

One-Vote Veto: Semi-Supervised Learning for Low-Shot Glaucoma Diagnosis

Rui Fan¹, Senior Member, IEEE, Christopher Bowd², Nicole Brye, Mark Christopher, Robert N. Weinreb, David J. Kriegman, Fellow, IEEE, and Linda M. Zangwill³

Abstract—Convolutional neural networks (CNNs) are a promising technique for automated glaucoma diagnosis from images of the fundus, and these images are routinely acquired as part of an ophthalmic exam. Nevertheless, CNNs typically require a large amount of well-labeled data for training, which may not be available in many biomedical image classification applications, especially when diseases are rare and where labeling by experts is costly. This article makes two contributions to address this issue: 1) It extends the conventional Siamese network and introduces a training method for low-shot learning when labeled data are limited and imbalanced, and 2) it introduces a novel semi-supervised learning strategy that uses additional unlabeled training data to achieve greater accuracy. Our proposed multi-task Siamese network (MTSN) can employ any backbone CNN, and we demonstrate with four backbone CNNs that its accuracy with limited training data approaches the accuracy of backbone CNNs trained with a dataset that is 50 times larger. We also introduce One-Vote Veto (OVV) self-training, a semi-supervised learning strategy that is designed specifically for MTSNs. By taking both self-predictions and contrastive predictions of the unlabeled training data into account, OVV self-training provides additional pseudo labels for fine-tuning a pre-trained MTSN. Using a large (imbalanced) dataset with 66,715 fundus pho-

tographs acquired over 15 years, extensive experimental results demonstrate the effectiveness of low-shot learning with MTSN and semi-supervised learning with OVV self-training. Three additional, smaller clinical datasets of fundus images acquired under different conditions (cameras, instruments, locations, populations) are used to demonstrate the generalizability of the proposed methods.

Index Terms—Convolutional neural networks, glaucoma diagnosis, low-shot learning, semi-supervised learning.

I. INTRODUCTION

GLAUCOMA is a prevalent and debilitating disease that can lead to progressive and irreversible vision loss through optic nerve damage [1]. The global incidence of glaucoma was estimated at 64.3 million in 2013, and due to aging populations, this number is expected to rise to 111.8 million by 2040 [2]. Improvement in the management of glaucoma would have a major human and socio-economic impact [3]. Early identification and intervention would significantly reduce the economic burden of late-stage disease [4]. In addition, visual impairment in glaucoma patients has been associated with decreased physical activity and mental health [5], [6] and increased risk of motor vehicle accidents [7], [8].

With the recent advances in machine learning, convolutional neural networks (CNNs), trained via supervised learning, have shown promise in diagnosing glaucoma from fundus images (photographs of the back of eyes) [9]. However, this requires large amounts of empirical data for supervised training [10]. In this study, we use 66,715 fundus photographs from the Ocular Hypertension Treatment Study (OHTS) [11], [12], [13], which is a 22-site multi-center, longitudinal (phase 1 and 2, 1994–2008) randomized clinical trial of 1,636 subjects (3,272 eyes). The primary goal of the OHTS was to determine if topical ocular hypotensive medications could delay or prevent the onset of glaucoma in eyes with high intraocular pressure [11]. Conversion to glaucoma was decided by a masked endpoint committee of three glaucoma specialists using fundus photographs and visual fields. Owing to its well-characterized ground-truth labels, the OHTS dataset provides us a basis to explore an effective way of training CNNs to diagnose glaucoma with low-shot learning when only a small quantity of labeled data is available, and/or semi-supervised learning when raw data is abundant, but labeling resources are scarce, costly, require strong expertise, or are just unavailable.

Manuscript received 3 May 2023; accepted 15 August 2023. Date of publication 23 August 2023; date of current version 30 November 2023. This work was supported in part by the Fundamental Research Funds for the Central Universities; in part by the Awards from the National Eye Institute under Grant EY027510, Grant R214278211, Grant P30EY022589, Grant K99EY030942, Grant EY026574, Grant EY11008, and Grant EY19869; in part by the National Center on Minority Health and Health Disparities, National Institutes of Health, under Grant R01EY09341 and Grant R01EY09307; in part by the Horncrest Foundation; in part by the Department of Ophthalmology and Visual Sciences, Washington University in St. Louis; in part by the NIH Vision Core under Grant P30EY02687; in part by Merck Research Laboratories, Pfizer Inc., White House Station, NJ, USA; and in part by the Unrestricted Grants from Research to Prevent Blindness Inc., New York, NY, USA. (Corresponding author: Linda M. Zangwill.)

Rui Fan is with the College of Electronics and Information Engineering, Shanghai Research Institute for Intelligent Autonomous Systems, State Key Laboratory of Intelligent Autonomous Systems, and the Frontiers Science Center for Intelligent Autonomous Systems, Tongji University, Shanghai 201804, China (e-mail: rui.fan@ieee.org).

Christopher Bowd, Nicole Brye, Mark Christopher, Robert N. Weinreb, and Linda M. Zangwill are with the Hamilton Glaucoma Center, Viterbi Family Department of Ophthalmology, and the Shiley Eye Institute, University of California at San Diego, La Jolla, CA 92093 USA (e-mail: cbowd@health.ucsd.edu; nbrye@health.ucsd.edu; mac157@health.ucsd.edu; rweinreb@ucsd.edu; lzungwill@health.ucsd.edu).

David J. Kriegman is with the Department of Computer Science and Engineering, University of California at San Diego, La Jolla, CA 92093 USA (e-mail: kriegman@ucsd.edu).

Digital Object Identifier 10.1109/TMI.2023.3307689

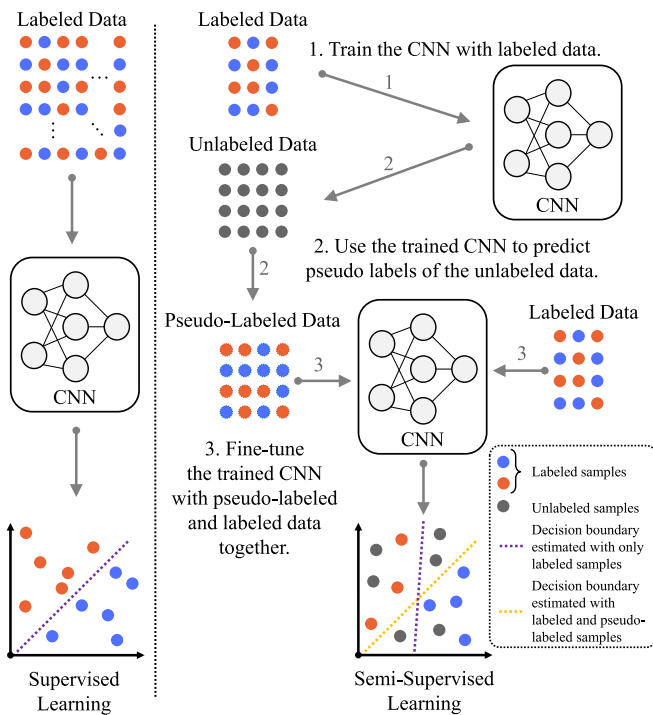


Fig. 1. Supervised learning v.s. semi-supervised learning.

However, as shown in Fig. 1, conventional semi-supervised learning approaches typically require a reliable pre-trained CNN (using a small sample) as prior knowledge, which is often challenging due to the over-fitting problem. Moreover, there is also a strong motivation to design a feasible semi-supervised learning strategy capable of determining confident predictions and generating pseudo labels for unlabeled data. We focus specifically on fundus images and glaucoma diagnosis in this article because we have sufficient data to accurately characterize the effectiveness of our methods. The same techniques could also be applied to tasks where there is limited data, such as rare diseases or where limited labels are available (e.g., asthma and diabetes prediction from fundus images). Therefore, this article aims to answer the following questions:

- 1) Can a CNN be developed to accurately diagnose glaucoma, compared to the expert graders of the OHTS? Will the model be generalizable to other datasets?
- 2) Is it necessary to train CNNs with thousands of labeled fundus images to diagnose glaucoma, or can diagnosis be achieved using only one image per patient (approximately 1.1K fundus images in the OHTS training set)?
- 3) Can the performance of a CNN trained using a small sample be improved further by fine-tuning it with additional unlabeled training data?

To answer these questions, we first evaluate the performance of state-of-the-art (SoTA) glaucoma diagnosis algorithms, including six supervised learning algorithms [3], [14], [15], [16], [17], [18], one low-shot learning algorithm [19], and two semi-supervised learning algorithms [20], [21], on the OHTS dataset. Their generalizabilities are further validated on three additional clinical datasets of fundus images: (a)




Fig. 2. An illustration of our MTSN for joint learning of fundus image similarity measurement (primary task) and glaucoma diagnosis (secondary task) in a low-shot manner.

ACRIMA (Spain) [22], (b) Large-Scale Attention-Based Glaucoma (LAG, China) [9], and (c) the UCSD-based Diagnostic Innovations in Glaucoma Study and African Descent and Glaucoma Evaluation Study (DIGS/ADAGES, US) [23].

Furthermore, we propose a novel extension to the conventional Siamese network, referred to as the **Multi-Task Siamese Network (MTSN)**, as depicted in Fig. 2. By minimizing a novel **Combined Weighted Cross-Entropy (CWCE) Loss**, the MTSN can simultaneously perform two tasks: measuring the similarity of a given pair of images (primary task) and classifying them as healthy or glaucoma (secondary task). With a small training set of approximately 1.1K fundus images, we explore the feasibility of training an MTSN for glaucoma diagnosis. Although the MTSN may not provide complementary information, it effectively performs a type of “data augmentation” by generating $C(N, 2)$ pairs of fundus images for training instead of using N independent fundus images. The visual features learned from these two tasks prove to be more informative for glaucoma diagnosis when the training set is small. Our experimental results demonstrate that the MTSN greatly reduces over-fitting and achieves an accuracy on a small training set comparable to a large training set, which contains approximately 53K fundus images.

Moreover, we propose a novel semi-supervised learning strategy, referred to as **One-Vote Veto (OVV) Self-Training**, which generates reliable pseudo labels for the unlabeled training data and incorporates them into the labeled training data to fine-tune the MTSN for improved performance and generalizability. Our extensive experiments show that the MTSN fine-tuned with OVV self-training achieves similar performance to the corresponding backbone CNN trained via supervised learning on the OHTS dataset, and achieves higher area under the receiver operating characteristic curve (AUROC) scores on the additional fundus image datasets. The fine-tuned MTSN also outperforms SoTA semi-supervised glaucoma diagnosis approaches [20], [21], and in some cases, even outperforms SoTA supervised approaches. Additionally, we compare our proposed OVV self-training approach with four SoTA general-purpose semi-supervised learning methods, including FreeMatch [24], SoftMatch [25], FixMatch [26], and FlexMatch [27], all of which utilize vision Transformer [28] as their backbone network. The results demonstrate that

our proposed OVV self-training approach outperforms these methods on the OHTS dataset and demonstrates better generalizability on three additional fundus image test sets.

We also conduct two additional few-shot biomedical image classification experiments (chest X-ray image classification [29], [30] and lung histopathological image classification [31]) to further validate the effectiveness of the MTSN on other types of image data. The promising results indicate that our proposed algorithms have the potential to solve a variety of biomedical image classification problems.

II. RELATED WORKS

Most SoTA glaucoma diagnosis algorithms are developed based on supervised fundus image classification. For example, Judy et al. [16] trained an AlexNet [32] to diagnose glaucoma. As VGG architectures [33] can learn more complicated image features than AlexNet, Gómez-Valverde et al. [15] employed a VGG-19 [33] model to diagnose glaucoma. Nevertheless, VGG architectures [33] consist of hundreds of millions of parameters, making them very memory-consuming. In contrast, GoogLeNet [34] and Inception-v3 [35] have lower computational complexities. Hence, Ahn et al. [36] and Li et al. [14] utilized transfer learning to re-train an Inception-v3 [35] model (pre-trained on the ImageNet [37] database) for glaucoma diagnosis, while Serener and Serte [17] re-trained a pre-trained GoogLeNet [34] model to diagnose glaucoma. However, with the increase of network depth, accuracy gets saturated and then degrades rapidly due to vanishing gradients [38]. To tackle this problem, the residual neural network (ResNet) [38] was developed. Due to its robustness, ResNet-50 [38] has been extensively used for biomedical image analysis, and it is a popular choice [3], [39], [40], [41], [42], [43] for fundus image classification. Additionally, developing low-cost and real-time embedded glaucoma diagnosis systems [18], [44], [45], *e.g.*, based on MobileNet-v2 [46], for mobile devices is also an emerging area.

Machine/deep learning has achieved compelling performance in data-intensive applications, but it is often challenging for these algorithms to yield comparable performance when only a limited amount of labeled training data is available [47]. Low-shot and semi-supervised learning can address these issues. Unfortunately, they are rarely discussed in the field of glaucoma diagnosis. To the best of our knowledge, [19] is the only published low/few-shot glaucoma diagnosis algorithm. This algorithm employs a conventional Siamese network to compare two groups of (negative and positive) fundus images. The Siamese network utilizes two identical CNNs to learn visual embeddings. A bi-directional long short-term memory [48] component is then trained over the CNN outputs for glaucoma diagnosis. However, the training process is complicated since different types of losses are minimized, and the achieved glaucoma diagnosis results are unsatisfactory since each sub-network is only fed with one type of fundus images (either negative or positive). The lack of same-class comparisons leads to a performance bottleneck when compared to the MTSN proposed in this article.

A thorough search of the relevant literature yielded only two published studies on semi-supervised learning specifically for glaucoma diagnosis [20], [21]. Díaz Pinto et al. [21] utilized a deep convolutional generative adversarial network (DCGAN) [49] for semi-supervised learning of glaucoma diagnosis, where the discriminator is trained to classify healthy and glaucomatous optic neuropathy (GON) fundus images, while also distinguishing between real and fake fundus images. The classifier for the former task is then employed for glaucoma diagnosis. On the other hand, Al Ghamdi et al. [20] developed a glaucoma diagnosis approach based on self-training [50], which is a typical semi-supervised learning approach that uses a pre-trained model (typically yielded via supervised learning) to produce pseudo labels of the unlabeled data. However, producing reliable pseudo labels is a significant challenge in self-training, and the pseudo labels generated by a single pre-trained CNN are usually not trustworthy enough for CNN fine-tuning [51]. Additionally, training a reliable pre-trained classifier with only a small amount of labeled data is notably demanding. In this article, we combine semi-supervised learning with low-shot learning to address these issues using glaucoma diagnosis as an example case. Specifically, our proposed OVV self-training strategy, as discussed in Sect. III-B, is inspired by the mechanism of *learning with external memory* (LwEM), used in low-shot learning [52], where the labels of unlabeled training data are predicted by a classifier trained via low-shot learning on a small collection of fundus images with ground-truth labels.

III. METHODOLOGY

A. Multi-Task Siamese Network

As illustrated in Fig. 1, conventional semi-supervised learning methods initialize a network by pre-training it with a small number of fundus images for glaucoma diagnosis. However, we observed that such approaches are highly sensitive to noise. As a result, we design a novel MTSN specifically for our semi-supervised learning approach, which requires not only predicting the label of a given fundus image but also determining the similarity between a pair of given fundus images to generate pseudo labels through a voting process.

Conventional Siamese networks have become a common choice for metric learning and few/low-shot image recognition tasks [53]. These networks comprise two identical sub-networks, as depicted in Fig. 2. Each pair of fundus images \mathbf{x}_i and \mathbf{x}_j are separately fed into these sub-networks, producing two 1D embeddings (features) \mathbf{h}_i and \mathbf{h}_j , respectively. Another 1D embedding $\mathbf{h}_{i,j}$ is generated by $\Phi(\cdot)$. $\mathbf{h}_{i,j}$ is then passed through a fully connected (FC) layer to produce a scalar $q(\mathbf{x}_i, \mathbf{x}_j) \in [0, 1]$ indicating the similarity between \mathbf{x}_i and \mathbf{x}_j . If \mathbf{x}_i and \mathbf{x}_j are dissimilar, $q(\mathbf{x}_i, \mathbf{x}_j)$ approaches 1, and vice versa. The ground-truth labels of \mathbf{x}_i and \mathbf{x}_j are represented by $y_i \in \{0, 1\}$ and $y_j \in \{0, 1\}$, respectively, where 0 denotes a healthy image, and 1 denotes a GON image.

However, a conventional Siamese network can only determine whether \mathbf{x}_i and \mathbf{x}_j belong to the same category, rather than predicting their independent categories. A straightforward solution is to connect \mathbf{h}_i and \mathbf{h}_j to separate FC layers, producing two scalars $p(\mathbf{x}_i)$ and $p(\mathbf{x}_j)$ indicating the probabilities

that \mathbf{x}_i and \mathbf{x}_j are GON images, respectively. Refer to Fig. 2 and note that the two FC layers connected to \mathbf{h}_i and \mathbf{h}_j use the same weights. In this article, we refer to the network architecture in Fig. 2 as an MTSN, which can simultaneously measure the similarity of a given pair of fundus images and classify them as either healthy or GON. These two tasks are dependent yet not directly deducible from one another. A well-trained glaucoma diagnosis network can be employed to compare differences between given pairs of fundus images, but a well-trained fundus image similarity measurement network cannot directly output the category of a given fundus image.

In addition, the visual features learned from the primary and secondary tasks are distinct from one another. For the primary task, the network learns the visual features to classify same-class and different-class fundus image pairs. On the other hand, for the secondary task, the network learns the visual features to classify GON and healthy fundus images. Although this network architecture may not provide complementary information, it effectively performs a type of “data augmentation” by producing $C(N, 2)$ pairs of fundus images for training, rather than using N independent fundus images. The visual features learned from these two tasks prove to be more informative for glaucoma diagnosis when the training set is small. Furthermore, multi-task learning is effective because requiring an algorithm to perform well on a related task induces regularization, which can be superior to uniform complexity penalization for preventing over-fitting. This idea has been explored in many Siamese neural network works, such as [54], [55] and [56].

In this article, we use n_0 and n_1 to respectively denote the numbers of healthy and GON fundus images used to train the MTSN, with $n = n_0 + n_1$. n_0 is usually much greater than n_1 , because there are fewer patients with glaucomatous disease than healthy patients, resulting in a severely imbalanced dataset. Therefore, the MTSN is trained by minimizing a CWCE loss as follows:

$$\mathcal{L} = \mathcal{L}_{\text{sim}} + \lambda \mathcal{L}_{\text{cla}}, \quad (1)$$

where

$$\begin{aligned} \mathcal{L}_{\text{sim}} &= -\frac{n_0(n_0 - 1) + n_1(n_1 - 1)}{n(n - 1)} |y_i - y_j| \log(q(\mathbf{x}_i, \mathbf{x}_j)) \\ &\quad - \frac{2n_0n_1}{n(n - 1)} (1 - |y_i - y_j|) \log(1 - q(\mathbf{x}_i, \mathbf{x}_j)), \end{aligned} \quad (2)$$

$$\begin{aligned} \mathcal{L}_{\text{cla}} &= -\frac{1}{n} \left(n_0 (y_i \log(p(\mathbf{x}_i)) + y_j \log(p(\mathbf{x}_j))) \right. \\ &\quad \left. + n_1 ((1 - y_i) \log(1 - p(\mathbf{x}_i)) + (1 - y_j) \log(1 - p(\mathbf{x}_j))) \right). \end{aligned} \quad (3)$$

The hyper-parameter λ balances the primary task loss \mathcal{L}_{sim} and the secondary task loss \mathcal{L}_{cla} . The choice of λ and $\Phi(\cdot)$ is discussed in Sect. IV-B. The motivations for using such a CWCE loss function instead of the commonly used triplet loss [57] or contrastive loss [58] to train the MTSN are:

- 1) Most datasets for rare disease diagnosis are imbalanced. As detailed in Sect. IV-A, the OHTS training set is severely imbalanced, with 50,208 healthy images and only 2,416 GON images for supervised learning, and 995 healthy images and 152 GON images for low-shot learning. Learning from such an imbalanced dataset without weights on different classes can result in many incorrect predictions, with most GON images likely to be predicted as healthy images. To address this issue, a higher weight should be used for the minority class to prevent the CNN from predicting all fundus images as the majority class.
- 2) In multi-task learning, weighing different types of losses, such as regression and classification, is typically challenging [59]. Assigning an incorrect weight may cause one task to perform poorly, even when other tasks converge to satisfactory results. Therefore, formulating \mathcal{L}_{sim} as a weighted cross-entropy loss function is a simple but effective solution. However, due to the dataset imbalance problem, the cross-entropy losses have to be weighted.
- 3) As shown in Fig. 3, OVV self-training requires both labels and probabilities (of being GON images), predicted by a pre-trained model, to produce pseudo labels for unlabeled data. Such network architecture and training loss can efficiently and effectively provide both “self-predicted” and “contrastively-predicted” labels and probabilities, as described in Sect. III-B.

It should be noted here that using the fundus images from the same patient as an image pair for MTSN training is not necessary.

B. One-Vote Veto Self-Training

As discussed in Sect. I, self-training aims to improve the performance of a pre-trained model by incorporating reliable predictions of the unlabeled data to obtain useful additional information that can be used for model fine-tuning. A feasible strategy to determine such reliable predictions is, therefore, key to the success of self-training [60].

In conventional semi-supervised learning algorithms, a pre-trained image classification model (obtained through supervised learning) can be fine-tuned by assessing the reliability of unlabeled images. To determine the reliability of an unlabeled image, its probability distribution for the most likely class is compared to a pre-determined threshold. If the probability surpasses this threshold, the prediction is considered a pseudo label. Subsequently, the image and its corresponding pseudo label are utilized to fine-tune the pre-trained model.

However, relying solely on probability distributions to generate pseudo labels is often insufficient [50]. Drawing inspiration from LwEM [61], we introduce *One-Vote Veto self-training* in this paper, as illustrated in Fig. 3.¹ Similar to LwEM [61], we use a collection of m reference (labeled) fundus images $\{\mathbf{x}_1^r, \dots, \mathbf{x}_m^r\} \in \mathcal{X}^r$ to provide “*contrastive predictions*” to the target (unlabeled) fundus images $\{\mathbf{x}_1^t, \dots, \mathbf{x}_m^t\} \in \mathcal{X}^t$. The contrastive predictions subsequently

¹The superscripts r and t denote “reference” and “target”, respectively.

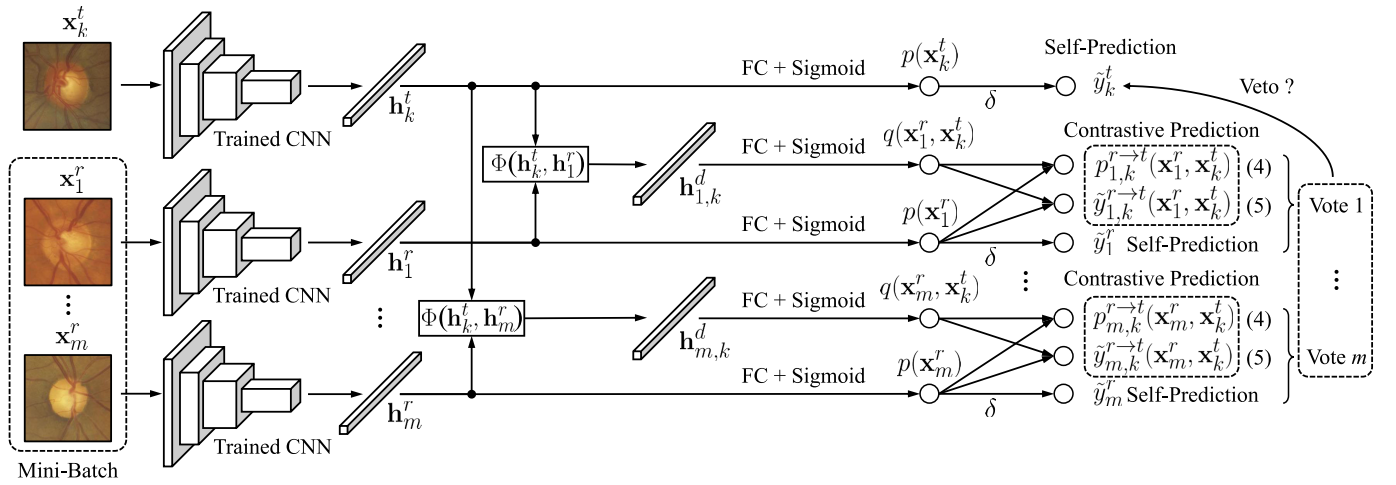


Fig. 3. An illustration of our One-Vote Veto Self-Training strategy. $\mathbf{h}_{1,k}^d$ and $\mathbf{h}_{m,k}^d$ are two 1D embeddings, followed by an FC layer to produce scalars indicating the similarities between the given pairs of reference and target fundus images. The reference fundus images having ground-truth labels are used to train the MTSN by minimizing (1). The contrastive predictions are obtained using (4) and (5). The pseudo labels of the target fundus photographs are generated using One-Vote Veto Self-Training strategy, as detailed in Algorithm 1.

vote to veto the unreliable “self-predictions” $\{\tilde{y}_1^t, \dots, \tilde{y}_m^t\}$ produced by the MTSN. Our OVV self-training is detailed in Algorithm 1, where the target model updates its parameters during self-training but the reference model does not.

When fine-tuning an MTSN pre-trained through low-shot learning, each mini-batch contains a discrete set of m reference fundus images $\{\mathbf{x}_1^r, \dots, \mathbf{x}_m^r\} \in \mathcal{X}^r$, their ground-truth labels $\{y_1^r, \dots, y_m^r\} \in \mathcal{Y}^r$, and an equal number of m target fundus images $\{\mathbf{x}_1^t, \dots, \mathbf{x}_m^t\} \in \mathcal{X}^t$ without labels. \mathbf{h}_k^r and \mathbf{h}_k^t represent the 1D embeddings learned from \mathbf{x}_k^r and \mathbf{x}_k^t ($k \in [1, m] \cap \mathbb{Z}$), respectively. Given a pair of reference and target fundus images, \mathbf{x}_l^r and \mathbf{x}_k^t , the pre-trained MTSN can “self-predict”:

- the scalars $p(\mathbf{x}_l^r)$ and $p(\mathbf{x}_k^t)$ which indicate the probabilities that \mathbf{x}_l^r and \mathbf{x}_k^t are GON images, respectively;
- their labels $\tilde{y}_l^r = \delta(p(\mathbf{x}_l^r))$ and $\tilde{y}_k^t = \delta(p(\mathbf{x}_k^t))$ using its fundus image classification functionality ($\delta(p) = 1$ when $p > 0.5$, and $\delta(p) = 0$ otherwise).

$p(\mathbf{x}_l^r)$ is then used to determine whether the reference fundus image \mathbf{x}_l^r is qualified to veto unreliable predictions. If $|p(\mathbf{x}_l^r) - y_l^r| > \kappa_2$, its vote will be omitted, where κ_2 is a threshold used to select qualified reference fundus images (step 6 in Algorithm 1). In the meantime, the pre-trained MTSN can also “contrastively-predict” the scalar

$$p_{l,k}^{r \rightarrow t}(\mathbf{x}_l^r, \mathbf{x}_k^t) = |p(\mathbf{x}_l^r) - q(\mathbf{x}_l^r, \mathbf{x}_k^t)| \quad (4)$$

indicating the GON probability as well as the label

$$\tilde{y}_{l,k}^{r \rightarrow t}(\mathbf{x}_l^r, \mathbf{x}_k^t) = |\delta(p(\mathbf{x}_l^r)) - \delta(q(\mathbf{x}_l^r, \mathbf{x}_k^t))| \quad (5)$$

of \mathbf{x}_k^t from \mathbf{x}_l^r using its input similarity measurement functionality.² To determine the reliability of $\tilde{y}_{l,k}^{r \rightarrow t}$ and whether it can be used as the pseudo label of \mathbf{x}_k^t , all the reference fundus images $\{\mathbf{x}_1^r, \dots, \mathbf{x}_m^r\} \in \mathcal{X}^r$ in the mini-batch are used to provide additional judgements. Each pair of contrastively-predicted scalar (indicating GON probability) and label form a vote ($p_{l,k}^{r \rightarrow t}(\mathbf{x}_l^r, \mathbf{x}_k^t)$, $\tilde{y}_{l,k}^{r \rightarrow t}(\mathbf{x}_l^r, \mathbf{x}_k^t)$). With all votes collected from

²In rare cases, $\tilde{y}_{l,k}^{r \rightarrow t}(\mathbf{x}_l^r, \mathbf{x}_k^t)$ might not be equivalent to $\delta(p_{l,k}^{r \rightarrow t}(\mathbf{x}_l^r, \mathbf{x}_k^t))$.

Algorithm 1 One-Vote Veto Self-Training Strategy.

Data: Reference fundus images \mathcal{X}^r and their labels \mathcal{Y}^r , and target fundus images \mathcal{X}^t

- 1 **while** Training **do**
- 2 Given a mini-batch consisting of $\{\mathbf{x}_1^r, \dots, \mathbf{x}_m^r\} \in \mathcal{X}^r$, $\{y_1^r, \dots, y_m^r\} \in \mathcal{Y}^r$ and $\{\mathbf{x}_1^t, \dots, \mathbf{x}_m^t\} \in \mathcal{X}^t$;
- 3 Initialize an empty set \mathcal{P} to store reliable target fundus images and their pseudo labels;
- 4 **for** Each target fundus image \mathbf{x}_k^t **do**
- 5 **if** The self-prediction of \mathbf{x}_k^t is reliable **then**
- 6 **for** Any qualified reference image \mathbf{x}_l^r **do**
- 7 Compute the self-prediction and contrastive prediction;
- 8 **end**
- 9 **if** The criteria to generate pseudo labels are satisfied **then**
- 10 Update the set \mathcal{P} ;
- 11 **end**
- 12 **end**
- 13 **end**
- 14 Fine-tune the target model using $\text{unique}(\mathcal{P})$;
- 15 **end**
- 16 **if** The target model outperforms the reference model **then**
- 17 Update the reference model parameters;
- 18 **end**

the qualified reference fundus images, the OVV self-training algorithm determines whether \tilde{y}_k^t should be used as the pseudo label for \mathbf{x}_k^t based on the following criteria (step 9 in Algorithm 1):

- Identical to the process of determining qualified reference fundus images, if any $p_{l,k}^{r \rightarrow t}(\mathbf{x}_l^r, \mathbf{x}_k^t)$ ($l \in [1, m] \cap \mathbb{Z}$) or $p(\mathbf{x}_k^t)$ is not close to either 0 (healthy) or 1 (GON), as evaluated by the threshold κ_2 , \tilde{y}_k^t will not be assigned to \mathbf{x}_k^t .
- If a minority of more than κ_1 qualified reference fundus images disagree with the majority of the qualified reference fundus images, \tilde{y}_k^t will not be assigned to \mathbf{x}_k^t .

As discussed in Sect. IV, $\kappa_1 = 0$ (all qualified reference images vote for the same category) achieves the best overall performance. Therefore, the aforementioned strategy is named “*One-Vote Veto*” in this paper. Since each target fundus image is required to be compared with all the reference fundus images in the same mini-batch, the proposed self-training strategy has a computational complexity of $\mathcal{O}(n^2)$, which is relatively memory-consuming. The reliable target fundus images and their pseudo labels are then included into the low-shot training data to fine-tune the pre-trained MTSN with supervised learning by minimizing a CWCE loss. The OVV self-training performance with respect to different κ_1 , κ_2 , and m values is discussed in Sect. IV.

IV. EXPERIMENTS

A. Datasets and Experimental Setups

The datasets utilized in our experiments were collected by various clinicians in different countries using distinct fundus cameras. The ACRIMA [22] and LAG [9] datasets are publicly available, while the OHTS [11], [12] and DIGS/ADAGES [23] datasets are available upon request after appropriate data use agreements are initiated. Their details are as follows:

- The **OHTS** [11], [12] is the only multi-center longitudinal study that has precise information on the dates/timing of the development of glaucoma (the enrolled subjects did not have glaucoma at study entry) using standardized assessment criteria by an independent Optic Disc Reading Center and confirmed by three glaucoma specialist endpoint committee members. In our experiments, a square region centered on the optic nerve head was first extracted from each raw fundus image using a well-trained DeepLabv3+ [62] model. A small part of the raw data are stereoscopic fundus images, each of which was split to produce two individual fundus images. Through this image pre-processing approach, a total of 74,678 fundus images were obtained. Moreover, ENPOAGDISC (endpoint committee attributable to primary open angle glaucoma based on optic disc changes from photographs) [3] labels are used as the classification ground truth. The fundus images are divided into a training set (50,208 healthy images and 2,416 GON images), a validation set (7,188 healthy images and 426 GON images), and a test set (13,780 healthy images and 660 GON images) by participant. Splitting by participant (instead of by image) ensured that the validation and test sets did not contain images from any eyes or individuals used to train the model. More details on dataset preparation and baseline supervised learning experiments are provided in our recent publications [3], [10]. Additionally, we select one image (from only one eye) from each patient in the training set to create the low-shot training set (995 healthy images and 152 GON images). If both eyes of a patient do not convert to glaucoma in the study, the first captured fundus photograph is selected. If either eye of a patient converts to glaucoma in the study, the first glaucoma fundus photograph is selected.




Fig. 4. Comparisons of dataset visualizations produced by t-SNE, where \bullet and \bullet represent the healthy and GON images in the OHTS test set, respectively; \bullet and \bullet in (a) represent the healthy and GON images in the ACRIMA dataset, respectively; \bullet and \bullet in (b) represent the healthy and GON images in the LAG dataset, respectively; \bullet and \bullet in (c) represent the healthy and GON images in the DIGS/ADAGES test set, respectively.

- The **ACRIMA** [22] dataset consists of 309 healthy images and 396 GON images. It was collected as part of an initiative by the government of Spain. Classification was based on the review by a single experienced glaucoma expert. Images were excluded if they did not provide a clear view of the optic nerve head region [43].
- The **LAG** [9] dataset contains 3,143 healthy images and 1,711 GON images,³ collected by Beijing Tongren Hospital. Similar to the OHTS dataset, we also use the well-trained DeepLabv3+ [62] model to extract a square region centered on the optic nerve head from each fundus image.
- The **DIGS** and **ADAGES** [23] are longitudinal studies designed to detect and monitor glaucoma based on optical imaging and visual function testing that, when combined, have generated tens of thousands of test results from over 4,000 healthy, glaucoma suspect, or glaucoma eyes. In our experiments, we utilize the DIGS/ADAGES test set (5,184 healthy images and 4,289 GON images) to evaluate the generalizability of our proposed methods.

Visualizations of the four test sets using t-SNE [63] are provided in Fig. 4. Since healthy and GON images are distributed similarly between the OHTS and LAG datasets, we expect models to perform similarly on these datasets. Dissimilar distributions in the ACRIMA and DIGS/ADAGES datasets led us to believe the performance of models on these datasets would be somewhat worse. Using these four datasets, we conduct three experiments:

- 1) **Supervised learning experiment:** We employ transfer learning [64] to train ResNet-50 [38], MobileNet-v2 [46], DenseNet [65], and EfficientNet [66] (pre-trained on the ImageNet database [37]), on the entire OHTS training set (including $\sim 53K$ fundus images). The best-performing models are selected using the OHTS validation set. Their performance is subsequently evaluated on the OHTS test set, the ACRIMA dataset, the LAG dataset, and the DIGS/ADAGES test set.
- 2) **Low-shot learning experiment:** The four pre-trained models mentioned above are used as the MTSN backbones and trained on the OHTS low-shot training set

³The number of fundus images being published is fewer than what was reported in publication [9].

TABLE I
COMPARISON BETWEEN SUPERVISED AND LOW-SHOT LEARNING (BOTH ON THE LOW-SHOT OHTS TRAINING SET).
 λ IS SET TO 0.3. THE BEST RESULTS ARE SHOWN IN BOLD FONT

Test set	Experiments	Training strategy	ResNet-50			MobileNet-v2		
			Accuracy (%)	F1-score (%)	AUROC	Accuracy (%)	F1-score (%)	AUROC
ACRIMA [22]	Baseline	Supervised learning	57.163	56.734	0.625	73.333	78.341	0.779
	EWAD	Low-shot learning	67.092	63.175	0.758	70.355	67.797	0.820
	EWSD	Low-shot learning	49.504	25.523	0.437	66.241	59.107	0.823
LAG [9]	Baseline	Supervised learning	64.318	57.445	0.714	65.122	63.268	0.781
	EWAD	Low-shot learning	79.028	65.908	0.841	79.007	69.482	0.843
	EWSD	Low-shot learning	78.039	68.179	0.826	78.430	61.830	0.846
DIGS/ADAGES [23]	Baseline	Supervised learning	59.639	59.708	0.648	61.478	66.128	0.669
	EWAD	Low-shot learning	67.745	60.754	0.743	69.176	68.145	0.748
	EWSD	Low-shot learning	65.736	55.722	0.700	68.120	64.028	0.740

TABLE II
EVALUATION OF OUR OVV SELF-TRAINING W.R.T. DIFFERENT κ_1 , κ_2 , AND m . THE BEST RESULTS ARE SHOWN IN BOLD FONT. \uparrow INDICATES THAT SEMI-SUPERVISED LEARNING OUTPERFORMS LOW-SHOT LEARNING

Dataset	κ_1	κ_2	m	ResNet-50			MobileNet-v2		
				Accuracy (%)	F1-score (%)	AUROC	Accuracy (%)	F1-score (%)	AUROC
OHTS [11], [12]	0	0.01	20	91.415 \uparrow	41.148 \uparrow	0.898 \uparrow	90.609 \uparrow	36.960 \uparrow	0.887 \uparrow
	0	0.01	15	92.113 \uparrow	41.316 \uparrow	0.898 \uparrow	93.470 \uparrow	36.976 \uparrow	0.893 \uparrow
	0	0.01	10	94.199 \uparrow	43.759 \uparrow	0.890 \uparrow	93.742 \uparrow	37.779 \uparrow	0.863 \uparrow
	0	0.10	20	90.516 \uparrow	38.139 \uparrow	0.898 \uparrow	88.825 \uparrow	34.351 \uparrow	0.878 \uparrow
	2	0.01	20	90.717 \uparrow	35.818 \uparrow	0.885 \uparrow	93.463 \uparrow	35.204 \uparrow	0.862 \uparrow
	2	0.10	20	92.656 \uparrow	32.017 \uparrow	0.851 \downarrow	90.772 \uparrow	32.616 \uparrow	0.858 \uparrow
	4	0.01	20	92.610 \uparrow	29.668 \downarrow	0.854 \downarrow	92.268 \uparrow	31.852 \uparrow	0.854 \uparrow
	4	0.10	20	92.672 \uparrow	28.463 \downarrow	0.842 \downarrow	90.803 \uparrow	32.690 \uparrow	0.859 \uparrow
ACRIMA [22]	0	0.01	20	59.858 \downarrow	49.192 \downarrow	0.775 \uparrow	72.340 \uparrow	70.229 \uparrow	0.840 \uparrow
	0	0.01	15	60.426 \downarrow	49.365 \downarrow	0.751 \downarrow	63.404 \downarrow	54.895 \downarrow	0.814 \downarrow
	0	0.01	10	54.610 \downarrow	35.743 \downarrow	0.721 \downarrow	61.986 \downarrow	51.273 \downarrow	0.826 \uparrow
LAG [9]	0	0.01	20	80.882 \uparrow	66.304 \uparrow	0.881 \uparrow	79.625 \uparrow	65.262 \downarrow	0.851 \uparrow
	0	0.01	15	76.864 \downarrow	56.252 \downarrow	0.825 \downarrow	76.670 \downarrow	55.906 \downarrow	0.841 \downarrow
	0	0.01	10	76.638 \downarrow	56.518 \downarrow	0.826 \downarrow	75.834 \downarrow	51.748 \downarrow	0.866 \uparrow
DIGS/ADAGES [23]	0	0.01	20	67.813 \uparrow	58.315 \downarrow	0.763 \uparrow	69.653 \uparrow	63.258 \downarrow	0.777 \uparrow
	0	0.01	15	63.045 \downarrow	44.727 \downarrow	0.753 \uparrow	66.383 \downarrow	52.888 \downarrow	0.773 \uparrow
	0	0.01	10	61.819 \downarrow	41.340 \downarrow	0.727 \downarrow	63.965 \downarrow	44.198 \downarrow	0.789 \uparrow

(containing 1,147 images) to validate the effectiveness of our proposed low-shot glaucoma diagnosis algorithm. The validation and testing procedures are identical to those in the supervised learning experiment.

- 3) **Semi-supervised learning experiment:** The MTSNs trained on the low-shot training set are fine-tuned on the entire OHTS training set without using additional ground-truth labels. The fine-tuned MTSNs are referred to as MTSN+OVV. The validation and testing procedures are identical to those in the supervised learning experiment.

The fundus images are resized to 224×224 pixels. The initial learning rate is set to 0.001, which decays gradually after the 100th epoch. Due to the dataset imbalance problem, F1-score is utilized to select the best-performing models during the validation stage. Moreover, we adopt an early stopping mechanism during the validation stage to reduce over-fitting (the training will be terminated if the achieved F1-score has not increased for 10 epochs). We use three metrics: (1) accuracy, (2) F1-score, and (3) AUROC to quantify the performances of the trained models. While accuracy is generally reported in image classification papers, F1-score and AUROC are more

comprehensive and reasonable evaluation metrics when the dataset is severely imbalanced.

B. Ablation Study and Threshold Selection

We set λ in (1) to 0.1, 0.2, 0.3, 0.4, and 0.5, respectively, and compare the MTSN performance when Φ computes the element-wise absolute difference (EWAD) and element-wise squared difference (EWSD), respectively. The comparisons in terms of F1-score and AUROC on the OHTS test set are provided in Fig. 5. It can be seen that the MTSN achieves the best overall performance when $\lambda = 0.3$. This is reasonable, as a higher λ weighs more on the image classification task, easily resulting in over-fitting. Additionally, the MTSN in which $\Phi(\cdot)$ computes the EWSD between \mathbf{h}_i and \mathbf{h}_j performs better when using ResNet-50 as the backbone CNN but slightly worse when using MobileNet-v2 as the backbone CNN. Therefore, we further evaluate their generalizability on three additional test sets, as shown in Table I. When $\Phi(\cdot)$ computes the EWAD, the MTSN generally performs better or very similarly on the additional test sets, especially when testing the MTSN assembled with ResNet-50 on the ACRIMA dataset. EWAD is, therefore, used in the following experiments. Furthermore,

TABLE III
AUROC (SHOWN ALONG WITH 95% CI) AND TRAINING TIME t (MIN) PER EPOCH OF SUPERVISED, LOW-SHOT, AND SEMI-SUPERVISED GLAUCOMA DIAGNOSIS

Backbone	Method	Training strategy	OHTS [11], [12]	ACRIMA [22]	LAG [9]	DIGS/ADAGES [23]	t (min)
ResNet-50	Baseline	Supervised learning	0.904 (95% CI: 0.865, 0.935)	0.736 (95% CI: 0.698, 0.771)	0.794 (95% CI: 0.780, 0.807)	0.744 (95% CI: 0.696, 0.792)	52.1
	MTSN	Low-shot learning	0.869 (95% CI: 0.833, 0.901)	0.758 (95% CI: 0.723, 0.792)	0.841 (95% CI: 0.829, 0.853)	0.743 (95% CI: 0.683, 0.795)	1.8
	MTSN+OVV	Semi-supervised learning	0.898 (95% CI: 0.857, 0.928)	0.775 (95% CI: 0.741, 0.808)	0.881 (95% CI: 0.870, 0.891)	0.763 (95% CI: 0.695, 0.820)	203.7
MobileNet-v2	Baseline	Supervised learning	0.893 (95% CI: 0.845, 0.932)	0.794 (95% CI: 0.760, 0.825)	0.856 (95% CI: 0.844, 0.867)	0.786 (95% CI: 0.728, 0.835)	42.9
	MTSN	Low-shot learning	0.859 (95% CI: 0.813, 0.896)	0.820 (95% CI: 0.786, 0.850)	0.843 (95% CI: 0.831, 0.855)	0.748 (95% CI: 0.689, 0.802)	1.2
	MTSN+OVV	Semi-supervised learning	0.887 (95% CI: 0.850, 0.920)	0.840 (95% CI: 0.808, 0.867)	0.851 (95% CI: 0.838, 0.862)	0.777 (95% CI: 0.718, 0.826)	125.4
DenseNet	Baseline	Supervised learning	0.898 (95% CI: 0.867, 0.927)	0.810 (95% CI: 0.778, 0.841)	0.784 (95% CI: 0.771, 0.798)	0.743 (95% CI: 0.688, 0.789)	122.7
	MTSN	Low-shot learning	0.854 (95% CI: 0.811, 0.894)	0.753 (95% CI: 0.716, 0.786)	0.853 (95% CI: 0.842, 0.865)	0.732 (95% CI: 0.675, 0.785)	5.7
	MTSN+OVV	Semi-supervised learning	0.896 (95% CI: 0.861, 0.926)	0.783 (95% CI: 0.748, 0.817)	0.831 (95% CI: 0.818, 0.843)	0.746 (95% CI: 0.678, 0.800)	324.0
EfficientNet	Baseline	Supervised learning	0.768 (95% CI: 0.684, 0.834)	0.633 (95% CI: 0.590, 0.672)	0.650 (95% CI: 0.634, 0.667)	0.658 (95% CI: 0.611, 0.702)	48.7
	MTSN	Low-shot learning	0.863 (95% CI: 0.818, 0.899)	0.845 (95% CI: 0.815, 0.873)	0.845 (95% CI: 0.833, 0.856)	0.719 (95% CI: 0.659, 0.774)	1.5
	MTSN+OVV	Semi-supervised learning	0.886 (95% CI: 0.845, 0.918)	0.792 (95% CI: 0.758, 0.824)	0.850 (95% CI: 0.837, 0.861)	0.749 (95% CI: 0.690, 0.800)	159.8

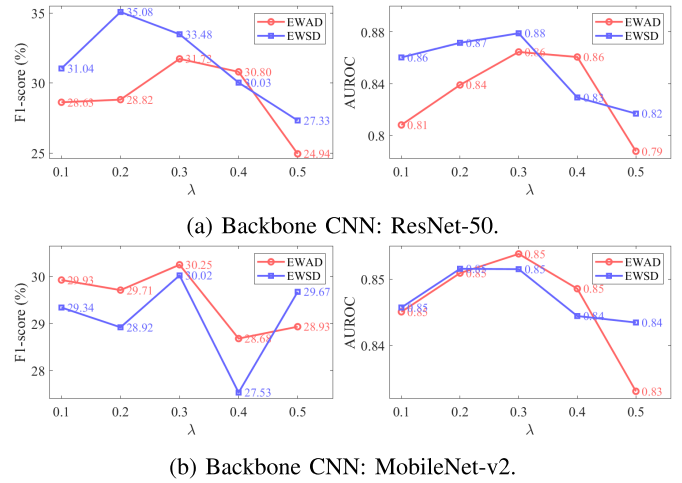


Fig. 5. MTSN performance on the OHTS test set with respect to different λ and Φ (.).

Table I provides the results of a baseline supervised learning experiment conducted on the low-shot training set. The results suggest that low-shot learning performs much better than supervised learning when the training size is small.

Furthermore, we discuss the selection of the thresholds κ_1 and κ_2 (used to select reliable “self-predictions” and “contrastive predictions” in our OVV self-training) as well as the impact of different mini-batch sizes $2m$ on OVV self-training (each mini-batch contains m pairs of reference and target fundus images). Table II shows the MTSN performances with respect to different κ_1 , κ_2 , and m . When evaluated on the OHTS test set, it can be seen that accuracy and F1-score increase slightly, but AUROC almost remains the same, with the decrease of m . Moreover, with the increase of κ_1 and κ_2 , the standard to determine reliable predictions becomes lower, making the semi-supervised learning performance degrade. Based on this experiment, we believe OVV self-training benefits from smaller κ_1 and κ_2 .

Additionally, MTSNs, trained under different m , are evaluated on the three additional test sets, as shown in Table II. It can be seen that the network trained with a larger m typically shows better results. When m decreases, the generalizability of MTSN degrades dramatically, especially for F1-score (decreases by around 9-19%). Therefore, increasing the mini-batch size can improve the MTSN generalizability, as more reference fundus images are used to provide contrastive predictions for the target fundus images, which can veto more unreliable predictions on the unlabeled data. Hence, we increase m to 30 to further improve OVV self-training when comparing it with other published SoTA algorithms, as shown in Sect. IV-D. Since our threshold selection experiments cover a very limited number of discrete sets of κ_1 , κ_2 , and m , we believe better performance can be achieved when more values are tested.

C. Comparison of Supervised, Low-Shot, and Semi-Supervised Glaucoma Diagnosis

Comparisons of supervised learning, low-shot learning, and semi-supervised learning (w.r.t. four backbone CNNs:

TABLE IV

COMPARISONS OF SUPERVISED LEARNING, LOW-SHOT LEARNING, AND SEMI-SUPERVISED LEARNING W.R.T. DIFFERENT PERCENTAGES OF LABELED TRAINING DATA. IN THE BASELINE EXPERIMENTS (2.0% OF TRAINING DATA), ONE FUNDUS PHOTOGRAPH IS SELECTED FROM EACH PATIENT; IN THE EXPERIMENTS WITH 0.5% AND 1.0% OF TRAINING DATA, A SUBSET IS CREATED BY CHOOSING PATIENT IDS AT REGULAR INTERVALS

Percentage of training data	Supervised learning			Low-shot learning			Semi-supervised learning		
	Accuracy (%)	F1-score (%)	AUROC	Accuracy (%)	F1-score (%)	AUROC	Accuracy (%)	F1-score (%)	AUROC
0.5	80.402	18.066	0.720	84.769	20.228	0.759	88.856	25.351	0.797
1.0	79.557	21.573	0.799	88.538	25.654	0.806	88.453	32.472	0.857
2.0 (baseline)	84.141	26.743	0.838	87.150	31.726	0.865	91.415	41.148	0.898
10.0	85.651	32.288	0.890	89.988	38.612	0.891	91.446	39.760	0.899
50.0	92.950	40.188	0.907	94.223	40.826	0.887	92.067	38.187	0.889
90.0	89.556	37.582	0.905	92.897	41.878	0.887	93.331	43.421	0.898

(a) Backbone CNN: ResNet-50.

Percentage of training data	Supervised learning			Low-shot learning			Semi-supervised learning		
	Accuracy (%)	F1-score (%)	AUROC	Accuracy (%)	F1-score (%)	AUROC	Accuracy (%)	F1-score (%)	AUROC
0.5	83.033	19.035	0.720	81.954	18.035	0.745	87.918	29.246	0.826
1.0	67.659	16.458	0.748	91.764	20.628	0.772	89.407	31.357	0.841
2.0 (baseline)	78.144	23.547	0.840	86.018	30.252	0.854	90.609	36.960	0.887
10.0	90.429	34.873	0.884	89.562	36.808	0.897	89.624	36.887	0.888
50.0	92.438	43.361	0.906	91.043	39.050	0.888	93.230	40.653	0.889
90.0	93.754	42.401	0.908	93.750	37.519	0.890	91.857	41.602	0.896

(b) Backbone CNN: MobileNet-v2.

ResNet-50, MobileNet-v2, DenseNet, and EfficientNet) for glaucoma diagnosis are provided in Table III. First, these results suggest that the MTSNs fine-tuned with OVV self-training that requires a small number of labeled fundus images perform similarly (AUROC 95% CI overlaps considerably) and, in some cases, significantly better (AUROC 95% CI does not overlap) than the backbone CNNs trained with a large number of labeled fundus images (50 times larger) under full supervision.

Specifically, when using ResNet-50, MobileNet-v2, or DenseNet as the backbone CNN, semi-supervised learning performs similarly to supervised learning on the OHTS and DIGS/ADAGES test sets, and in most cases, significantly better than supervised learning on the ACRIMA and LAG datasets. Although EfficientNet trained through supervised learning performs unsatisfactorily on all four test sets, it shows considerable compatibility with MTSN in the low-shot and semi-supervised learning experiments. Second as expected, the AUROC scores achieved by low-shot learning are in most, but not all, cases slightly lower than those achieved by the backbone CNNs, when evaluated on the OHTS test set. However, low-shot learning shows better generalizability than supervised learning on the ACRIMA and LAG datasets. Moreover, since low-shot learning uses only a small amount of training data, training an MTSN is much faster than supervised learning. As MTSNs assembled with ResNet-50 and MobileNet-v2 typically demonstrate better performances than the ones assembled with DenseNet and EfficientNet, we only use the former two CNNs for the following experiments.

We also employ Grad-CAM++ [67] to explain the models' decision-making, as shown in Fig. 6. These results suggest that the optic nerve head areas impact model decisions most. The neuroretinal rim areas are identified as

most important, and the periphery contributed comparatively little to model decisions for both healthy and GON eyes [42].

We also carry out a series of experiments with respect to different percentages of training data, as shown in Table IV, to further validate the effectiveness of our proposed low-shot and semi-supervised learning algorithms. The backbone CNNs trained via supervised learning on the small subsets generally perform worse than the MTSNs trained via low-shot and semi-supervised learning. With the labeled training data increase, the models' performance gets saturated. When using less labeled training data (0.5% and 1.0%), the MTSN performance degrades. However, its performance can still be greatly improved with OVV self-training (accuracy, F1-score, and AUROC can be improved by up to 6%, 11%, and 0.08, respectively). In addition, when using over 10% of the entire training data, the MTSN performance saturates, and the OVV self-training can bring very limited improvements on MTSNs.

D. Comparisons With Other SoTA Glaucoma Diagnosis Approaches

Table V provides comprehensive comparisons with nine SoTA glaucoma diagnosis algorithms.⁴ The results suggest that (a) for low-shot learning, the MTSNs trained by minimizing our proposed CWCE loss perform significantly better than the SoTA low-shot glaucoma diagnosis approach [19] on all four datasets (accuracy, F1-score, and AUROC are up to 5%, 15%, and 0.08 higher, respectively), and (b) for semi-supervised learning, the MTSNs fine-tuned with OVV self-training also achieve the superior performances over another two SoTA semi-supervised glaucoma diagnosis approaches [20], [21] (accuracy, F1-score, and AUROC are up to 6%, 11%, and

⁴Our recent work [3] provides the baseline supervised learning results.

TABLE V
COMPARISONS WITH OTHER SOTA GLAUCOMA DIAGNOSIS ALGORITHMS. THE BEST RESULTS FOR EACH TRAINING STRATEGY ARE SHOWN IN BOLD FONT

Training strategy	Method	Accuracy (%)	F1-score (%)	AUROC
Supervised learning	Li <i>et al.</i> [14]	93.812	42.590	0.886
	Gómez-Valverde <i>et al.</i> [15]	95.068	48.039	0.903
	Judy <i>et al.</i> [16]	94.188	42.835	0.908
	Serener and Serte [17]	90.492	41.227	0.912
	Thakur <i>et al.</i> [18]	94.145	44.115	0.896
	Fan <i>et al.</i> [3] (Baseline Result)	93.261	43.016	0.904
Low-shot learning	Kim <i>et al.</i> [19]	84.203	23.161	0.786
	MTSN (Backbone: ResNet-50) (Ours)	87.150	31.726	0.865
	MTSN (Backbone: MobileNet-v2) (Ours)	86.018	30.252	0.854
Semi-supervised learning	Al Ghamdi <i>et al.</i> [20]	84.808	27.579	0.830
	Diaz-Pinto <i>et al.</i> [21]	76.619	20.721	0.748
	MTSN (Backbone: ResNet-50) + OVV Self-Training (Ours)	90.244	38.454	0.899
	MTSN (Backbone: MobileNet-v2) + OVV Self-Training (Ours)	89.360	36.599	0.891

(a) OHTS [12] test set.

Training strategy	Method	Accuracy (%)	F1-score (%)	AUROC
Supervised learning	Li <i>et al.</i> [14]	60.142	46.272	0.813
	Gómez-Valverde <i>et al.</i> [15]	63.546	53.358	0.826
	Judy <i>et al.</i> [16]	57.872	41.650	0.824
	Serener and Serte [17]	54.326	36.364	0.675
	Thakur <i>et al.</i> [18]	65.106	58.020	0.794
	Fan <i>et al.</i> [3] (Baseline Result)	53.333	31.601	0.736
Low-shot learning	Kim <i>et al.</i> [19]	64.965	58.627	0.844
	MTSN (Backbone: ResNet-50) (Ours)	67.092	63.175	0.758
	MTSN (Backbone: MobileNet-v2) (Ours)	70.355	67.797	0.820
Semi-supervised learning	Al Ghamdi <i>et al.</i> [20]	68.511	67.257	0.794
	Diaz-Pinto <i>et al.</i> [21]	59.149	44.828	0.818
	MTSN (Backbone: ResNet-50) + OVV Self-Training (Ours)	64.539	57.627	0.801
	MTSN (Backbone: MobileNet-v2) + OVV Self-Training (Ours)	72.340	70.939	0.835

(b) ACRIMA [22] dataset.

Training strategy	Method	Accuracy (%)	F1-score (%)	AUROC
Supervised learning	Li <i>et al.</i> [14]	76.535	53.756	0.855
	Gómez-Valverde <i>et al.</i> [15]	80.202	62.417	0.883
	Judy <i>et al.</i> [16]	78.348	60.174	0.860
	Serener and Serte [17]	77.379	66.545	0.806
	Thakur <i>et al.</i> [18]	80.305	65.882	0.856
	Fan <i>et al.</i> [3] (Baseline Result)	75.052	50.267	0.794
Low-shot learning	Kim <i>et al.</i> [19]	74.619	65.000	0.805
	MTSN (Backbone: ResNet-50) (Ours)	79.028	65.908	0.841
	MTSN (Backbone: MobileNet-v2) (Ours)	79.007	69.482	0.843
Semi-supervised learning	Al Ghamdi <i>et al.</i> [20]	79.028	72.382	0.860
	Diaz-Pinto <i>et al.</i> [21]	65.554	56.662	0.701
	MTSN (Backbone: ResNet-50) + OVV Self-Training (Ours)	81.644	69.929	0.879
	MTSN (Backbone: MobileNet-v2) + OVV Self-Training (Ours)	80.470	68.692	0.849

(c) LAG [9] dataset.

Training strategy	Method	Accuracy (%)	F1-score (%)	AUROC
Supervised learning	Li <i>et al.</i> [14]	65.293	46.955	0.780
	Gómez-Valverde <i>et al.</i> [15]	65.157	45.729	0.795
	Judy <i>et al.</i> [16]	63.556	41.144	0.760
	Serener and Serte [17]	69.108	57.478	0.757
	Thakur <i>et al.</i> [18]	70.606	60.322	0.786
	Fan <i>et al.</i> [3] (Baseline Result)	62.500	38.042	0.744
Low-shot learning	Kim <i>et al.</i> [19]	63.862	52.949	0.687
	MTSN (Backbone: ResNet-50) (Ours)	67.745	60.754	0.743
	MTSN (Backbone: MobileNet-v2) (Ours)	69.176	68.145	0.748
Semi-supervised learning	Al Ghamdi <i>et al.</i> [20]	64.850	54.816	0.716
	Diaz-Pinto <i>et al.</i> [21]	64.441	59.908	0.677
	MTSN (Backbone: ResNet-50) + OVV Self-Training (Ours)	66.281	55.486	0.747
	MTSN (Backbone: MobileNet-v2) + OVV Self-Training (Ours)	70.402	66.015	0.776

(d) DIGS/ADAGES [23] test set.

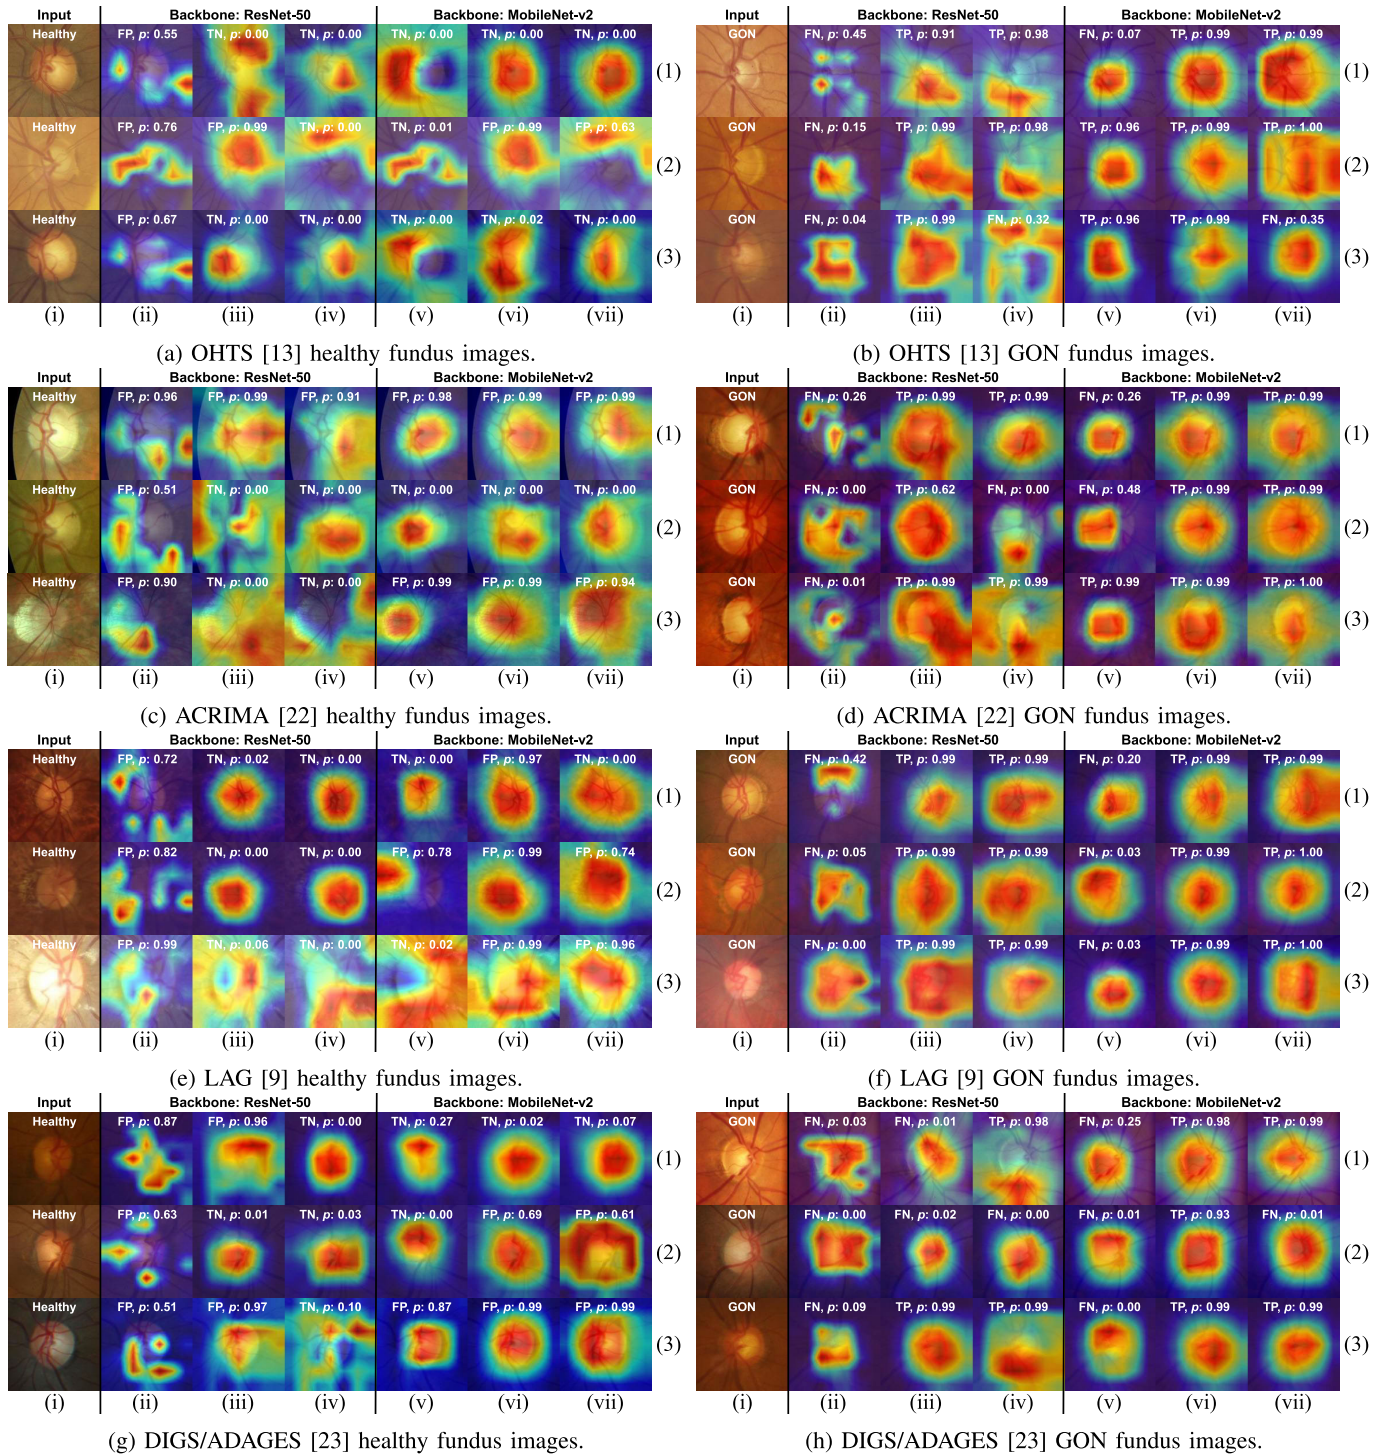


Fig. 6. Examples of Grad-CAM++ [67] results: (i) fundus images; (ii) and (v) show the class activation maps of (i), obtained by the backbone CNNs trained through supervised learning on the entire training set (containing $\sim 53\text{K}$ fundus images); (iii) and (vi) show the class activation maps of (i), obtained by MTSNs trained through low-shot learning on a small training set (containing 1,147 fundus images); (iv) and (vii) show the class activation maps of (i), obtained by MTSNs fine-tuned with our proposed OVV self-training on the entire training set (containing $\sim 53\text{K}$ fundus images without ground-truth labels).

0.07 higher, respectively). Compared with the SoTA supervised approaches, the fine-tuned MTSNs demonstrate similar performance on the OHTS test set and better generalizability on three additional test sets. Therefore, we believe that MTSN with our proposed OVV self-training is an effective technique for semi-supervised glaucoma diagnosis.

E. Comparisons With SoTA General-Purpose Semi-Supervised Learning Approaches

Table VI provides a comprehensive comparison of our proposed OVV self-training approach with four SoTA general-purpose semi-supervised learning methods: FreeMatch [24], SoftMatch [25], FixMatch [26], and FlexMatch [27], which

TABLE VI

COMPARISONS WITH SoTA GENERAL-PURPOSE SEMI-SUPERVISED LEARNING METHODS THAT USE THE VISION TRANSFORMER [28] AS THE BACKBONE NETWORK. THE BEST RESULTS OF EACH DATASET ARE SHOWN IN BOLD FONT

Test Set	Method	Accuracy (%)	F1-score (%)	AUROC
OHTS [11], [12]	FreeMatch [24]	86.510	27.745	0.784
	SoftMatch [25]	85.679	25.718	0.768
	FixMatch [26]	84.695	27.825	0.807
	FlexMatch [27]	92.715	26.331	0.696
	OVV Self-Training (Ours)	90.244	38.454	0.899
ACRIMA [22]	FreeMatch [24]	77.872	77.778	0.791
	SoftMatch [25]	78.582	79.115	0.795
	FixMatch [26]	78.582	79.622	0.792
	FlexMatch [27]	60.426	47.850	0.644
	OVV Self-Training (Ours)	64.539	57.627	0.801
LAG [9]	FreeMatch [24]	78.471	67.087	0.793
	SoftMatch [25]	79.110	67.972	0.798
	FixMatch [26]	81.314	73.394	0.827
	FlexMatch [27]	71.302	35.360	0.654
	OVV Self-Training (Ours)	81.644	69.929	0.879
DIGS/ADAGES [23]	FreeMatch [24]	65.429	56.716	0.685
	SoftMatch [25]	65.395	56.986	0.679
	FixMatch [26]	68.835	65.090	0.731
	FlexMatch [27]	57.050	26.558	0.594
	OVV Self-Training (Ours)	66.281	55.486	0.747

all employ vision Transformer [28] as their backbone network. Our results demonstrate that the proposed OVV self-training approach outperforms these methods in terms of F1-score and AUROC on the OHTS dataset. Specifically, we observe improvements in F1-score ranging from approximately 11% to 13%, and improvements in AUROC ranging from around 9% to 20%. Furthermore, our method demonstrates better generalizability in terms of AUROC across three additional fundus image test sets. Although their results are inferior to ours, particularly in terms of AUROC, it may be unjust to compare them as they were not specifically designed for the diagnosis of glaucoma or other diseases.

F. MTSN and CWCE Loss for Few-Shot Multi-Class Biomedical Image Classification

We conduct two additional few-shot multi-class lung disease diagnosis experiments: (a) chest X-ray image classification for COVID-19 and viral pneumonia detection [29], [30], and (b) lung histopathological image classification for lung cancer diagnosis [31], to validate the effectiveness of our proposed MTSN and CWCE loss. The first experiment has three classes of images: (1) healthy, (2) viral pneumonia, and (3) COVID-19 (an example of each class is shown in Fig. 7), while the second experiment also has three classes of images: (1) benign tissue, (2) adenocarcinoma, and (3) squamous cell carcinoma (an example of each class is shown in Fig. 7). In these two experiments, we only select a few images from each class for MTSN training. The numbers of images used for training, validation, and testing are given in Table VII, where it can be observed that the training set is much smaller than the validation and test sets. Since (1) can only be used for binary image classification problems, we extend it here to tackle multi-class image classification problems.

Let us consider the chest X-ray image classification task as an example. Each image \mathbf{x} is assigned a pair of two labels (r, s) with the following values:

- $r = 1$ and $s = 0$ when \mathbf{x} is a healthy image;
- $r = 2$ and $s = 1$ when \mathbf{x} is a viral pneumonia image;
- $r = 3$ and $s = 3$ when \mathbf{x} is a COVID-19 image.

The numbers of healthy images (class 1), viral pneumonia images (class 2), and COVID-19 images (class 3) are denoted as $n_1, n_2,$ and $n_3,$ respectively. The total number of images is $N = n_1 + n_2 + n_3.$ The weight $\omega_{\text{cla},e}$ used in the image classification loss \mathcal{L}_{cla} w.r.t. class e is given by:

$$\omega_{\text{cla},e} = \frac{N - n_e}{2N}. \tag{6}$$

$\sum_{e=1}^3 \omega_{\text{cla},e} = 1.$ Therefore, \mathcal{L}_{cla} can be written as follows:

$$\mathcal{L}_{\text{cla}} = - \sum_{e=1}^3 \omega_{\text{cla},e} \left(k_{e,i} \log(p_e(\mathbf{x}_i)) + k_{e,j} \log(p_e(\mathbf{x}_j)) \right), \tag{7}$$

where $p_e(\mathbf{x}) \in [0, 1]$ indicates the probability that \mathbf{x} belongs to class $e.$ $\sum_{e=1}^3 p_e(\mathbf{x}) = 1.$ $k_{e,i} = 1$ when $e = r_i,$ and $k_{e,i} = 0$ otherwise. Given a pair of images \mathbf{x}_i and \mathbf{x}_j with ground-truth labels (r_i, s_i) and $(r_j, s_j),$ respectively, there are four cases:

- case 1: $|s_i - s_j| = 1$ (healthy v.s. viral pneumonia);
- case 2: $|s_i - s_j| = 2$ (viral pneumonia v.s. COVID-19);
- case 3: $|s_i - s_j| = 3$ (COVID-19 v.s. healthy);
- case 4: $|s_i - s_j| = 0$ (two images are of the same class).

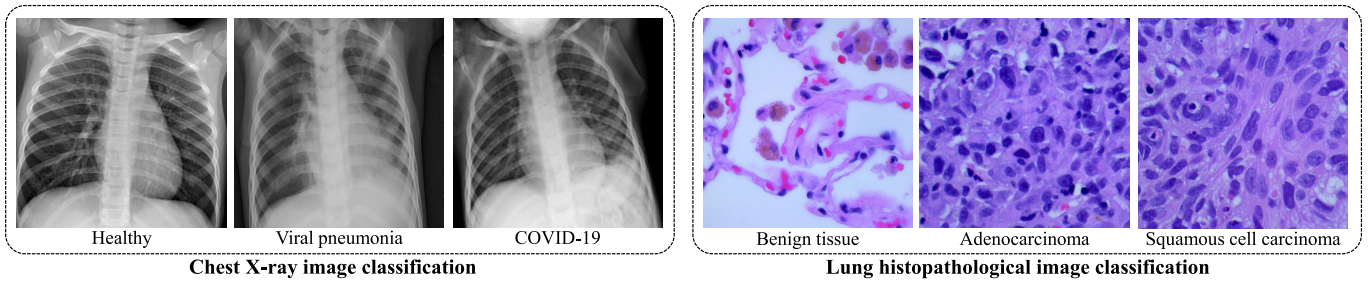


Fig. 7. Examples of images used in two few-shot multi-class lung disease diagnosis tasks.

		Ground truth			PRE
		Healthy	Viral pneumonia	COVID-19	ACC
Prediction	Healthy	622 32.7%	44 2.3%	9 0.5%	92.1% 7.9%
	Viral pneumonia	31 1.6%	605 31.8%	27 1.4%	91.3% 8.7%
	COVID-19	4 0.2%	10 0.5%	552 29.0%	97.5% 2.5%
REC		94.7% 5.3%	91.8% 8.2%	93.9% 6.1%	93.4% 6.6%

Chest X-ray image classification

		Ground truth			PRE
		Benign tissue	Adenocarcinoma	Squamous cell carcinoma	ACC
Prediction	Benign tissue	2462 32.9%	142 1.9%	8 0.1%	94.3% 5.7%
	Adenocarcinoma	33 0.4%	2149 28.7%	388 5.2%	83.6% 16.4%
	Squamous cell carcinoma	0 0.0%	204 2.7%	2099 28.0%	91.1% 8.9%
REC		98.7% 1.3%	86.1% 13.9%	84.1% 15.9%	89.6% 10.4%

Lung histopathological image classification

Fig. 8. Experimental results of two few-shot lung disease diagnosis tasks. REC: recall; PRE: precision; ACC: accuracy.

The weight $\omega_{\text{sim},c}$ used in the image similarity comparison loss \mathcal{L}_{sim} with respect to case c is given by:

$$\omega_{\text{sim},c} = \begin{cases} \frac{1}{3} \left(1 - \frac{2n_1n_2}{N(N-1)}\right) & \text{if } c = 1 \\ \frac{1}{3} \left(1 - \frac{2n_2n_3}{N(N-1)}\right) & \text{if } c = 2 \\ \frac{1}{3} \left(1 - \frac{2n_1n_3}{N(N-1)}\right) & \text{if } c = 3 \\ \frac{1}{3} \left(1 - \frac{\sum_{w=1}^3 n_w(n_w-1)}{N(N-1)}\right) & \text{if } c = 4, \end{cases} \quad (8)$$

$\sum_{c=1}^4 \omega_{\text{sim},c} = 1$. Therefore, \mathcal{L}_{sim} can be written as follows:

$$\mathcal{L}_{\text{sim}} = - \sum_{c=1}^4 \omega_{\text{sim},c} h_{c,(i,j)} \log(q_c(\mathbf{x}_i, \mathbf{x}_j)), \quad (9)$$

where $q_c(\mathbf{x}_i, \mathbf{x}_j) \in [0, 1]$ indicates the similarity between \mathbf{x}_i and \mathbf{x}_j under case c . $\sum_{c=1}^4 q_c(\mathbf{x}_i, \mathbf{x}_j) = 1$. $h_{c,(i,j)} = 1$ when

TABLE VII

TRAINING, VALIDATION, AND TEST SAMPLE SIZES IN THE TWO FEW-SHOT BIOMEDICAL IMAGE CLASSIFICATION EXPERIMENTS

Class	Training	Validation	Test
Healthy	27	657	657
Viral pneumonia	27	659	659
COVID-19	24	588	588

(a) Chest X-ray image classification.

Class	Training	Validation	Test
Benign tissue	10	2,495	2,495
Adenocarcinoma	10	2,495	2,495
Squamous cell carcinoma	10	2,495	2,495

(b) Lung histopathological image classification.

\mathbf{x}_i and \mathbf{x}_j belong to case c , and $h_{c,(i,j)} = 0$ otherwise. The hyper-parameter λ is empirically set to 0.3.

The experimental results of these two multi-task biomedical image classification tasks are presented in Fig. 8 with two confusion matrices. These results demonstrate that our proposed MTSN can be effectively trained with very few images to solve multi-class biomedical image classification problems. Specifically, the achieved accuracy values for chest X-ray image classification (~ 25 -shot learning) and lung histopathological image classification (10-shot learning) are 93% and 90%, respectively. The chest X-ray image classification result compares favorably with the accuracy range of 82%-93% achieved by supervised methods (2,520 images for training, 840 images for validation, and 840 images for testing) using all available training data [68]. Although the accuracy achieved by MTSN for lung histopathological image classification is lower than the accuracy of over 97% reported in [69] by supervised approaches using the full training set (8,250 images for training, $\sim 3,000$ images for validation, and 3,744 images for testing), we believe that our proposed low-shot learning method can achieve comparable results when a small amount of additional images are incorporated for MTSN training.

V. DISCUSSION

Extensive experiments demonstrate the effectiveness and efficiency of training an MTSN by minimizing our proposed CWCE loss. Such a low-shot learning approach significantly reduces over-fitting and achieves an accuracy on a small training set (1,147 fundus images) comparable to a large training set (approximately 53K fundus images). We also demonstrate its effectiveness on two additional multi-class

few-shot biomedical image classification tasks. Additionally, the MTSNs fine-tuned with OVV self-training outperform the SoTA semi-supervised glaucoma diagnosis algorithms [20], [21] as well as general-purpose semi-supervised learning algorithms [24], [25], [26], [27] trained for glaucoma diagnosis. They perform similarly, and in some cases, better than SoTA supervised algorithms. However, our proposed method has two limitations:

- In the OVV self-training, each target fundus image must be compared with all the reference fundus images in the same mini-batch, resulting in a computational complexity of $\mathcal{O}(n^2)$. As the mini-batch size increases, OVV self-training becomes relatively memory-consuming. The high computational complexity of OVV self-training may reduce the feasibility of this method in clinical practice for now. Therefore, we plan to improve the OVV self-training strategy in the future by adaptively selecting only a limited number of reference fundus photographs for semi-supervised glaucoma diagnosis, which can reduce the computational complexity and make the method more practical in clinical settings.
- Our proposed OVV self-training strategy is developed for binary image classification and may not be directly applicable to multi-class image classification problems. Therefore, we plan to extend the contrastive prediction procedure to handle multi-class image classification problems in future work. More hyper-parameter tuning can always be done, but it is so easy to over-fit with limited data.

VI. CONCLUSION

The main contributions of this paper include: (1) a multi-task Siamese network that can learn glaucoma diagnosis from very limited labeled training data; (2) an effective semi-supervised learning strategy, referred to as One-Vote Veto self-training, which can produce pseudo labels for the unlabeled data to fine-tune a pre-trained multi-task Siamese network. Extensive experiments conducted on four fundus image datasets demonstrated the effectiveness of these proposed techniques. The low-shot learning reduces over-fitting and achieves an accuracy on a small training set comparable to that of a large training set. Furthermore, with One-Vote Veto self-training, the multi-task Siamese networks perform similarly to their backbone CNNs (trained via supervised learning on the full training set) on the OHTS test set and show better generalizability on three additional test sets. The methods introduced in this paper can also be applied to other few-shot multi-class biomedical image classification problems, e.g., COVID-19 and lung cancer diagnosis, and other diseases in which only a small quantity of ground-truth labels are available for network training.

REFERENCES

- [1] R. N. Weinreb and P. T. Khaw, "Primary open-angle glaucoma," *Lancet*, vol. 363, no. 9422, pp. 1711–1720, May 2004.
- [2] Y.-C. Tham, X. Li, T. Y. Wong, H. A. Quigley, T. Aung, and C.-Y. Cheng, "Global prevalence of glaucoma and projections of glaucoma burden through 2040: A systematic review and meta-analysis," *Ophthalmology*, vol. 121, no. 11, pp. 2081–2090, 2014.
- [3] R. Fan et al., "Detecting glaucoma in the ocular hypertension study using deep learning," *J. Amer. Med. Assoc. Ophthalmol.*, vol. 140, no. 4, pp. 383–391, Apr. 2022.
- [4] C. E. Traverso, "Direct costs of glaucoma and severity of the disease: A multinational long term study of resource utilisation in Europe," *Brit. J. Ophthalmol.*, vol. 89, no. 10, pp. 1245–1249, Oct. 2005.
- [5] W. Huang, K. Gao, Y. Liu, M. Liang, and X. Zhang, "The adverse impact of glaucoma on psychological function and daily physical activity," *J. Ophthalmol.*, vol. 2020, Apr. 2020, Art. no. 9606420.
- [6] R. K. Parrish et al., "Visual function and quality of life among patients with glaucoma," *Arch. Ophthalmol.*, vol. 115, no. 11, pp. 1447–1455, Nov. 1997.
- [7] M. Kwon, C. Huisinigh, L. A. Rhodes, G. McGwin, J. M. Wood, and C. Owsley, "Association between glaucoma and at-fault motor vehicle collision involvement among older drivers," *Ophthalmology*, vol. 123, no. 1, pp. 109–116, Jan. 2016.
- [8] G. McGwin, C. Huisinigh, S. G. Jain, C. A. Girkin, and C. Owsley, "Binocular visual field impairment in glaucoma and at-fault motor vehicle collisions," *J. Glaucoma*, vol. 24, no. 2, pp. 138–143, Feb. 2015.
- [9] L. Li, M. Xu, X. Wang, L. Jiang, and H. Liu, "Attention based glaucoma detection: A large-scale database and CNN model," in *Proc. IEEE/CVF Conf. Comput. Vis. Pattern Recognit. (CVPR)*, Jun. 2019, pp. 10563–10572.
- [10] R. Fan et al., "Detecting glaucoma from fundus photographs using deep learning without convolutions: Transformer for improved generalization," *Ophthalmol. Sci.*, vol. 3, no. 1, Oct. 2022, Art. no. 100233.
- [11] M. O. Gordon and M. A. Kass, "The ocular hypertension treatment study: Design and baseline description of the participants," *Arch. Ophthalmol.*, vol. 117, no. 5, pp. 573–583, May 1999.
- [12] M. A. Kass et al., "The ocular Hypertension Treatment study: A randomized trial determines that topical ocular hypotensive medication delays or prevents the onset of primary open-angle glaucoma," *Arch. Ophthalmol.*, vol. 120, no. 6, pp. 701–713, Jun. 2002.
- [13] M. O. Gordon et al., "Assessment of the impact of an endpoint committee in the ocular hypertension treatment study," *Amer. J. Ophthalmol.*, vol. 199, pp. 193–199, Mar. 2019.
- [14] Z. Li, Y. He, S. Keel, W. Meng, R. T. Chang, and M. He, "Efficacy of a deep learning system for detecting glaucomatous optic neuropathy based on color fundus photographs," *Ophthalmology*, vol. 125, no. 8, pp. 1199–1206, Aug. 2018.
- [15] J. J. Gómez-Valverde et al., "Automatic glaucoma classification using color fundus images based on convolutional neural networks and transfer learning," *Biomed. Opt. Exp.*, vol. 10, no. 2, pp. 892–913, 2019.
- [16] D. Judy, "Automated identification of glaucoma from fundus images using deep learning techniques," *Eur. J. Mol. Clin. Med.*, vol. 7, no. 2, pp. 5449–5458, Nov. 2020.
- [17] A. Serener and S. Serte, "Transfer learning for early and advanced glaucoma detection with convolutional neural networks," in *Proc. Med. Technol. Congr. (TIPEKNO)*, Oct. 2019, pp. 1–4.
- [18] A. Thakur, M. Goldbaum, and S. Yousefi, "Predicting glaucoma before onset using deep learning," *Ophthalmol. Glaucoma*, vol. 3, no. 4, pp. 262–268, Jul. 2020.
- [19] M. Kim, J. Zuallaert, and W. De Neve, "Few-shot learning using a small-sized dataset of high-resolution FUNDUS images for glaucoma diagnosis," in *Proc. 2nd Int. Workshop Multimedia Pers. Health Health Care*, Oct. 2017, pp. 89–92.
- [20] M. Al Ghamdi, M. Li, M. Abdel-Mottaleb, and M. A. Shousha, "Semi-supervised transfer learning for convolutional neural networks for glaucoma detection," in *Proc. IEEE Int. Conf. Acoust., Speech Signal Process. (ICASSP)*, May 2019, pp. 3812–3816.
- [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 Trans. Med. Imag.*, vol. 38, no. 9, pp. 2211–2218, Sep. 2019.
- [22] A. Diaz-Pinto, S. Morales, V. Naranjo, T. Köhler, J. M. Mossi, and A. Navea, "CNNs for automatic glaucoma assessment using fundus images: An extensive validation," *Biomed. Eng. OnLine*, vol. 18, no. 1, pp. 1–19, Dec. 2019.
- [23] P. Sample et al., "The African descent and glaucoma evaluation study (ADAGES): Design and baseline data," *Arch. Ophthalmol.*, vol. 127, no. 9, pp. 1136–1145, Sep. 2009.
- [24] Y. Wang et al., "FreeMatch: Self-adaptive thresholding for semi-supervised learning," in *Proc. Int. Conf. Learn. Represent. (ICLR)*, 2023, pp. 1–20.
- [25] H. Chen et al., "SoftMatch: Addressing the quantity-quality trade-off in semi-supervised learning," in *Proc. Int. Conf. Learn. Represent. (ICLR)*, 2023, pp. 1–21.

- [26] K. Sohn et al., "FixMatch: Simplifying semi-supervised learning with consistency and confidence," in *Proc. Adv. Neural Inf. Process. Syst.*, vol. 33, 2020, pp. 596–608.
- [27] B. Zhang et al., "FlexMatch: Boosting semi-supervised learning with curriculum pseudo labeling," in *Proc. Adv. Neural Inf. Process. Syst.*, vol. 34, 2021, pp. 18408–18419.
- [28] A. Dosovitskiy et al., "An image is worth 16×16 words: Transformers for image recognition at scale," in *Proc. Int. Conf. Learn. Represent. (ICLR)*, 2021, pp. 1–22.
- [29] M. E. H. Chowdhury et al., "Can AI help in screening viral and COVID-19 pneumonia?" *IEEE Access*, vol. 8, pp. 132665–132676, 2020.
- [30] T. Rahman et al., "Exploring the effect of image enhancement techniques on COVID-19 detection using chest X-ray images," *Comput. Biol. Med.*, vol. 132, May 2021, Art. no. 104319.
- [31] A. A. Borkowski, M. M. Bui, L. B. Thomas, C. P. Wilson, L. A. DeLand, and S. M. Mastorides, "Lung and colon cancer histopathological image dataset (LC25000)," 2019, *arXiv:1912.12142*.
- [32] A. Krizhevsky, "One weird trick for parallelizing convolutional neural networks," 2014, *arXiv:1404.5997*.
- [33] K. Simonyan and A. Zisserman, "Very deep convolutional networks for large-scale image recognition," in *Proc. Int. Conf. Learn. Represent. (ICLR)*, 2015, pp. 1–14.
- [34] C. Szegedy et al., "Going deeper with convolutions," in *Proc. IEEE Conf. Comput. Vis. Pattern Recognit. (CVPR)*, Jun. 2015, pp. 1–9.
- [35] C. Szegedy, V. Vanhoucke, S. Ioffe, J. Shlens, and Z. Wojna, "Rethinking the inception architecture for computer vision," in *Proc. IEEE Conf. Comput. Vis. Pattern Recognit. (CVPR)*, Jun. 2016, pp. 2818–2826.
- [36] J. M. Ahn, S. Kim, K.-S. Ahn, S.-H. Cho, K. B. Lee, and U. S. Kim, "A deep learning model for the detection of both advanced and early glaucoma using fundus photography," *PLoS One*, vol. 13, no. 11, Nov. 2018, Art. no. e0207982.
- [37] J. Deng, W. Dong, R. Socher, L.-J. Li, K. Li, and L. Fei-Fei, "ImageNet: A large-scale hierarchical image database," in *Proc. IEEE Conf. Comput. Vis. Pattern Recognit.*, Jun. 2009, pp. 248–255.
- [38] K. He, X. Zhang, S. Ren, and J. Sun, "Deep residual learning for image recognition," in *Proc. IEEE Conf. Comput. Vis. Pattern Recognit. (CVPR)*, Jun. 2016, pp. 770–778.
- [39] S. Liu et al., "A deep learning-based algorithm identifies glaucomatous discs using monoscopic fundus photographs," *Ophthalmol. Glaucoma*, vol. 1, no. 1, pp. 15–22, Jul. 2018.
- [40] A. R. Ran et al., "Detection of glaucomatous optic neuropathy with spectral-domain optical coherence tomography: A retrospective training and validation deep-learning analysis," *Lancet Digit. Health*, vol. 1, no. 4, pp. 172–182, Aug. 2019.
- [41] F. A. Medeiros, A. A. Jammal, and E. B. Mariottoni, "Detection of progressive glaucomatous optic nerve damage on fundus photographs with deep learning," *Ophthalmology*, vol. 128, no. 3, pp. 383–392, Mar. 2021.
- [42] M. Christopher et al., "Performance of deep learning architectures and transfer learning for detecting glaucomatous optic neuropathy in fundus photographs," *Sci. Rep.*, vol. 8, no. 1, pp. 1–13, Nov. 2018.
- [43] M. Christopher et al., "Effects of study population, labeling and training on glaucoma detection using deep learning algorithms," *Transl. Vis. Sci. Technol.*, vol. 9, no. 2, p. 27, Apr. 2020.
- [44] D. Jain, T. Swedish, B. Shen, D. Y. Kim, S. Mukai, and R. Raskar, "Open-source, ultra-low-cost smartphone attachment for non-mydratric fundus photography—Open indirect ophthalmoscope," *Investigative Ophthalmol. Vis. Sci.*, vol. 57, no. 12, p. 1685, Sep. 2016.
- [45] E. M. Lawson and R. Raskar, "Smart phone administered fundus imaging without additional imaging optics," *Investigative Ophthalmol. Vis. Sci.*, vol. 55, no. 13, p. 1609, Apr. 2014.
- [46] M. Sandler, A. Howard, M. Zhu, A. Zhmoginov, and L.-C. Chen, "MobileNetV2: Inverted residuals and linear bottlenecks," in *Proc. IEEE/CVF Conf. Comput. Vis. Pattern Recognit.*, Jun. 2018, pp. 4510–4520.
- [47] Y. Wang, Q. Yao, J. T. Kwok, and L. M. Ni, "Generalizing from a few examples: A survey on few-shot learning," *ACM Comput. Surveys*, vol. 53, no. 3, pp. 1–34, May 2021.
- [48] P. Zhou et al., "Attention-based bidirectional long short-term memory networks for relation classification," in *Proc. 54th Annu. Meeting Assoc. Comput. Linguistics*, vol. 2, 2016, pp. 207–212.
- [49] A. Radford, L. Metz, and S. Chintala, "Unsupervised representation learning with deep convolutional generative adversarial networks," in *Proc. Int. Conf. Learn. Represent. (ICLR)*, 2015, pp. 1–16.
- [50] I. Triguero, S. García, and F. Herrera, "Self-labeled techniques for semi-supervised learning: Taxonomy, software and empirical study," *Knowl. Inf. Syst.*, vol. 42, no. 2, pp. 245–284, 2015.
- [51] J. E. Van Engelen and H. H. Hoos, "A survey on semi-supervised learning," *Mach. Learn.*, vol. 109, no. 2, pp. 373–440, Feb. 2020.
- [52] A. Miller, A. Fisch, J. Dodge, A.-H. Karimi, A. Bordes, and J. Weston, "Key-value memory networks for directly reading documents," in *Proc. Conf. Empirical Methods Natural Lang. Process.*, 2016, pp. 1400–1409.
- [53] G. Koch, "Siamese neural networks for one-shot image recognition," in *Proc. ICML Deep Learn. Workshop*, vol. 2, Lille, France, 2015, pp. 1–30.
- [54] J. Jang and C. O. Kim, "Siamese network-based health representation learning and robust reference-based remaining useful life prediction," *IEEE Trans. Ind. Informat.*, vol. 18, no. 8, pp. 5264–5274, Aug. 2022.
- [55] Q. Wang, Z. Teng, J. Xing, J. Gao, W. Hu, and S. Maybank, "Learning attentions: Residual attentional Siamese network for high performance online visual tracking," in *Proc. IEEE/CVF Conf. Comput. Vis. Pattern Recognit.*, Jun. 2018, pp. 4854–4863.
- [56] W. Wang, J. Yang, J. Xiao, S. Li, and D. Zhou, "Face recognition based on deep learning," in *Proc. 1st Int. Conf. Human Centered Comput. (HCC)*. Cham, Switzerland: Springer, Jan. 2015, pp. 812–820.
- [57] E. Hoffer and N. Ailon, "Deep metric learning using triplet network," in *Proc. Int. Workshop Similarity-Based Pattern Recognit.* Cham, Switzerland: Springer, Nov. 2015, pp. 84–92.
- [58] P. Khosla et al., "Supervised contrastive learning," in *Proc. Adv. Neural Inf. Process. Syst.*, vol. 33, 2020, pp. 18661–18673.
- [59] R. Cipolla, Y. Gal, and A. Kendall, "Multi-task learning using uncertainty to weigh losses for scene geometry and semantics," in *Proc. IEEE/CVF Conf. Comput. Vis. Pattern Recognit.*, Jun. 2018, pp. 7482–7491.
- [60] D. McClosky, E. Charniak, and M. Johnson, "Effective self-training for parsing," in *Proc. Main Conf. Human Lang. Technol. Conf. North Amer. Chapter Assoc. Comput. Linguistics*, 2006, pp. 152–159.
- [61] S. Sukhbaatar, A. Szlam, J. Weston, and R. Fergus, "End-to-end memory networks," in *Proc. Adv. Neural Inf. Process. Syst. (NeurIPS)*, 2015, pp. 2440–2448.
- [62] L.-C. Chen, G. Papandreou, F. Schroff, and H. Adam, "Rethinking atrous convolution for semantic image segmentation," 2017, *arXiv:1706.05587*.
- [63] L. van der Maaten and G. Hinton, "Visualizing data using t-SNE," *J. Mach. Learn. Res.*, vol. 9, pp. 2579–2605, Nov. 2008.
- [64] C. Tan, F. Sun, T. Kong, W. Zhang, C. Yang, and C. Liu, "A survey on deep transfer learning," in *Proc. Int. Conf. Artif. Neural Netw. (ICANN)*. Cham, Switzerland: Springer, Sep. 2018, pp. 270–279.
- [65] G. Huang, Z. Liu, L. Van Der Maaten, and K. Q. Weinberger, "Densely connected convolutional networks," in *Proc. IEEE Conf. Comput. Vis. Pattern Recognit. (CVPR)*, Jul. 2017, pp. 2261–2269.
- [66] M. Tan and Q. Le, "EfficientNet: Rethinking model scaling for convolutional neural networks," in *Proc. 36th Int. Conf. Mach. Learn. (ICML)*, 2019, pp. 6105–6114.
- [67] A. Chattopadhyay, A. Sarkar, P. Howlader, and V. N. Balasubramanian, "Grad-CAM++: Generalized gradient-based visual explanations for deep convolutional networks," in *Proc. IEEE Winter Conf. Appl. Comput. Vis. (WACV)*, Mar. 2018, pp. 839–847.
- [68] S. Jadon, "COVID-19 detection from scarce chest X-ray image data using few-shot deep learning approach," *Proc. SPIE*, vol. 11601, Feb. 2021, Art. no. 116010X.
- [69] M. A. Abbas, S. U. K. Bukhari, A. Syed, and S. S. H. Shah, "The histopathological diagnosis of adenocarcinoma & squamous cells carcinoma of lungs by artificial intelligence: A comparative study of convolutional neural networks," *MedRxiv*, pp. 1–13, May 2020, doi: [10.1101/2020.05.02.20044602](https://doi.org/10.1101/2020.05.02.20044602).