-OCT and 85.15% for SS-OCT. These results demonstrate the efficacy of the segmentation process for real OCT B-scans. Furthermore, we extended our assessment to an additional dataset containing 500 EDI-OCT and 500 SS-OCT B-scans not previously encountered, and achieved a DC value of 73.45% for EDI-OCT and 79.40% for SS-OCT. These findings underscore the generalizability and robustness of the proposed approach across diverse data.

IV. Discussion

In this study, we proposed an image synthesis technique based on GANs to enhance the segmentation of the choroid layer in OCT images while safeguarding patient privacy. Our

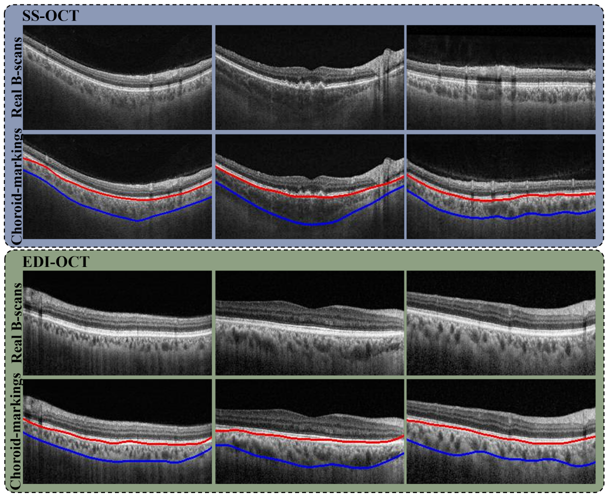

Fig. 11. Pix2Pix-GAN output for choroid Segmentation on real EDI-OCT and SS-OCT B-scans, unseen during any phase of training. Note that segmentation is performed using the model trained only on synthetic OCT B-scans.

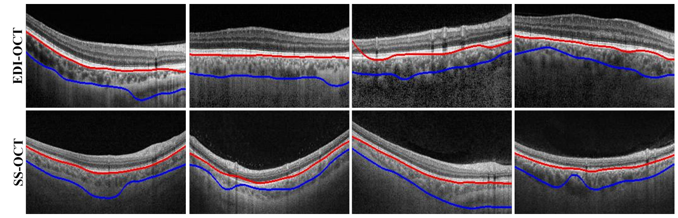

Fig. 12. Some of the spuriously segmented EDI-OCT and SS-OCT B-scans resulted from respective segmentation models developed using only synthesized B-scans.

approach comprised three main steps, *Step-1*: the generation of choroid-labeled OCT B-scans through a standard GAN network; *Step-2*: the acquisition of corresponding unlabeled OCT B-scans using Pix2Pix-GAN; and *Step-3*: the training of a Pix2Pix choroid segmentation model solely using synthesized data. By creating synthesized choroid-labeled B-scans in step-1, we not only eliminated the demanding and laborious task of manually annotating the choroid layer but also generated realistic choroid representations, ultimately contributing to the development of highly accurate choroid segmentation models.

To demonstrate the generalizability of the proposed approach, we developed two separate synthesis models for EDI-OCT and SS-OCT B-scan data modalities. Performance evaluation is carried out, based on subjective and objective metrics, to assess multiple aspects including (i) closeness of synthetic B-scans vis-à-vis real B-scans, and (ii) accuracy of choroid-layer segmentation model. In particular, our subjective grading scores of distinguishing real-synthesized OCT B-scans resulted in an accuracy of 94% and 98%, respectively, for EDI-OCT and SS-OCT B-scans substantiated that the synthesized scans do not directly copy the information from real B-scans and thereby ensuring the protection of patient-specific information. Further, favorable Dice coefficient values of 73.45% and 79.40%, respectively, for segmenting EDI-OCT and SS-OCT real B-scans, indicate the efficacy of the segmentation model developed using only synthetic images. While the overall segmentation performance is commendable, we acknowledge that there exist instances where the boundaries are detected with some sporadic inconsistencies as depicted in Fig. 12. Note that, even in these cases, the model's segmentation remains accurate for significant portions of the boundary. We aim to leverage the information from neighboring scans to correct the sporadic boundaries and enhance the robustness of the segmentation model.

This work establishes proof of concept of the proposed methodology on a limited dataset. In the future, we plan to incorporate a more diverse and larger set of OCT images, In particular, we plan to include a greater variety of anatomical variations and pathological conditions to enhance the generalizability and robustness of the image synthesis model which in turn may improve the segmentation accuracy. We also plan to automate the selection of good quality synthesized choroid-labeled scans part in Step-1 of the methodology.

We observed that the overall quality of synthesized B-scans is poor when compared to that of original B-scans which may have also played a part in the high subjective grading scores accurately distinguishing real and synthesized images. However, ideally, the quality of the synthesized B-scans should be comparable to that of the original scans. To this end, we plan to perform more experiments to make synthesis specific to OCT-imaging. We also plan to extend our work to segment other disease features such as hard exudates and retinal cysts using other imaging modalities including color fundus (CF) photography. Indeed, the proposed methodology is generalizable and can be easily adapted to even non-ophthalmic images such as magnetic resonance imaging (MRI) or computed tomography (CT) scans.

References

[1] F. H. Adler, P. L. Kaufman, L. A. Levin, and A. Alm, *Adler's Physiology of the Eye*. Elsevier Health Sciences, 2011.

[2] D. S. Dhoot, S. Huo, A. Yuan, D. Xu, S. Srivistava, J. P. Ehlers, E. Traboulsi, and P. K. Kaiser, "Evaluation of choroidal thickness in retinitis pigmentosa using enhanced depth imaging optical coherence tomography," *British Journal of Ophthalmology*, vol. 97, no. 1, pp. 66–69, 2013.

[3] R. Agrawal, J. Chhablani, K.-A. Tan, S. Shah, C. Sarvaiya, and A. Banker, "Choroidal vascularity index in central serous chorioretinopathy," *Retina*, vol. 36, no. 9, pp. 1646–1651, 2016.

[4] S. R. Singh, K. K. Vupparaboina, A. Goud, K. K. Dansingani, and J. Chhablani, "Choroidal imaging biomarkers," *Survey of Ophthalmology*, vol. 64, no. 3, pp. 312–333, 2019.

[5] V. Kajić, M. Esmaeelpour, B. Považay, D. Marshall, P. L. Rosin, and W. Drexler, "Automated choroidal segmentation of 1060 nm oct in healthy and pathologic eyes using a statistical model," *Biomedical optics express*, vol. 3, no. 1, pp. 86–103, 2012.

[6] Y. Guo, Y. Liu, T. Georgiou, and M. S. Lew, "A review of semantic segmentation using deep neural networks," *International journal of multimedia information retrieval*, vol. 7, no. 2, pp. 87–93, 2018.

[7] V. Kiran Kumar, T. R. Chandra, S. Jana, A. Richhariya, and J. Chhablani, "3d visualization and mapping of choroid thickness based on optical coherence tomography: A step-by-step geometric approach," in *International Conference on 3D Imaging (IC3D)*. IEEE, 2013, pp. 1–8.

[8] K. Lee, M. Niemeijer, M. K. Garvin, Y. H. Kwon, M. Sonka, and M. D. Abramoff, "Segmentation of the optic disc in 3-d oct scans of the optic nerve head," *Medical Imaging, IEEE Transactions on*, vol. 29, no. 1, pp. 159–168, 2010.

[9] H. Lu, N. Boonarpha, M. T. Kwong, and Y. Zheng, "Automated segmentation of the choroid in retinal optical coherence tomography images," in *35th Annual International Conference of the Engineering in Medicine and Biology Society (EMBC)*. IEEE, 2013, pp. 5869–5872.

[10] Q. Chen, S. Niu, W. Fang, Y. Shuai, W. Fan, S. Yuan, and Q. Liu, "Automated choroid segmentation of three-dimensional sd-oct images by incorporating edi-oct images," *Computer methods and programs in biomedicine*, vol. 158, pp. 161–171, 2018.

[11] B. Al-Bander, B. M. Williams, M. A. Al-Taee, W. Al-Nuaimy, and Y. Zheng, "A novel choroid segmentation method for retinal diagnosis using deep learning," in *2017 10th International Conference on Developments in eSystems Engineering (DeSE)*. IEEE, 2017, pp. 182–187.

[12] X. Sui, Y. Zheng, B. Wei, H. Bi, J. Wu, X. Pan, Y. Yin, and S. Zhang, "Choroid segmentation from optical coherence tomography with graph-edge weights learned from deep convolutional neural networks," *Neurocomputing*, vol. 237, pp. 332–341, 2017.

[13] H. Zhang, J. Yang, K. Zhou, Z. Chai, J. Cheng, S. Gao, and J. Liu, "Bionet: Infusing biomarker prior into global-to-local network for choroid segmentation in optical coherence tomography images," *arXiv preprint arXiv:1912.05090*, 2019.

[14] J. Kugelman, D. Alonso-Caneiro, S. A. Read, J. Hamwood, S. J. Vincent, F. K. Chen, and M. J. Collins, "Automatic choroidal segmentation in oct images using supervised deep learning methods," *Scientific Reports*, vol. 9, no. 1, 2019.

[15] T. T. Khaing, T. Okamoto, C. Ye, M. A. Mannan, H. Yokouchi, K. Nakano, P. Aimmanee, S. S. Makhanov, and H. Haneishi, "Choroidnet: A dense dilated u-net model for choroid layer and vessel segmentation in optical coherence tomography images," *IEEE Access*, vol. 9, pp. 150 951–150 965, 2021.

[16] Z. Hu, X. Wu, Y. Ouyang, Y. Ouyang, and S. R. Sadda, "Semiautomated segmentation of the choroid in spectral-domain optical coherence tomography volume scans," *Investigative Opthalmology Visual Science*, vol. 54, no. 3, 2013.

[17] M. Chen, J. Wang, I. Oguz, B. L. VanderBeek, and J. C. Gee, *Automated Segmentation of the Choroid in EDI-OCT Images with Retinal Pathology Using Convolution Neural Networks*, ser. Lecture Notes in Computer Science, 2017, book section Chapter 20, pp. 177–184.

[18] F. He, R. K. M. Chun, Z. Qiu, S. Yu, Y. Shi, C. H. To, X. Chen, and D. Song, "Choroid segmentation of retinal oct images based on cnn classifier and l2-lq fitter," *Computational and Mathematical Methods in Medicine*, vol. 2021, pp. 1–13, 2021.

[19] K. K. Vupparaboina, A. Selvam, S. Suthaharan, M. N. Ibrahim, S. Jana, J.-A. Sahel, K. K. Dansingani, and J. Chhablani, "Automated choroid layer segmentation based on wide-field ss-oct images using deep residual encoder-decoder architecture," *Investigative Ophthalmology & Visual Science*, vol. 62, no. 8, pp. 2162–2162, 2021.

[20] T. Zhang, H. Fu, Y. Zhao, J. Cheng, M. Guo, Z. Gu, B. Yang, Y. Xiao, S. Gao, and J. Liu, "Skrgan: Sketching-rendering unconditional generative adversarial networks for medical image synthesis," in *Medical Image Computing and Computer Assisted Intervention–MICCAI 2019: 22nd International Conference, Shenzhen, China, October 13–17, 2019, Proceedings, Part IV 22*. Springer, 2019, pp. 777–785.

[21] A. Diaz-Pinto, A. Colomer, V. Naranjo, S. Morales, Y. Xu, and A. F. Frangi, "Retinal image synthesis and semi-supervised learning for glaucoma assessment," *IEEE transactions on medical imaging*, vol. 38, no. 9, pp. 2211–2218, 2019.

[22] X. Wang, M. Xu, L. Li, Z. Wang, and Z. Guan, "Pathology-aware deep network visualization and its application in glaucoma image synthesis," in *Medical Image Computing and Computer Assisted Intervention–MICCAI 2019: 22nd International Conference, Shenzhen, China, October 13–17, 2019, Proceedings, Part I 22*. Springer, 2019, pp. 423–431.

[23] Z. Yu, Q. Xiang, J. Meng, C. Kou, Q. Ren, and Y. Lu, "Retinal image synthesis from multiple-landmarks input with generative adversarial networks," *Biomedical engineering online*, vol. 18, no. 1, pp. 1–15, 2019.

[24] A. You, J. K. Kim, I. H. Ryu, and T. K. Yoo, "Application of generative adversarial networks (gan) for ophthalmology image domains: a survey," *Eye and Vision*, vol. 9, no. 1, pp. 1–19, 2022.

[25] S. Xun, D. Li, H. Zhu, M. Chen, J. Wang, J. Li, M. Chen, B. Wu, H. Zhang, X. Chai *et al.*, "Generative adversarial networks in medical image segmentation: A review," *Computers in biology and medicine*, vol. 140, p. 105063, 2022.

[26] Y. Zhou, K. Yu, M. Wang, Y. Ma, Y. Peng, Z. Chen, W. Zhu, F. Shi, and X. Chen, "Speckle noise reduction for oct images based on image style transfer and conditional gan," *IEEE Journal of Biomedical and Health Informatics*, vol. 26, no. 1, pp. 139–150, 2021.

[27] Z. Chen, Z. Zeng, H. Shen, X. Zheng, P. Dai, and P. Ouyang, "Dn-gan: Denoising generative adversarial networks for speckle noise reduction in optical coherence tomography images," *Biomedical Signal Processing and Control*, vol. 55, p. 101632, 2020.

[28] P. Jeihouni, O. Dehzangi, A. Amireskandari, A. Rezai, and N. M. Nasrabadi, "Gan-based super-resolution and segmentation of retinal layers in optical coherence tomography scans," in *2021 IEEE International Conference on Image Processing (ICIP)*. IEEE, 2021, pp. 46–50.

[29] M. Ibrahim, S. B. Bashar, M. Rasheed, A. Selvam, V. Sant, J. Sahel, J. Chhablani, K. Vupparaboina, and S. Jana, "Volumetric quantification of choroid and haller's sublayer using oct scans: An accurate and unified approach based on stratified smoothing," *Computerized Medical Imaging and Graphics*, vol. 99, p. 102086, 2022.

[30] P. Isola, J.-Y. Zhu, T. Zhou, and A. A. Efros, "Image-to-image translation with conditional adversarial networks," in *Proceedings of the IEEE conference on computer vision and pattern recognition*, 2017, pp. 1125–1134.

[31] K. K. Vupparaboina, S. C. Bollepalli, S. R. Manne, J. Sahel, and J. Chhablani, "Choroid layer segmentation using oct b-scans: An image translation approach based on pix2pix generative adversarial networks," *Investigative Ophthalmology & Visual Science*, vol. 64, no. 8, pp. 1123–1123, 2023.

[32] W. S. Cleveland and S. J. Devlin, "Locally weighted regression: an approach to regression analysis by local fitting," *Journal of the American statistical association*, vol. 83, no. 403, pp. 596–610, 1988.