

A Novel Deep Learning Pipeline for Retinal Vessel Detection In Fluorescein Angiography

Li Ding, *Student Member, IEEE*, Mohammad H. Bawany, Ajay E. Kuriyan, Rajeev S. Ramchandran, Charles C. Wykoff, and Gaurav Sharma, *Fellow, IEEE*

Abstract—While recent advances in deep learning have significantly advanced the state of the art for vessel detection in color fundus (CF) images, the success for detecting vessels in fluorescein angiography (FA) has been stymied due to the lack of labeled ground truth datasets. We propose a novel pipeline to detect retinal vessels in FA images using deep neural networks (DNNs) that reduces the effort required for generating labeled ground truth data by combining two key components: cross-modality transfer and human-in-the-loop learning. The cross-modality transfer exploits concurrently captured CF and fundus FA images. Binary vessels maps are first detected from CF images with a pre-trained neural network and then are geometrically registered with and transferred to FA images via robust parametric chamfer alignment to a preliminary FA vessel detection obtained with an unsupervised technique. Using the transferred vessels as initial ground truth labels for deep learning, the human-in-the-loop approach progressively improves the quality of the ground truth labeling by iterating between deep-learning and labeling. The approach significantly reduces manual labeling effort while increasing engagement. We highlight several important considerations for the proposed methodology and validate the performance on three datasets. Experimental results demonstrate that the proposed pipeline significantly reduces the annotation effort and the resulting deep learning methods outperform prior existing FA vessel detection methods by a significant margin. A new public dataset, RECOVERY-FA19, is introduced that includes high-resolution ultra-widefield images and accurately labeled ground truth binary vessel maps.

Index Terms—Fluorescein angiography, generative adversarial networks, vessel detection, retinal image analysis, deep learning

Manuscript received June 19, 2019; revised January 9, 2020 and March 27, 2020; accepted April 22, 2020. The work was supported in part by a University of Rochester Research Award, by a distinguished researcher award from the New York State funded Rochester Center of Excellence in Data Science under contract CoE #3B C160189, by an unrestricted grant to the Department of Ophthalmology from Research to Prevent Blindness, and by the National Institutes of Health under Grant P30EY001319-35. Preliminary reports on the research culminating in this paper have been presented as [1]–[3]. The associate editor coordinating the review of this manuscript and approving it for publication was Dr. Christophoros Nikou. (*Corresponding author: Gaurav Sharma.*)

L. Ding and G. Sharma are with the Department of Electrical and Computer Engineering, University of Rochester, Rochester, NY 14627, USA (e-mail: {l.ding, gaurav.sharma}@rochester.edu).

A. E. Kuriyan is with Retina Service, Wills Eye Hospital, Philadelphia, PA 19107 & the University of Rochester Medical Center, University of Rochester, Rochester, NY 14642, USA (e-mail: ajay_kuriyan@urmc.rochester.edu).

M. H. Bawany and R. S. Ramchandran are with the University of Rochester Medical Center, University of Rochester, Rochester, NY 14642, USA (e-mail: {mohammad_bawany, rajeev_ramchandran}@urmc.rochester.edu).

C. C. Wykoff is with Retina Consultants of Houston and Blanton Eye Institute, Houston Methodist Hospital & Weill Cornell Medical College, Houston, TX 77030, USA (e-mail: ccwmd@houstonretina.com).

This article has supplementary downloadable material available at <http://ieeexplore.ieee.org>, provided by the authors.

Digital Object Identifier 10.1109/TIP.2020.2991530



Fig. 1: Sample fluorescein angiography (FA) images. left: fundus FA. Middle: ultra-widefield FA. Right: enlarged view of the cyan rectangle (top and bottom: the original and the contrast-enhanced views, respectively). For a larger version of this figure see Fig. 1H in the Supplementary Material.

I. INTRODUCTION

RECENTLY deep learning based image processing algorithms have shown compelling improvement in the analysis of color fundus (CF) images [4], [5]. The CF images are color images of the retina captured under white light illumination using a fundus camera that consists of a specialized microscope equipped with a camera. The images mimic what physicians see with ophthalmoscopy and are the predominant form of retinal images [6]. A DNN can detect retinal vessels in CF imagery with high accuracy and robustness [7], [8] and achieve performance close to human experts [9]. Manually labeled ground truth datasets are a key ingredient in the success of these techniques. Three commonly used datasets that provide CF images and corresponding manually labeled pixel-wise binary vessel maps include DRIVE [10] (forty 584×565 pixel images), STARE [11] (twenty 605×700 pixel images), and the high resolution HRF [12] (forty-five 3504×2336 images) datasets. The datasets provide a modest number of images and are used for training in combination with data augmentation techniques [13].

The detection of retinal vessels is also of interest for alternative imaging modalities that are of independent diagnostic utility in the clinic. For instance, fluorescein angiography (FA) and optical coherence tomography angiography (OCT-A) are used for assessing retinal non-perfusion. FA provides a larger field of imaging beyond the macula, while commercially available OCT-A provides more detailed imaging of the macular micro-vasculature. FA images are captured after intravenous injection of sodium fluorescein dye. Blue illumination, over the wavelength range from 465 to 490 nm, causes the dye to fluoresce and emit photons in the 520-530 nm green-yellow wavelength band. The spatial pattern of fluorescence intensity is captured as an FA image, in which, the vessels

with blood flowing through them appear brighter because of the fluorescent dye in the blood [14]. Although, conceptually, one could redeploy the DNN architectures that are successful in CF imagery to these alternative modalities, the fundamental differences between the modalities require fresh training and the lack of ground truth labeled data becomes a key obstacle to such reuse. Specifically, for FA images, only one dataset is available: VAMPIRE [15] which provides eight ultra-widefield FA (UWFFA) images (3072×3900 pixels, each) along with limited accuracy ground truth binary vessel maps. Manually annotating vessel maps for training a DNN is not a trivial task. Specifically, UWFFA images have high resolution and exhibit variations in contrast between the background and the vasculature, which pose a significant challenge for manual annotation. Fig. 1 shows sample FA images and highlights the particular challenge of contrast variations. The patch labeled in cyan in the middle UWFFA image is shown in an enlarged view on the right, as captured and with contrast enhanced. From the contrast enhanced view, one can appreciate that the region corresponding to the patch contains a large number of fine vessels that are rather difficult to see without contrast enhancement. In particular, ophthalmologists normally have difficulty in identifying fine vessels in the peripheral region without image enhancement because of the low contrast and brightness. High-quality annotation requires carefully adjusting image contrast for the entire FA image and labeling both major and minor vessels, making it a tedious, time-consuming, and labor-intensive process.

In this paper, we propose a novel pipeline that enables accurate vessel detection in FA images using DNNs by significantly reducing manual annotation effort. The proposed pipeline integrates the following novel elements:

- an unsupervised method for preliminary retinal vessel detection that is based on multiple scales and orientations morphological analysis,
- a cross-modality approach that transfers vessel maps from CF to FA images using robust chamfer alignment [16] in an Expectation-Maximization (EM) framework, and
- an efficient and effective human-in-the-loop iterative deep learning process for detection of retinal vessels in FA imagery that significantly reduces the tedium of generating labeled data.

We demonstrate the utility of the proposed pipeline by developing the first set of DNNs for detection of retinal vessels in FA images and evaluating the performance on alternative network architectures. The best performing method provides remarkably accurate results (maximum Dice coefficient of 0.854) and offers very significant improvements over the prior methods. Results demonstrate that the approach adapts particularly well to the contrast variations that are typical in FA imagery. To facilitate further development of vessel detection in FA images, we also release a new dataset of UWFFA images from the RECOVERY trial [17] along with ground truth labeled vessels from our pipeline. In addition to the innovative pipeline for the generation of training data, demonstration of the first deep learning approaches, evaluation of alternative architectures, and the new ground truth labeled

datasets are also contributions of the present work.

The proposed pipeline is also significant from a clinical perspective. FA is a well-established method that provides a useful imaging modality for visualizing, assessing and understanding the impact of diseases on the vascular system. Retinal vasculature changes assessed via FA imagery play a key role in the clinical assessment of vasculature changes caused by multiple common diseases, including diabetes, hypertension, and atherosclerosis, and also for eye-specific diseases, such as retinal venous occlusive diseases and retinal vasculitis. In current clinical practice, ophthalmologists manually review FA images to access disease conditions in retinal vasculature. These examinations are typically qualitative and subjective due to the limited time available during the clinical visits. Quantitative analysis of FA images, although highly desirable, requires inordinate time and patience to be performed manually and thus is not feasible in clinical settings. The proposed pipeline for detecting vessels in FA images offers an automated approach to examine retinal vasculature, which is a key component of computer-assisted retinal image analysis and diagnosis systems. Details of fine vessels are of particular diagnostic significance as changes are often first observed in the fine vessels [18]–[20]; a key strength of the method developed is the ability to reliably detect fine vessels, which are often not seen with non-FA modalities and, even for the FA modality, require significant iterative contrast manipulations for visual detection. Using the proposed pipeline, the results of retinal vessel detection achieve a level of accuracy that enables reliable computation of “digital biomarkers” from FA imagery that unlock the potential for improving clinical care, speeding up clinical trials, defining new endpoints of clinical relevance, and characterizing inter-individual variations. Preliminary work demonstrating how the analysis presented here can relate to clinical attributes of interest is being concurrently submitted in a companion paper [21].

The rest of this paper is organized as follows. Section II summarizes the existing works on retinal vessel detection. Section III provides an overview of the proposed pipeline. In Section IV, we describe the cross-modality transfer for generating ground truth data. In Section V, we introduce the human-in-the-loop learning approach for both vessel detection and manual annotation. We present the experimental results in Section VI and summarize concluding remarks in Section VII.

II. RELATED WORK

Prior work on detection of vessels in FA imagery is rather limited and due the paucity of ground truth labeled data has been primarily focused on unsupervised techniques. These methods, which are generally rule-based, include hand-crafted matched-filtering [15], active contour models [22], and morphological analysis [1], [23]. The unsupervised methods, however, offer limited accuracy (Dice coefficient of 0.634 compared to 0.854 for the best performing method benchmarked here).

Detection of retinal vessels in CF imagery has been extensively studied. For broad context, we refer the readers to a survey [24] and a recent paper [25] that categorize and

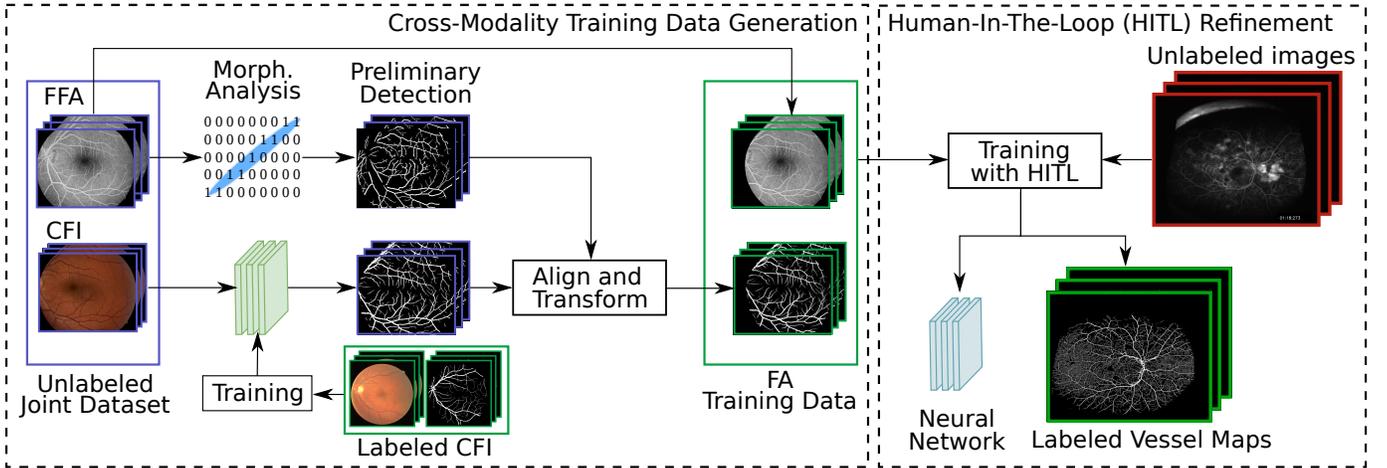


Fig. 2: Overview of the proposed pipeline for vessel detection in FA images. CFI: color fundus images; FFA: fundus fluorescein angiography. The cross-modality transfer (left block) generates the FA training data by aligning vessel maps from CF images with the preliminary vessel maps in FA images. The human-in-the-loop approach (right block) refines the neural network and significantly reduces manual annotation effort.

compare the existing methods. For our discussion, we focus on supervised methods based on deep learning which have significantly advanced the current state of the art for vessel detection in CF images. Various network architectures have been exploited, including per-pixel classifiers [7], [26], fully convolutional networks [9], [27], [28], generative adversarial networks [29], and graphical convolutional network [30]. In addition to the network architectures, several works focus on new loss terms that are particularly attuned to vessel detection [31]–[33]. The basic idea is to incorporate prior knowledge of the topology of vasculature into loss functions.

Recent work in [34] proposes a self-supervised domain adaption work to generate FA images from CF images using a CNN. While this method aims to alleviate the tedium of creating labeled data by utilizing both CF and FA images, the generated pseudo-FA images do not represent actual FA images and normally contain artifacts. In contrast, the proposed pipeline uses a cross-modality approach that directly transfers the vessel map from CF images to FA via robust chamfer registration in an EM framework, and thus is more robust and reliable than the synthesis-based approach.

III. OVERVIEW OF THE PROPOSED METHOD

The proposed pipeline, illustrated in Fig. 2, has two key components: (1) cross-modality transfer for generating an initial training dataset for FA images from CF images, and (2) a human-in-the-loop learning approach that iteratively refines DNNs and expedites the manual annotation process.

The cross-modality transfer exploits the availability of near concurrently captured CF and FA images in combination with existing deep learning methods for detection of vessels in CF imagery, for which, multiple ground truth annotated datasets are available. Specifically, we use the publicly available DR-IsfahanCFnFA (Diabetic Retinopathy Isfahan Color Fundus and Fluorescein Angiography) dataset [35] (“Unlabeled Joint Dataset” in Fig. 2) that contains pairs of CF and FA images

captured at the same clinical visit but with varying capture viewpoints. A DNN (green in Fig. 2) is trained on existing labeled CF images to extract vessel maps from unlabeled CF images. The detected vessel maps are geometrically aligned with and transferred to FA images via robust chamfer alignment [16] to a preliminary FA vessel map obtained with morphological analysis [1]. The co-aligned pairs of FA and transformed vessel map (“FA Training Data” in Fig. 2) are used as initial labeled data to train a DNN for vessel detection in FA images.

The human-in-the-loop learning approach is motivated by the synergistic relationship between deep learning and labeling. A well-trained DNN model can accurately detect vessel maps from FA images. Manually refinement of the predicted vessel map is much less time-consuming than labeling the entire image from scratch. The model performance improves with an enlarged training dataset. Thus, the training and the labeling make each other more effective. We initialize the approach with a DNN trained on the (approximate ground truth) labeled data generated from the cross-modality transfer. A human annotator then manually refines one or more of the predicted vessel maps to generate improved vessel map labels, which, in the next iteration, are incorporated in the training data to improve the DNN performance. We repeat this human-in-the-loop iterative process till the network performance improves significantly and the manual labeling introduces few changes. The end result is a trained DNN (shown in blue in Fig. 2) and a set of accurately labeled vessel maps.

Both the cross-modality transfer and the iterative learning approach reduce the burden of manual labeling significantly and engage the annotators more effectively. Instead of requiring a large number of images to be annotated before improvements are realized, in the proposed iterative approach, the annotator sees improvements in the DNN performance from iteration to iteration as an immediate reward them for their effort. A by product of this engagement and reduction of

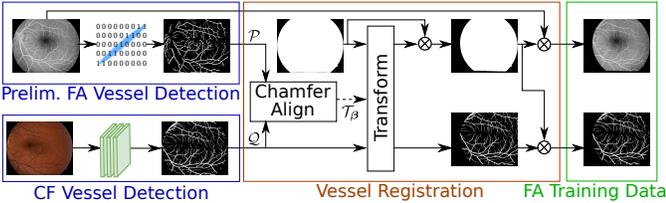


Fig. 3: Overview of cross-modality ground truth transfer. The bottom-left shows the vessel detection in unlabeled CF image with neural networks pre-trained on existing CFI dataset. The upper-left shows the preliminary vessel detection in FA obtained with unsupervised morphological analysis. The detected vessels from CF image are transformed to FA via parametric chamfer alignment with vessel maps detected from FA. The overlapping area between CFI and FFA is also estimated. The green block shows the generated training data that includes FA and co-aligned vessel maps that remains in the overlapping area.

tedium is that the images are labeled much more accurately than other studies that annotated the images from scratch (see Section VI-D).

IV. CROSS-MODALITY GROUND TRUTH TRANSFER

The cross-modality ground truth transfer, illustrated in Fig. 3, generates a training dataset for FA vessel detection from CF images. This approach consists of three steps: (1) vessel detection in CF images using a DNN, (2) preliminary vessel detection in FA for anchoring, and (3) vessel registration by parametric chamfer alignment.

A. Vessel Detection in CF Images

To detect vessels in CF images, we adopt an existing DNN proposed in [29] that exploits adversarial learning. The model is trained on DRIVE dataset [10] which scores an Area Under the Receiver Operating Characteristic Curve (AUC ROC) of 0.9803, an Area Under the Precision-Recall curve (AUC PR) of 0.915, and a Dice coefficient of 0.829. The pre-trained network is applied to overlapping patches of CF images in the DRISfahanCFnFA dataset. The final CF binary vessel map is obtained by thresholding the probability map obtained from the generator using Otsu thresholding [36].

B. Preliminary Vessel Detection In FA Images For Anchoring

A preliminary detection of vessels in FA imagery is obtained using an unsupervised method based on multiple scales and orientations morphological analysis that is attuned to the variations in directions and widths of retinal vessel structure [1]. The preliminary detection need not be particularly precise; as noted in the next section, a low false positive rate is preferable even at the cost of a higher rate of missed detections. An overview of the approach is included here and additional detail, including specific parameter settings used, are provided in Section S.IV of the Supplementary Material.

The input FA image is decomposed into multiple resolutions represented by an image pyramid [37]. Images at each scale

are processed independently and the resulting vessel maps at different scales are then combined together to generate a binary vessel map. A Gaussian pyramid expansion is used to resize vessel maps from each scale to the size of input FA image. Pixels where vessels are detected at any scale collectively comprise the estimated vessel map.

The key component in the preliminary vessel detection are morphological operators that extract locally linear patterns in terms of which the curvilinear network of interconnected vessels can be approximated. To detect vessel pixels at each scale, we choose a set of linear structuring elements S_α with the same length but oriented along different angles α , ranging from 0° to 180° . We apply the top-hat operator to the FA images using the structuring elements S_α . The conventional top-hat operator [38, pp. 557], which is defined as the difference between original and the corresponding morphological opening image, is sensitive to noise. Therefore, we adopt a modified top-hat filtering [39] to improve the robustness of vessel detection. The modified top-hat operator \odot is defined as

$$\mathbf{X} \odot \mathbf{S}_\alpha = \mathbf{X} - \min((\mathbf{X} \bullet \mathbf{S}_\alpha) \circ \mathbf{S}_\alpha, \mathbf{X}), \quad (1)$$

where \mathbf{X} is the input image, and \bullet and \circ indicate the morphological operators of image closing and opening, respectively.

Each top hat filtering operation yields a response image in which pixel locations for vessels with a matching orientation are invariably high and those for other locations are usually low. The results of the top-hat filters across different orientations are combined by taking the maximum, resulting in an overall map where high and low values are likely for vessel and background pixels, respectively. This soft vessel segmentation is converted into a binary vasculature map by locally adaptive thresholding [40]. Typically, binary vessel maps obtained by this process have a few disconnected components. As a post-processing step, we therefore perform an area opening operation to remove all small segments from the vessel map.

C. Vessel Registration By Chamfer Alignment

To precisely transfer the vessel maps in CF images to the corresponding FA images, we use parametric chamfer alignment in an EM framework [2]. Let $\mathcal{P} = \{\mathbf{p}_i\}_{i=1}^{N_P}$ and $\mathcal{Q} = \{\mathbf{q}_j\}_{j=1}^{N_Q}$ be two sets of reference and targets points corresponding to the coordinates of the vessel pixels in FA and CF images, respectively, where $\mathbf{p}_i = (x_i, y_i)^\top$ and $\mathbf{q}_j = (u_j, v_j)^\top$. Because the geometry of the image capture and of the retinal surface are unavailable, an elastic registration transform is more appropriate than a non-elastic one. An empirical evaluation of alternative geometric transformations (see Section S.III of the Supplementary Material and also [41]) indicated that a second-order polynomial transformation offers significant improvements over alternative non-elastic transforms and that higher order transforms offer little additional improvement. Therefore, we adopt a second-order polynomial transformation to align the two sets of coordinate vectors for points corresponding to detected vessels. Specifically, the

coordinate vector \mathbf{q}_j for the j^{th} point is mapped to the coordinate vector

$$\mathcal{T}_{\beta}(\mathbf{q}_j) = \begin{bmatrix} \beta_1 \\ \beta_7 \end{bmatrix} + \begin{bmatrix} \beta_2 & \beta_3 \\ \beta_8 & \beta_9 \end{bmatrix} \begin{bmatrix} u_j \\ v_j \end{bmatrix} + \begin{bmatrix} \beta_4 & \beta_5 & \beta_6 \\ \beta_{10} & \beta_{11} & \beta_{12} \end{bmatrix} \begin{bmatrix} u_j^2 \\ u_j v_j \\ v_j^2 \end{bmatrix}, \quad (2)$$

where $\beta = \{\beta_i\}_{i=1}^{12}$ are the transformation parameters and \mathcal{T}_{β} denotes the geometric transformation. The alignment error $d_j(\beta)$ for the j^{th} point under the geometric transformation \mathcal{T}_{β} , is quantified as the minimum squared Euclidean distance between the transformed location $\mathcal{T}_{\beta}(\mathbf{q}_j)$ and the nearest point from \mathcal{P} , viz.,

$$d_j(\beta) = \min_i \|\mathbf{p}_i - \mathcal{T}_{\beta}(\mathbf{q}_j)\|^2. \quad (3)$$

In the absence of outliers, the parameters β can be estimated by minimizing the average of the errors in (3), which corresponds to conventional chamfer minimization [16]. The method is, however, sensitive to outliers, that are inevitable in the detection process due to stochastic variations and noise in the imaging processes and due to differences in the FOV between the modalities. Particularly, vessel pixels in \mathcal{Q} that do not have corresponding points in \mathcal{P} inevitably cause the chamfer minimization to converge to a poor local minima, resulting in poor registration. To tackle this issue, we adopt a probabilistic formulation of chamfer alignment in an EM framework. Specifically, we introduce latent binary variables $W_j \in \{0, 1\}$ to assess putative correspondence between vessel pixel \mathbf{q}_j in CF images and vessel pixels \mathcal{P} in FA, where $W_j = 1$ indicates that \mathbf{q}_j has corresponding points in \mathcal{P} and thus is not an outlier point, and $W_j = 0$ otherwise. The prior probability of latent variable W_j follows a Bernoulli distribution with parameter $\pi = p(W_j = 1)$. Under the assumption that the points correspond, the transformed inlier vessel pixels in CF image should be located in close proximity to the vessel pixels in FA. Therefore, the alignment error is modeled as an exponential distribution with parameter λ . For outlier points, we model the alignment error as an uniform distribution over the interval $[0, D_{max}]$, where D_{max} is a free parameter. Specifically, conditioned on the latent variable and the parameters $\theta = \{\pi, \lambda, \beta\}$, the distribution of the random variable D_j corresponding to the squared distance in (3) is modeled as

$$p_{D_j|W_j, \theta}(d_j|w_j, \theta) = \begin{cases} \lambda e^{-\lambda d_j}, & \text{if } w_j = 1 \\ \frac{1}{D_{max}}, & \text{if } w_j = 0 \end{cases} \quad (4)$$

The EM algorithm seeks to obtain a maximum likelihood estimate of the parameters θ via an iterative procedure comprising two steps: an expectation (E) step and a maximization (M) step. At the $(l+1)^{\text{th}}$ iteration, the E-step computes the expectation $Q(\theta, \hat{\theta}^{(l)})$ of the complete-data log-likelihood

$$L_c(\theta) = \sum_{j=1}^{N_Q} \log p(d_j, w_j|\theta), \quad (5)$$

given the current estimate $\hat{\theta}^{(l)}$ of the parameters. In the M-step, the updated parameters $\hat{\theta}^{(l+1)}$ are determined by

maximizing $Q(\theta, \hat{\theta}^{(l)})$. For our specific setting, the E-step reduces to a computation of the posterior probabilities $p_j^{(l)} = p(W_j = 1|d_j, \theta^{(l)})$, which are obtained as

$$p_j^{(l)} = \frac{\pi^{(l)} \lambda^{(l)} e^{-\lambda^{(l)} d_j}}{\pi^{(l)} \lambda^{(l)} e^{-\lambda^{(l)} d_j} + (1 - \pi^{(l)}) \frac{1}{D_{max}}}, \quad (6)$$

The updates in the M-step become

$$\hat{\pi}^{(l+1)} = \frac{\sum_{j=1}^{N_Q} p_j^{(l)}}{N_Q}, \quad \hat{\lambda}^{(l+1)} = \frac{\sum_{j=1}^{N_Q} p_j^{(l)}}{\sum_{j=1}^{N_Q} p_j^{(l)} d_j}, \quad (7)$$

with the updated transformation parameter $\beta^{(l+1)}$ given by

$$\hat{\beta}^{(l+1)} = \arg \min_{\beta} \frac{1}{N_Q} \sum_{j=1}^{N_Q} p_j^{(l)} d_j(\beta). \quad (8)$$

By examining (8), we see that the optimal parameters are obtained by minimizing the weighted average chamfer distance where the weighting for each datapoint equals the posterior probability that it is not an outlier. This makes intuitive sense, with the EM framework, the weighting concentrates on non-outliers and discounts the impact of outliers, making the parameter estimates much more robust than direct (non-probabilistic) chamfer minimization.

The optimization problem in (8) can be solved using the iterative Levenberg-Marquardt (LM) non-linear least squares algorithm [42] in combination with the distance transform methodology [43] that significantly simplifies the computation of the objective function and required gradients with respect to the parameters β . Detailed derivations of the parameter update equations listed above are provided in Section S.II in the Supplementary Material.

The LM algorithm based transformation parameter updates in (8) can get trapped in poor local minima. This is because the LM algorithm strongly depends on the initial parameter $\hat{\beta}^{(0)}$. Thus, a good initialization is important to obtain a good solution. Instead of estimating all 12 parameters from scratch, the optimization in (8) is further performed in progressive steps that use Euclidean, similarity, affine, projective (homography), and second-order polynomial transformations, in sequence. The EM iterations are terminated when the changes in the updates become smaller than a tolerance threshold and the final estimates $\hat{\beta}$ for the transformation parameters are set to the values from the last iteration.

The binary vessel maps in CF images are registered to the corresponding FA images by applying the transformation $\mathcal{T}_{\hat{\beta}}$. To select common region where retina surface is captured in both CF and FA images, we first generate a binary mask for original CF, which is then transformed using the same transformation used for the binary vessel map. The mask for the overlapping area can be readily obtained as the intersection of the transformed mask and the original one. Only pixels remaining in the common area are selected as the inferred training data for initiating the next stage of the pipeline.

Parametric chamfer alignment is an ideal tool for registering images from different modalities. First, given the asymmetry of the chamfer distance, the preliminary vessel detector can be

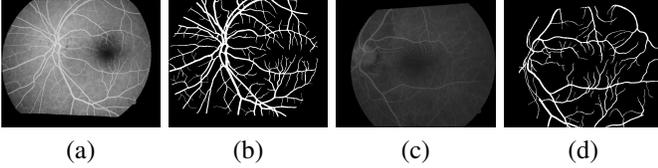


Fig. 4: Sample results of generated training data for FA imagery in DRIsFahanCFnFA dataset. (a) and (c) show two FA images, and (b) and (d) are the corresponding vessel maps. Notice that the generated vessel maps are robust under different contrast conditions. For a larger version of this figure see Fig. 4H in the Supplementary Material.

chosen to have a high specificity but a relatively low sensitivity. This means that the results of preliminary vessel detection have a low false positive rate, even though the corresponding true positive rate is low as well. In addition, the formulation uses a global matching of the detected vessels rather than relying on a set of key points with feature descriptors, which is beneficial for the polynomial parametric mapping. Finally, the incorporation of EM framework for parameter estimation significantly enhances the robustness of the registration by mitigating the effects of outlier vessel points.

As a method for generating training data for FA vessel detection, the proposed cross-modality transfer has the benefit of contrast invariance because the inferred vessels are transformed from those detected in CF images. Figures 4(a) and 4(b) show two FA images in DRIsFahanCFnFA dataset with significant variation in contrast. The corresponding vessel maps, which are shown in Figs. 4(b) and 4(d), respectively, provide consistent detection, regardless of image contrast, and capture both major and minor vessels.

V. HUMAN-IN-THE-LOOP ITERATIVE LEARNING/LABELING

Although the cross-modality transfer allows generation of a reasonable labeled dataset for training DNNs for detecting vessels in FA images, the accuracy of the labeling is limited by the differences between the modalities and the performance limitations of the CF vessel detection. The network performance can be significantly improved by providing additional better labeled ground truth data.

As indicated in Section I, manually annotating a high-resolution UWFFA image is particularly tedious and time-consuming. In this section, we present the human-in-the-loop learning approach that aims to further refine the DNN by incorporating more training data and to facilitate and expedite the manual annotation process. Figure 5 contrasts the conventional approach to annotation of training data against the proposed human-in-the-loop approach. For conventional approach, the annotation and the training are carried out in separate sequential phases, meaning that all images in the dataset are first annotated and then used for the training stage. The human-in-the-loop approach, however, is an iterative process that exploits the synergistic relationship between deep learning and labeling. The process is initialized with a trained DNN trained to detect vessels in FA images using the

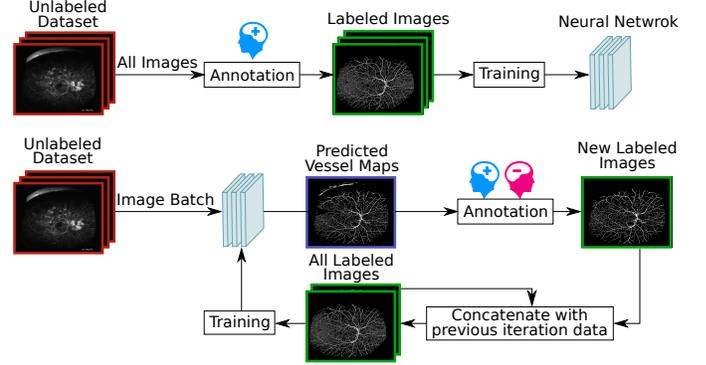


Fig. 5: Annotation and training pipelines. Top: conventional approach starts with manual annotation that generates ground truth for all images and then trains neural network with the ground truth data. Bottom: the proposed human-in-the-loop approach iterates between training neural network and manually correcting annotations generated for a batch of images using a trained network from the previous iteration.

training data obtained by the cross-modality transfer approach of Section IV. Estimated binary vessel maps that indicate the pixels corresponding to vessels are obtained for a small subset of images from an unlabeled (FA-only) dataset and used as the starting point for manual annotation. Specifically, the human annotator corrects the estimated binary vessel map by removing false positive detections and adding in false negative detections. The new labeled images are incorporated into the training dataset to refine the DNN in the next iteration. This process is repeated until all images are labeled.

The proposed human-in-the-loop approach radically reduces the effort required for annotating images (see the discussion in Section VI-B where the experiments are described). In addition to reducing the time and tedium for annotation, the approach also benefits from a psychological advantage that it provides. The annotators see the improvements in the trained network from iteration to iteration and feel immediately rewarded for their effort instead of having to label many images before seeing any machine generated annotations. This engages annotators much better than *de novo* labeling approaches, analogous to how gamification of learning and education generates better engagement [44], [45]. Our results indicate that the approach generates significantly better labeled data than the traditional *de novo* labeling approach.

A. Network Architecture

We trained and evaluated a number of alternative DNN architectures for vessel detection in FA images. In this section, we describe the best performing approach that exploits the recent concept of generative adversarial network (GAN) [46], which was also the architecture used for the human-in-the-loop labeling iterations. Detailed architectures for other neural networks are provided in Section S.IV in the Supplementary Material. To apply GAN to vessel detection, we formulate the problem as an image-to-image translation [47]. In this context, the network consists of a generator \mathcal{G} , which is trained to learn a mapping from the FA image X to the vessel map V and

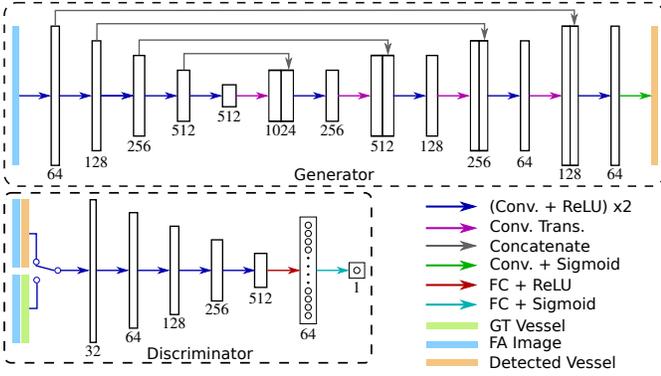


Fig. 6: Network architecture for the GAN network used with the proposed pipeline. The rectangular blocks are feature maps where heights indicate spatial dimensions. The last two blocks in the discriminator show the outputs from fully connected layers. The numbers below the rectangular block show the number of feature channels (or number of hidden units for fully connected layers).

a discriminator \mathcal{D} , which aims to distinguish between real pairs (\mathbf{X}, \mathbf{V}) and generated pairs $(\mathbf{X}, \mathcal{G}(\mathbf{X}))$ of FA images and vessel maps, where $\mathcal{G}(\mathbf{X})$ is the vessel probability map estimated from the generator and \mathbf{V} is the binary ground truth vessel map. The idea is to jointly train \mathcal{G} and \mathcal{D} to achieve the min max operating point where the vessel maps generated by \mathcal{G} minimize the maximum error for the discriminator \mathcal{D} in distinguishing between real and generated pairs.

The network architecture is visualized in Fig. 6. For the generator, we adopt the UNet [48] architecture, which comprises a downsampling path and an upsampling path. The key component in the UNet is the skip-connection that concatenates each upsampled feature map with the corresponding one in the downsampling path that has the same spatial resolution. The skip-connection is designed for detecting fine vessel structures. The discriminator receives either an image pair (\mathbf{X}, \mathbf{V}) (the blue and green bars) or $(\mathbf{X}, \mathcal{G}(\mathbf{X}))$ (the blue and yellow bars).

B. Training

The objective function for the GAN is defined as

$$\mathcal{L}_{GAN} = \mathbb{E}_{\mathbf{X}, \mathbf{V}} [\log \mathcal{D}(\mathbf{X}, \mathbf{V})] + \mathbb{E}_{\mathbf{X}} [\log (1 - \mathcal{D}(\mathbf{X}, \mathcal{G}(\mathbf{X})))] , \quad (9)$$

where minimization of the first and the second terms encourage correct classification by the discriminator \mathcal{D} of real pairs (\mathbf{X}, \mathbf{V}) sampled from training set and the pairs $(\mathbf{X}, \mathcal{G}(\mathbf{X}))$ generated by \mathcal{G} , respectively.

Inspired by the idea proposed in [47] that integrates a data loss (ℓ_1 loss) into the objective function, we combine the objective function in (9) with the binary-cross entropy loss commonly used for segmentation. Specifically, we use the segmentation loss

$$\mathcal{L}_s = -\mathbb{E}_{\mathbf{X}, \mathbf{V}} [\mathbf{V} \log \mathcal{G}(\mathbf{X}) + (1 - \mathbf{V}) \log (1 - \mathcal{G}(\mathbf{X}))] , \quad (10)$$

which penalizes the disagreement between the estimated vessel probability map $\mathcal{G}(\mathbf{X})$ and the binary ground truth vessel map \mathbf{V} .

The training procedure is then a min-max game [46] between the generator and the discriminator

$$\min_{\mathcal{G}} \max_{\mathcal{D}} \mathcal{L}_{GAN}(\mathcal{G}, \mathcal{D}) + \lambda \mathcal{L}_s(\mathcal{G}) , \quad (11)$$

where λ is the free parameter to control the relation between GAN loss and segmentation loss. The trained deep network \mathcal{G} obtained from this procedure is used to detect vessels in FA images.

VI. EXPERIMENTS

We begin by summarizing the implementation parameters, listing alternative vessel detection methods that we use as baselines for comparison, and defining the evaluation metrics that we use. We then structure our presentation of the results as follows. First, we highlight the operation and benefit of the proposed pipeline, illustrating how the cross-modality transfer and the human-in-the-loop approach reduce the burden of annotation and yield our accurately labeled RECOVERY-FA19 dataset. Next, we evaluate the performance of alternative network architectures on the UWFFA RECOVERY-FA19 dataset. Additionally, we demonstrate the broader utility of the trained networks for vessel detection in FA images, by quantifying the performance on the VAMPIRE [15] dataset and the DRIsFahanCFnFA [35] dataset, the first of which consists of UWFFA images from a source that is entirely independent of the RECOVERY-FA19 dataset and the second of fundus FA images.

A. Implementation, Baselines, and Evaluation Metrics

The preliminary vessel detection and chamfer registration discussed in Section IV are implemented in MATLABTM. Using the training data generated with the proposed pipeline, we assess the performance of several alternative DNN architectures for FA vessel detection. Specifically, we use the UNet [48], NestUNet [49], and GAN [46] architectures, where, as described in Section V the GAN uses UNet [48] as the generator. The DNNs are implemented using PyTorch [50] (Version 0.4.1). Detailed parameter settings and training protocol are provided in Section S.IV of the Supplementary Material. As baselines for performance comparisons, we use the following existing methods for vessel detection in FA images: SFAT [15], MSMA [1], and VDGAN [3].

For quantitative comparison, we use the Receiver Operating Characteristic (ROC) curve, the Precision-Recall (PR) curve, and the CAL metric [51] and its individual C , A , and L components. The ROC curve is plotted as the true positive rate (TPR, or recall) against the false positive rate (FPR) as the estimated vessel probability map from the DNN is binarized using a threshold τ ranging from 0 and 1, and the PR curve is similarly a plot of the precision versus the recall obtained by varying the threshold τ . We also report the area under curve (AUC) and the maximum Dice coefficient (DC, or F1 score) as summary measures. These metrics can be computed from

the numbers of true positive (TP), false positive (FP), true negative (TN), and false negative (FN) pixels as

$$\text{Recall} = \frac{TP}{TP + FN} \quad \text{FPR} = \frac{FP}{FP + TN}$$

$$\text{Precision} = \frac{TP}{TP + FP} \quad \text{DC} = \frac{2TP}{2TP + FP + FN}.$$

The CAL metric [51] is sensitive to anatomical features of retinal vasculature and provides better agreement with human visual judgments. CAL consists of three individual factors, C , A , and L , that quantify the consistency between the binary ground truth and the binary predicted vessel maps with regard to connectivity, overlapping area, and the corresponding skeletons (lengths), respectively. The overall CAL metric is defined as the product of C , A , and L factors, each of which ranges between 0 and 1, with 1 indicating complete consistency. The computation of the A and L factors makes use of morphological filtering operations that provide robustness against, respectively, (a) variations in the labeling of “peripheral” pixels that may be inherently uncertain because these pixels span both vessel and background regions and (b) minor perturbations in the skeletons that human observers would discount but direct pixelwise comparisons would not. The C factor equals one minus the difference between the number of connected components in the two vessel maps divided by the number of ground truth vessel pixels, truncated to zero in the unlikely scenario where the computation yields a negative value. The CAL metric computation is summarized in Section S.VIII of the Supplementary Material. The computation of the CAL metric requires a binary vessel map, which is obtained for the proposed methods by thresholding the estimated vessel probability map from the DNN with a threshold τ . We present as “CAL curves” plots of the CAL metric as a function of the threshold τ and also report the CAL value for the nominal vessel estimates obtained with a fixed threshold of $\tau = 0.5$. The computation of the C , A , and L values, and the overall CAL metric was performed using the code provided at [52].

B. Annotation of the RECOVERY-FA19 Dataset

Images for the RECOVERY-FA19 dataset were selected from those gathered for the Intravitreal Aflibercept for Retinal Non-Perfusion in Proliferative Diabetic Retinopathy trial (RECOVERY, ClinicalTrials.gov Identified: NCT02863354) [17]. The dataset comprises eight high resolution (3900×3072 pixels) UWFFA images in 8-bit TIFF format acquired using Optos California and 200Tx cameras with a 200° FOV of the retina [53]. Ground truth binary vessel map annotations were obtained for the images using the proposed pipeline described in Sections III–V. In each human-in-the-loop iteration, the network-predicted vessel map was refined by an annotator. The refinement annotations were performed using the Fiji distribution of ImageJ [54] with the segmentation editor plugin, which allows the (current estimate of the) vessel map to be overlaid on the UWFFA image to facilitate annotation. The brush tool, polygon selection, and freehand selection tools available in Fiji were used to add and remove pixels in the vessel map. The annotator adjusted the brightness and contrast of the UWFFA

images to accurately identify the vessels. The annotations were validated by consultation with two ophthalmologists who routinely use UWFFA images for diagnosis in their clinical practice and research.

To validate that the proposed pipeline can reduce the burden of annotation, at each iteration, we calculate the number of pixels changed from the network-predicted vessel map in the human-annotation process. Table I lists the number of pixels added and removed during the iterative annotation process for seven iterations. The traditional *de novo* labeling approach on average requires annotation of an estimated $1.1M$ pixels in each image. Using the proposed pipeline, in the first iteration, 36.6% ($292.4K$) pixels were added and 0.87% ($6.9K$) pixels were removed from the initial vessel map generated from the training data obtained using the cross-modality transfer approach of Section IV, which is *very significantly reduced compared with labeling from scratch*. This highlights the benefit of the cross-modality transfer approach, while also illustrating the need for improvement beyond what is achieved with that approach. Specifically, the FOV for the CF modality is smaller than for UWFFA and therefore the training dataset generated with the cross-modality transfer approach lacks fine vessel structure seen in the peripheral regions for UWFFA. As a result, in the first iteration the annotator added a significant number of pixels. As the human-in-the-loop iterations proceed, and newly labeled images are incorporated into the training dataset, the performance of DNN progressively improves requiring fewer and fewer annotation changes. In the last (7^{th}) iteration, only 2% ($19.3K$) pixels are added and only 1.4% ($14.1K$) pixels are removed. In addition to the number of changed pixels, we also estimated the time needed for annotation. Previous work [55] stated that approximately 18 hours are required to label one UWF fundus photograph, which is lower resolution and has a narrower field of view than the UWFFA images that are our primary focus. Our empirical estimation based on *de novo* annotation of two square 512×512 patches indicates that approximately 150 hours would be needed to annotate an entire UWFFA image from scratch. Using the proposed human-in-the-loop approach, the required time is decreased, very significantly, to about 25 hours per image, where most of the time is spent in validating the labeling. As noted in Section V, the progressive improvements in the network performance also have a positive psychological impact as the annotator realizes that the tedium of labeling is progressively decreasing.

Iteration	# images	# pixels added	# pixel removed
0*	-	1.1M (100%)	0.0
1	1	292.4K (36.6 %)	6.9K (0.87 %)
2	2	79.1K (13.0 %)	13.0K (0.99 %)
3	1	42.1K (3.8 %)	7.8K (0.7 %)
4	1	32.7K (2.9%)	14.1K (1.3%)
5	1	21.4K (1.7%)	9.1K (0.7%)
6	1	20.4K (1.5%)	26.2K (1.9%)
7	1	19.3K (2.0%)	14.1K (1.4%)

TABLE I: Number of pixels changed in each iteration in the proposed human-in-the-loop process. *The row labeled iteration 0 lists the estimated number of pixels that would need to be added to a vessel map, starting from scratch.

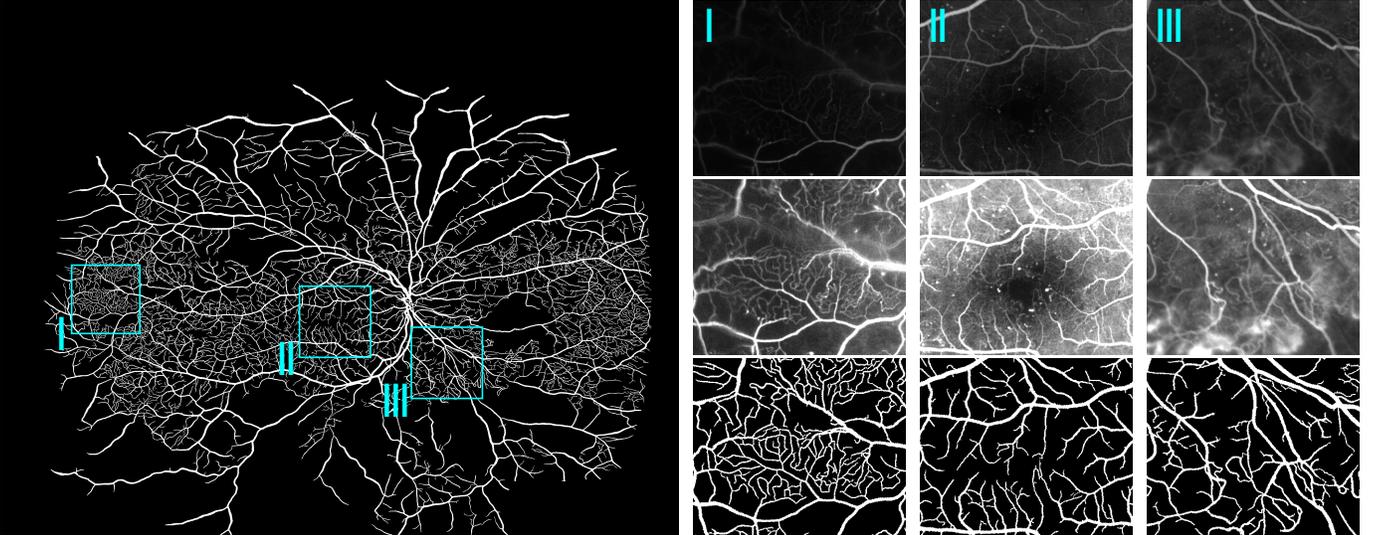


Fig. 7: Example of labeled ground truth vessel map from the RECOVERY-FA19 dataset. Enlarged views of cyan rectangles are shown on the right (top: original view; middle: contrast-enhanced view; bottom: labeled ground truth). The corresponding UWFFA image is shown in Fig. 1.

The annotated vessel maps obtained by the human-in-the-loop iterations along with the corresponding eight UWFFA images constitute a new labeled dataset for vessel detection in FA images, which we refer to as the RECOVERY-FA19 dataset [56]. The RECOVERY-FA19 dataset contains fine vessel branches, leakage, neo-vasculature, and retinal non-perfusion, which make the vessel detection more challenging. These attributes are of particular diagnostic significance [57] but are barely seen in the prior VAMPIRE dataset [15]. Figure 7 shows an example of labeled ground truth vessel map for the UWFFA image in Fig. 1. The ground truth annotations for RECOVERY-FA19 are also significantly better than for VAMPIRE, which we attribute primarily to the pipeline proposed in this paper, which significantly reduces the tedium of labeling and improves annotator engagement considerably.

C. Evaluations on the RECOVERY-FA19 Dataset

In the course of the human-in-the-loop iterations, labeled ground truth data is combined with prior iteration training data to generate training data for the next iteration. A limitation of this setting is that only the added ground truth data at each iteration is the “test” data independent of the training data. Therefore, we evaluate the performance of alternative network architectures on the RECOVERY-FA19 dataset by two alternative approaches. First, we use leave-one-out cross validation [58], where the model is trained on seven of the eight UWFFA images and the corresponding ground truth vessel map labels and tested on the remaining image. The performance of the model is then reported in terms of statistics of the eight evaluation metrics¹. Second, we also evaluate the approach using image patches that are labeled *de novo* and therefore completely independent of the training process.

¹The estimated vessel maps and code for computing the reported statistics is provided as a Code Ocean capsule [59].

For the leave-one-out cross validation, Fig. 8 shows the ROC, the PR, and the CAL curves for different methods and Table II summarizes the AUC for the ROC and the PR curves, the maximum DC, and the CAL. The best performing network (Prop. + GAN) achieves an AUC ROC of 0.987, an AUC PR of 0.930, the maximum DC of 0.854, and a CAL of 0.760. Using the proposed pipeline, all DNNs show significant improvement over traditional methods SFAT [15] and MSMA [1]. The performance is also significantly better than that obtained with the precursor to the present work [3], where only the cross-modality transfer was used. This highlights the benefit of the human-in-the loop iterations in the proposed pipeline. In Figure 9, we show qualitative results of different methods. Notice that the proposed pipeline is robust to contrast variations. Fine vessels are detected in the periphery that has extremely low contrast and brightness. Although these details in vasculature can be seen manually by repeatedly adjusting contrast and viewing different regions, the burden and the time requirement for doing this are prohibitive in typical clinical settings. The proposed pipeline also handles capillary leakage, neo-vasculature, and retinal non-perfusion, as shown in the enlarged views in Fig. 9.

Methods	AUC ROC	AUC PR	Max DC	CAL (C, A, L)
SFAT [15]	-	-	0.606	0.335 (0.999, 0.606, 0.550)
MSMA [1]	-	-	0.634	0.362 (0.999, 0.622, 0.579)
VDGAN [3]	0.981	0.883	0.800	0.687 (0.995, 0.844, 0.817)
Prop. + UNet	0.987	0.923	0.842	0.753 (0.996, 0.887, 0.853)
Prop. + NestUNet	0.955	0.900	0.817	0.698 (0.995, 0.858, 0.816)
Prop. + GAN	0.987	0.930	0.854	0.760 (0.999, 0.889, 0.856)

TABLE II: Quantitative results obtained from different methods on the RECOVERY-FA19 dataset. The best result is shown in bold. The individual *C*, *A*, and *L* values are listed parenthetically.

The classification into vessel and background categories is inherently uncertain for the edge pixels that span both vessel and background regions. For such pixels, the human-in-the-

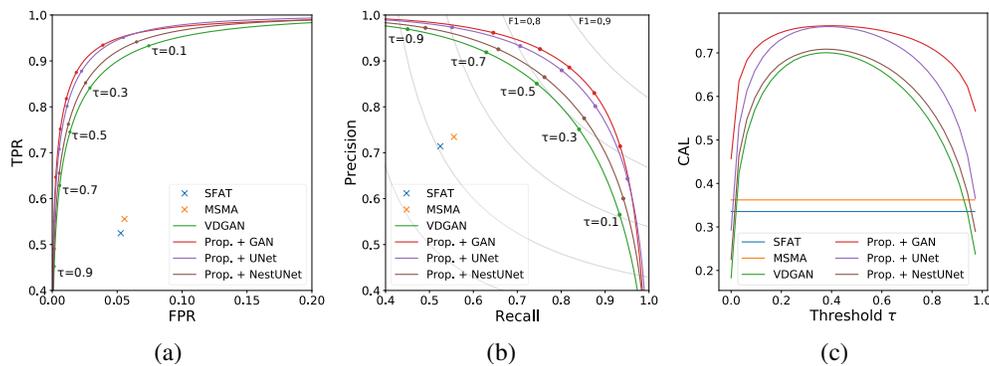


Fig. 8: (a) ROC, (b) PR, and (c) CAL curves for different methods on the RECOVERY-FA19 dataset. The gray curves in (b) represent the isolines of Dice coefficients. The small circular dots on the curves in (a) and (b) identify the corresponding values of the threshold τ .

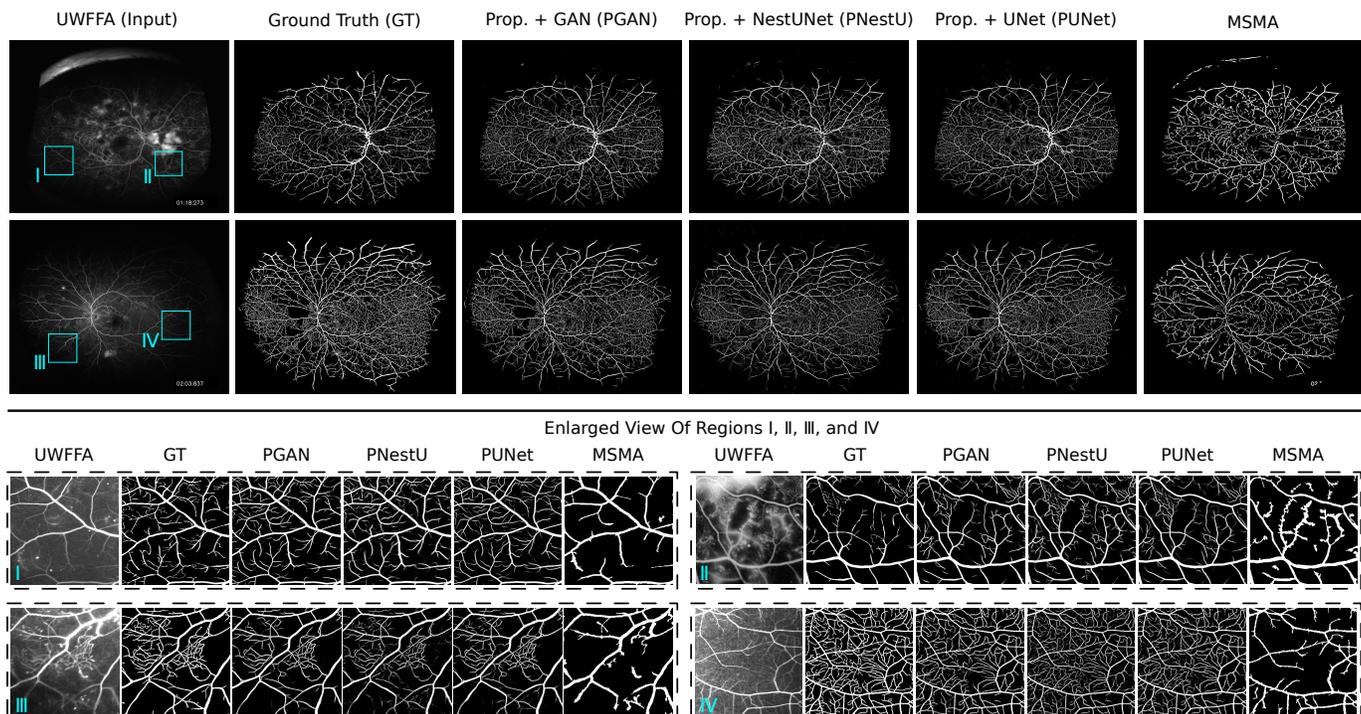


Fig. 9: Qualitative comparison of results obtained with different algorithms for images from the RECOVERY-FA19 dataset. For each full image, two contrast-enhanced enlarged views of the selected regions (shown by cyan rectangles) are also included.

loop labeling process is potentially subject to confirmation bias, wherein labels for these pixels are simply validated by the human observer instead of being critically re-evaluated. The CAL metrics are designed to be robust against such uncertainty. Therefore, the improvements in the CAL metrics for the proposed methods over prior alternatives shown in Fig. 8 and Table II represent actual improvements that are not impacted by the potential confirmation bias. On the other hand, the pixel-wise metrics (TPR (Recall)/FPR/Precision/Dice coefficient) may be impacted by the afore-mentioned confirmation bias. To address this potential concern, we also performed a second evaluation using a *de novo* labeled dataset. Because the time requirements for labeling entire images from scratch are prohibitive, the evaluation on *de novo* labeled data relied only on two square image patches of 512×512 pixels. Each

patch required about 10 hours for the *de novo* labeling, which translates to approximately 150 hours for labeling a full high resolution UWFFA image. The selected patches cover both central and peripheral retina and represent both major and minor vessel branches. Table III reports quantitative results for the evaluation performed using the *de novo* labeled data. The results, while slightly worse than those reported for the cross-validation based evaluation, reinforce the overall findings: the proposed approaches outperform the alternatives by relatively large margins. Importantly, the C , A , and L factors of the CAL metrics, that are designed to be robust against alternative classifications of the uncertain pixels, are comparable for the *de novo* and the cross-validation evaluations. As additional validation, we also assessed the consistency between the labelings for the same image patches obtained *de novo* and

using the human-in-the-loop approach. The results, presented in Section S.V of the Supplementary Material, illustrate that the level of consistency is comparable to that obtained between different human annotators.

Methods	AUC ROC	AUC PR	Max DC	CAL (C, A, L)
SFAT [15]	-	-	0.559	0.328 (0.999, 0.664, 0.495)
MSMA [1]	-	-	0.644	0.409 (0.999, 0.713, 0.574)
VDGAN [3]	0.954	0.849	0.747	0.710 (0.981, 0.871, 0.832)
Prop. + UNet	0.958	0.852	0.756	0.726 (0.986, 0.881, 0.836)
Prop. + NestUNet	0.955	0.855	0.761	0.687 (0.983, 0.846, 0.825)
Prop. + GAN	0.951	0.861	0.768	0.732 (0.995, 0.879, 0.837)

TABLE III: Quantitative results obtained from the different methods on the *de novo* labeled dataset. The best result is shown in bold. The individual *C*, *A*, and *L* values are listed parenthetically.

D. Evaluations on the VAMPIRE and DRIsFahanCFnFA Datasets

The FA imaging modality shares common physical characteristics across alternative imaging options and therefore the proposed methodology is useful for both UWFFA and fundus FA images. To demonstrate the broader utility of the networks (only) trained on RECOVERY-FA19 dataset, we test the vessel detection performance on two additional datasets: VAMPIRE [15] and DRIsFahanCFnFA [35] datasets.

The VAMPIRE dataset [15] provides eight high resolution (3900×3072 pixels) UWFFA images acquired using the OPTOS P200C camera [53] with a 200° FOV of the retina. There are two sequences of images in the VAMPIRE dataset representing a healthy retina (GER) and a retina with age-related macular degeneration (AMD). For each image, a binary vessel map that is manually annotated by ophthalmologists is provided as ground truth. We detected vessels in the UWFFA images from the VAMPIRE dataset using the best performing (Prop+GAN) network that was trained on the RECOVERY-FA19 dataset.

Our results reveal an issue with the VAMPIRE dataset: we notice that the vessel branches are not fully-labeled, especially in peripheral regions where the images have extremely low contrast. As mentioned in Section I, contrast and exposure normally pose a big challenge for manual annotation. To demonstrate the issue, we visually examine the result for the image “AMD2” in the VAMPIRE dataset, as shown in Fig. 10. Using the labeled vessel map provided with the VAMPIRE dataset as “ground truth”, we visualize true positive (black), false positive (red), false negative (blue), and true negative (white), as shown in the middle image in the first row of Fig. 10. After closely examining the vessel detection results, we observe that most “false positive” detections are indeed true vessels but are not annotated in the original labeling. For example, the second and the third rows of Fig. 10 show six rectangular regions where the true vessel branches are missed. This illustrates that quantitative comparisons using the original labeling for the VAMPIRE dataset are not reliable. To remedy the situation, we selected two images, “AMD2” and “GER4”, from the dataset and obtained (refined) ground truth vessel map annotations for these using the human-in-the-loop approach. The fourth row of Fig. 10 shows the

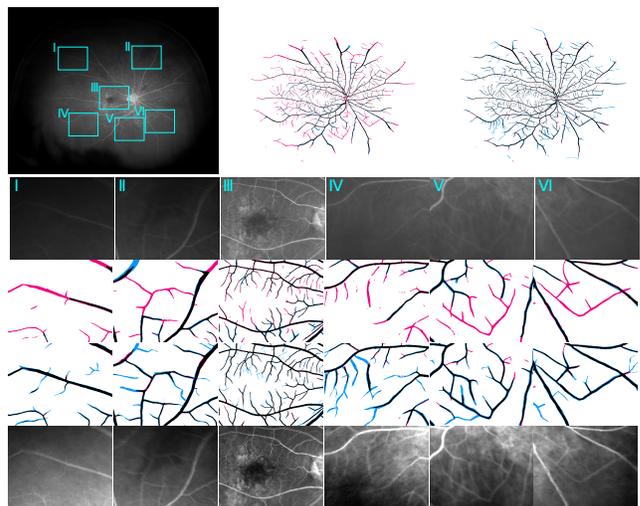


Fig. 10: Sample results of vessel detection on the VAMPIRE dataset [15]. The first row, from left to right: UWFFA, vessel map evaluated on the original VAMPIRE ground truth, and the vessel map evaluated on the refined ground truth. Black, red, and blue indicates true positive, false positive, and false negative, respectively. The second to the fourth rows show the enlarged views of six rectangular regions marked on the wide-field FA images and corresponding results, respectively. The “false positive” detections in the third row are actually true vessels that are not labeled in the VAMPIRE dataset. In the last row, we show the images after contrast enhancement for a better visualization.

same enlarged views as earlier, evaluated on the ground truth images. Compared with the evaluation using the original labeling (the third row of Fig. 10), the evaluation using the ground truth data indicates that the detected vessel map has much less false positives. On average, 73% of the original false positive detections becomes true positive if they are evaluated using the ground truth.

In Section S.VI of the Supplementary Material, we: (a) report complete quantitative evaluations performed with the original labeling and contrast these against evaluations over the two images with refined ground truth, and (b) present evaluations of the alternative methods over “trusted regions” where the VAMPIRE annotations were accurate, excluding “non-trusted” regions where our refined ground truth clearly identified vessels that were not labeled (correctly). Just like the experiment with the *de novo* labeling in Section VI-C, the latter evaluations address the issue of potential confirmation bias for edge pixels in the human-in-the-loop labeling process.

The DRIsFahanCFnFA dataset [35] contains 59 pairs of near concurrently captured CF and FA images. All images have the same resolution of 576×720 pixels. The ground truth binary vessel maps are obtained using the proposed pipeline described in Sections III–V. We report the quantitative results in Table IV. The best performing method achieves an AUC ROC of 0.974, an AUC PR 0.887, the maximum DC of 0.808, and a CAL of 0.783, outperforming other baseline methods. Visual results of detected vessel maps and the ROC and the

PR curves are shown in Section S.VII of the Supplementary Material.

Methods	AUC ROC	AUC PR	Max DC	Max CAL (C, A, L)
SFAT [15]	-	-	0.607	0.432 (0.991, 0.655, 0.656)
MSMA [1]	-	-	0.691	0.504 (0.999, 0.720, 0.688)
VDGAN [3]	0.965	0.851	0.776	0.728 (0.996, 0.868, 0.840)
Prop. + UNet	0.972	0.883	0.802	0.743 (0.997, 0.878, 0.847)
Prop. + NestUNet	0.972	0.882	0.804	0.761 (0.997, 0.889, 0.858)
Prop. + GAN	0.974	0.887	0.808	0.783 (0.997, 0.899, 0.872)

TABLE IV: Quantitative results obtained from different methods on the DRISfahanCFnFA dataset. The best result is shown in bold. The individual C , A , and L values are listed parenthetically.

VII. CONCLUSION

We proposed a novel deep learning pipeline for detecting retinal vessels in FA images. Using a cross-modality approach and a human-in-the-loop approach, our pipeline significantly reduces the effort required for generating labeled ground truth images. Experimental validations on three datasets, including a new RECOVERY-FA19 UWFFA dataset, demonstrate that the proposed pipeline significantly outperforms existing methods. To facilitate further development and evaluation of retinal vessel detection in FA images, we make publicly available the RECOVERY-FA19 dataset [56] and a Code Ocean capsule [59] for replicating the results in Table II.

The proposed pipeline provided a particularly useful methodology for generating labeled ground truth data. While our focus here was on labeling vessels in FA retinal images, the key underlying ideas could be applied in other situations. The registration approach that we describe in Section IV can also be used to facilitate identification and comparison of longitudinal vessel changes, preliminary results on which have been reported in [2]. The idea of cross-modality (label) transfer by registering observations of the same object captured with different modalities is potentially useful in speeding up other ground truth labeling tasks. Used in combination with the human-in-the-loop approach, such methods can significantly reduce tedium and improving engagement, and improve availability of datasets with accurately labeled ground truth, which is currently a key bottleneck in deploying deep learning solutions for a number of problems.

ACKNOWLEDGMENTS

The authors would like to thank the Center for Integrated Research Computing, University of Rochester, for providing access to computational resources and Shaun Lampen, Alex Rusakevich, and Brenda Zhou for collecting the UWFFA images for the RECOVERY trial. They also thank the anonymous reviewers for their comments and suggestions, which have significantly improved this article.

REFERENCES

- [1] L. Ding, A. Kuriyan, R. Ramchandran, and G. Sharma, "Multi-scale morphological analysis for retinal vessel detection in wide-field fluorescein angiography," in *Proc. IEEE Western NY Image and Signal Proc. Wksp. (WNYISPW)*, Rochester, NY, Nov. 2017, pp. 1–5.
- [2] L. Ding, A. Kuriyan, R. Ramchandran, and G. Sharma, "Quantification of longitudinal changes in retinal vasculature from wide-field fluorescein angiography via a novel registration and change detection approach," in *Proc. IEEE Intl. Conf. Acoustics Speech and Sig. Proc.*, Apr. 2018, pp. 1070–1074.
- [3] L. Ding, A. Kuriyan, R. Ramchandran, and G. Sharma, "Retinal vessel detection in wide-field fluorescein angiography with deep neural networks: A novel training data generation approach," in *IEEE Intl. Conf. Image Proc.*, Oct 2018, pp. 356–360.
- [4] V. Gulshan, L. Peng, M. Coram, M. C. Stumpe, D. Wu, A. Narayanaswamy, S. Venugopalan, K. Widner, T. Madams, J. Cuadros *et al.*, "Development and validation of a deep learning algorithm for detection of diabetic retinopathy in retinal fundus photographs," *JAMA*, vol. 316, no. 22, pp. 2402–2410, 2016.
- [5] R. Poplin, A. V. Varadarajan, K. Blumer, Y. Liu, M. V. McConnell, G. S. Corrado, L. Peng, and D. R. Webster, "Prediction of cardiovascular risk factors from retinal fundus photographs via deep learning," *Nature Biomed. Eng.*, vol. 2, no. 3, pp. 158–164, 2018.
- [6] M. D. Abramoff, M. K. Garvin, and M. Sonka, "Retinal imaging and image analysis," *IEEE Rev. Biomed. Eng.*, vol. 3, pp. 169–208, 2010.
- [7] P. Liskowski and K. Krawiec, "Segmenting retinal blood vessels with deep neural networks," *IEEE Trans. Med. Imaging*, vol. 35, no. 11, pp. 2369–2380, Nov 2016.
- [8] H. Fu, Y. Xu, S. Lin, D. W. K. Wong, and J. Liu, "DeepVessel: Retinal vessel segmentation via deep learning and conditional random field," in *Intl. Conf. Med. Image Computing and Computer-Assisted Intervention*, 2016, pp. 132–139.
- [9] K.-K. Maninis, J. Pont-Tuset, P. Arbeláez, and L. Van Gool, "Deep retinal image understanding," in *Intl. Conf. Med. Image Computing and Computer-Assisted Intervention*, 2016, pp. 140–148.
- [10] J. Staal, M. Abràmoff, M. Niemeijer, M. Viergever, and B. van Ginneken, "Ridge based vessel segmentation in color images of the retina," *IEEE Trans. Med. Imaging*, vol. 23, no. 4, pp. 501–509, 2004.
- [11] A. Hoover, V. Kouznetsova, and M. Goldbaum, "Locating blood vessels in retinal images by piecewise threshold probing of a matched filter response," *IEEE Trans. Med. Imaging*, vol. 19, no. 3, pp. 203–210, 2000.
- [12] A. Budai, R. Bock, A. Maier, J. Hornegger, and G. Michelson, "Robust vessel segmentation in fundus images," *Intl. J. of Biomed. imaging*, vol. 2013, 2013.
- [13] A. Krizhevsky, I. Sutskever, and G. E. Hinton, "Imagenet classification with deep convolutional neural networks," in *Adv. in Neural Info. Proc. Sys.*, 2012, pp. 1097–1105.
- [14] T. J. Bennett, D. A. Quillen, and D. Rolly Coronica, "Fundamentals of fluorescein angiography," *Curr Concepts Ophthalmology*, vol. 9, no. 3, pp. 43–9, 2001.
- [15] A. Perez-Rovira, K. Zutis, J. Hubschman, and E. Trucco, "Improving vessel segmentation in ultra-wide field-of-view retinal fluorescein angiograms," in *IEEE Intl. Conf. Eng. in Med. and Biol. Soc.*, Aug. 2011, pp. 2614–2617.
- [16] H. G. Barrow, J. M. Tenenbaum, R. C. Bolles, and H. C. Wolf, "Parametric correspondence and chamfer matching: Two new techniques for image matching," in *Proc. Int. Joint Conf. Artificial Intell.*, 1977, pp. 659–663.
- [17] C. C. Wykoff, "Intravitreal aflibercept for retinal non-perfusion in proliferative diabetic retinopathy (RECOVERY)," accessed 31 May 2019. [Online]. Available: <https://clinicaltrials.gov/ct2/show/NCT02863354>
- [18] M. S. AlSalhi, S. Devanesan, K. E. AlZahrani, M. AlShebly, F. Al-Qahtani, K. Farhat, and V. Masilamani, "Impact of diabetes mellitus on human erythrocytes: Atomic force microscopy and spectral investigations," *Intl. J. Environ. Res. and Public Health*, vol. 15, no. 11, p. 2368, 2018.
- [19] S. Ghosh, I. Chakraborty, M. Chakraborty, A. Mukhopadhyay, R. Mishra, and D. Sarkar, "Evaluating the morphology of erythrocyte population: An approach based on atomic force microscopy and flow cytometry," *Biochim. Biophys. Acta*, vol. 1858, no. 4, pp. 671–681, 2016.
- [20] A. Hall, "Recognising and managing diabetic retinopathy," *Community Eye Health*, vol. 24, no. 75, pp. 5–9, 2011.
- [21] M. H. Bawany, L. Ding, R. S. Ramchandran, G. Sharma, C. C. Wykoff, and A. E. Kuriyan, "Automated vessel density detection in fluorescein angiography images correlates with vision in proliferative diabetic retinopathy," submitted, under review.
- [22] Y. Zhao, L. Rada, K. Chen, S. P. Harding, and Y. Zheng, "Automated vessel segmentation using infinite perimeter active contour model with hybrid region information with application to retinal images," *IEEE Trans. Med. Imaging*, vol. 34, no. 9, pp. 1797–1807, 2015.

- [23] F. Zana and J. Klein, "Segmentation of vessel-like patterns using mathematical morphology and curvature evaluation," *IEEE Trans. Image Proc.*, vol. 10, no. 7, pp. 1010–1019, July 2001.
- [24] C. L. Srinidhi, P. Aparna, and J. Rajan, "Recent advancements in retinal vessel segmentation," *J. Med. Syst.*, vol. 41, p. 70, 2017.
- [25] Z. Fan, J. Lu, C. Wei, H. Huang, X. Cai, and X. Chen, "A hierarchical image matting model for blood vessel segmentation in fundus images," *IEEE Trans. Image Proc.*, vol. 28, no. 5, pp. 2367–2377, May 2019.
- [26] D. Maji, A. Santara, P. Mitra, and D. Sheet, "Ensemble of deep convolutional neural networks for learning to detect retinal vessels in fundus images," *arXiv:1603.04833*, 2016.
- [27] Q. Li, B. Feng, L. Xie, P. Liang, H. Zhang, and T. Wang, "A cross-modality learning approach for vessel segmentation in retinal images," *IEEE Trans. Med. Imaging*, vol. 35, no. 1, pp. 109–118, 2016.
- [28] A. Dasgupta and S. Singh, "A fully convolutional neural network based structured prediction approach towards the retinal vessel segmentation," in *IEEE Intl. Symp. Biomed. Imaging*, 2017, pp. 248–251.
- [29] J. Son, S. J. Park, and K.-H. Jung, "Towards accurate segmentation of retinal vessels and the optic disc in fundoscopic images with generative adversarial networks," *J. Digital Imaging*, vol. 32, no. 3, pp. 499–512, Jun 2019.
- [30] S. Y. Shin, S. Lee, I. D. Yun, and K. M. Lee, "Deep vessel segmentation by learning graphical connectivity," *Med. Image Analysis*, vol. 58, p. 101556, 2019.
- [31] Z. Yan, X. Yang, and K. T. Cheng, "A three-stage deep learning model for accurate retinal vessel segmentation," *J. Biomed. and Health Informatics*, vol. 23, no. 4, pp. 1427–1436, Sept 2018.
- [32] Z. Yan, X. Yang, and K. Cheng, "Joint segment-level and pixel-wise losses for deep learning based retinal vessel segmentation," *IEEE Trans. Biomed. Eng.*, vol. 65, no. 9, pp. 1912–1923, 2018.
- [33] A. Mosinska, P. Márquez-Neila, M. Koziński, and P. Fua, "Beyond the pixel-wise loss for topology-aware delineation," in *IEEE Intl. Conf. Comp. Vision, and Pattern Recog.*, June 2018, pp. 3136–3145.
- [34] Á. S. Hervella, J. Rouco, J. Novo, and M. Ortega, "Retinal image understanding emerges from self-supervised multimodal reconstruction," in *Intl. Conf. Med. Image Computing and Computer-Assisted Intervention*, 2018, pp. 321–328.
- [35] S. Hajeb-Mohammad-Alipour, H. Rabbani, and M. R. Akhlaghi, "Diabetic retinopathy grading by digital curvelet transform," *Computational and Mathematical Methods in Med.*, vol. 2012, 2012.
- [36] N. Otsu, "A threshold selection method from gray-level histograms," *IEEE Intl. Conf. Systems, Man and Cyber.*, vol. 9, no. 1, pp. 62–66, 1979.
- [37] P. Burt and E. Adelson, "The Laplacian pyramid as a compact image code," *IEEE Trans. Comm.*, vol. 31, no. 4, pp. 532–540, 1983.
- [38] R. C. Gonzales and P. Wintz, *Digital Image Processing (2nd Ed.)*. Boston, MA, USA: Addison-Wesley Longman Publishing Co., Inc., 1987.
- [39] A. M. Mendonca and A. Campilho, "Segmentation of retinal blood vessels by combining the detection of centerlines and morphological reconstruction," *IEEE Trans. Med. Imaging*, vol. 25, no. 9, pp. 1200–1213, Sept 2006.
- [40] D. Bradley and G. Roth, "Adaptive thresholding using the integral image," *J. Graphics Tools*, vol. 12, no. 2, pp. 13–21, 2007.
- [41] Y. Gavet, M. Fernandes, and J.-C. Pinoli, "Quantitative evaluation of image registration techniques in the case of retinal images," *J. Electronic Imaging*, vol. 21, no. 2, pp. 1 – 8, 2012.
- [42] J. Nocedal and S. Wright, *Numerical optimization*. New York, NY, USA: Springer, 2006.
- [43] G. Borgefors, "Distance transformations in digital images," *Comp. Vis., Graphics and Image Proc.*, vol. 34, no. 3, pp. 344–371, 1986.
- [44] K. M. Kapp, *The gamification of learning and instruction*. San Francisco: Wiley, 2012.
- [45] C. I. Muntean, "Raising engagement in e-learning through gamification," in *Intl. Conf. Virtual Learning*, vol. 1, 2011.
- [46] I. Goodfellow, J. Pouget-Abadie, M. Mirza, B. Xu, D. Warde-Farley, S. Ozair, A. Courville, and Y. Bengio, "Generative adversarial nets," in *Adv. in Neural Info. Proc. Sys.*, 2014, pp. 2672–2680.
- [47] P. Isola, J.-Y. Zhu, T. Zhou, and A. A. Efros, "Image-to-image translation with conditional adversarial networks," in *IEEE Intl. Conf. Comp. Vision, and Pattern Recog.*, July 2017, pp. 1125–1134.
- [48] O. Ronneberger, P. Fischer, and T. Brox, "U-Net: Convolutional networks for biomedical image segmentation," in *Intl. Conf. Med. Image Computing and Computer-Assisted Intervention*, 2015, pp. 234–241.
- [49] Z. Zhou, M. M. R. Siddiquee, N. Tajbakhsh, and J. Liang, "UNet++: A nested U-Net architecture for medical image segmentation," in *Deep Learning in Med. Image Analysis*, 2018, pp. 3–11.
- [50] "PyTorch," accessed April 2020. [Online]. Available: <https://pytorch.org/>
- [51] M. E. Gegundez-Arias, A. Aquino, J. M. Bravo, and D. Marin, "A function for quality evaluation of retinal vessel segmentations," *IEEE Trans. Med. Imaging*, vol. 31, no. 2, pp. 231–239, Feb 2012.
- [52] "Implementation of CAL metric," accessed 27 Nov 2019. [Online]. Available: <https://github.com/ZengqiangYan/SkeletalSimilarityMetric/blob/master/CAL.m>
- [53] *Optos California Tech Sheet*, Optos, 2015. [Online]. Available: <https://www.optos.com/globalassets/www.optos.com/products/california/california-brochure.pdf>
- [54] J. Schindelin, I. Arganda-Carreras, E. Frise, V. Kaynig, M. Longair, T. Pietzsch, S. Preibisch, C. Rueden, S. Saalfeld, B. Schmid *et al.*, "Fiji: an open-source platform for biological-image analysis," *Nature methods*, vol. 9, no. 7, p. 676, 2012.
- [55] E. Pellegrini, G. Robertson, E. Trucco, T. J. MacGillivray, C. Lupascu, J. van Hemert, M. C. Williams, D. E. Newby, E. J. van Beek, and G. Houston, "Blood vessel segmentation and width estimation in ultra-wide field scanning laser ophthalmoscopy," *Biomed. Opt. Express*, vol. 5, no. 12, pp. 4329–4337, Dec 2014.
- [56] L. Ding, M. H. Bawany, A. E. Kuriyan, R. S. Ramchandran, C. C. Wykoff, and G. Sharma, "RECOVERY-FA19: Ultra-widefield fluorescein angiography vessel detection dataset," IEEE Dataport, 2019. [Online]. Available: <http://dx.doi.org/10.21227/m9yw-xs04>
- [57] K. Kim, E. S. Kim, and S.-Y. Yu, "Optical coherence tomography angiography analysis of foveal microvascular changes and inner retinal layer thinning in patients with diabetes," *British Journal of Ophthalmology*, vol. 102, no. 9, pp. 1226–1231, 2018.
- [58] T. Hastie, R. Tibshirani, and J. Friedman, *The elements of statistical learning: Data Mining, Inference, and Prediction*, 2nd ed. New York: Springer, 2009.
- [59] L. Ding, M. H. Bawany, A. E. Kuriyan, R. S. Ramchandran, C. C. Wykoff, and G. Sharma, "(Code Ocean capsule): Deep vessel segmentation for fluorescein angiography and evaluation," <https://doi.org/10.24433/CO.1133548.v1>, Apr 2020.