

Harvard Glaucoma Detection and Progression: A Multimodal Multitask Dataset and Generalization-Reinforced Semi-Supervised Learning

Yan Luo* Min Shi* Yu Tian* Tobias Elze Mengyu Wang
Harvard Ophthalmology AI Lab, Harvard University

{yluo16, mshi6, ytian11, tobias.elze, mengyu.wang}@meei.harvard.edu

Abstract

Glaucoma is the number one cause of irreversible blindness globally. A major challenge for accurate glaucoma detection and progression forecasting is the bottleneck of limited labeled patients with the state-of-the-art (SOTA) 3D retinal imaging data of optical coherence tomography (OCT). To address the data scarcity issue, this paper proposes two solutions. First, we develop a novel generalization-reinforced semi-supervised learning (SSL) model called pseudo supervisor to optimally utilize unlabeled data. Compared with SOTA models, the proposed pseudo supervisor optimizes the policy of predicting pseudo labels with unlabeled samples to improve empirical generalization. Our pseudo supervisor model is evaluated with two clinical tasks consisting of glaucoma detection and progression forecasting. The progression forecasting task is evaluated both unimodally and multimodally. Our pseudo supervisor model demonstrates superior performance than SOTA SSL comparison models. Moreover, our model also achieves the best results on the publicly available LAG fundus dataset. Second, we introduce the Harvard Glaucoma Detection and Progression (Harvard-GDP) Dataset, a multimodal multitask dataset that includes data from 1,000 patients with OCT imaging data, as well as labels for glaucoma detection and progression. This is the largest glaucoma detection dataset with 3D OCT imaging data and the first glaucoma progression forecasting dataset that is publicly available. Detailed sex and racial analysis are provided, which can be used by interested researchers for fairness learning studies. Our released dataset is benchmarked with several SOTA supervised CNN and transformer deep learning models. The dataset and code are made publicly available via <https://ophai.hms.harvard.edu/datasets/harvard-gdp1000>.

1. Introduction

Glaucoma is the leading cause of irreversible blindness globally caused by retinal nerve fiber layer (RNFL)

*First three authors contribute equally.

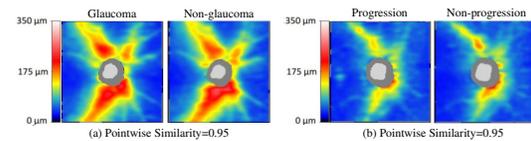


Figure 1: Examples of RNFLT maps with high similarities but from different label groups.

damage [74, 58, 59, 69]. The global prevalence of glaucoma for the population between 40 and 80 years old is 3.54% [74]. Timely clinical treatment in the early stage of glaucoma can significantly reduce the risk of further vision loss [18, 17]. However, most commonly, glaucoma patients are not aware of their early VF loss until the VF loss becomes severe enough to impair their daily activities such as reading and driving due to the brain and fellow eye compensation [4, 19, 34, 50, 51, 30].

It is desirable to develop automated glaucoma screen tools to reduce societal disease burden [77]. The clinical diagnosis of glaucoma is a holistic decision primarily based on RNFL damage and VF loss assessment supplemented with other patient clinical information and family disease history [29, 47]. The VF test is only available in the ophthalmology specialty. So existing automated screening tools are based on the idea of predicting glaucoma diagnosis from evaluating RNFL damage using retinal images [14, 57, 42, 5, 49, 37, 15, 35, 60, 81]. Before the massive clinical adoption of the 3D retinal imaging tool of optical coherence tomography (OCT) a decade ago [31, 12], 2D fundus photos were used to assess RNFL damage. Therefore, most prior glaucoma detection models are based on the 2D fundus photo [37, 57, 33, 14, 42, 1, 64, 15], which is not the state of the art (SOTA) clinical imaging modality anymore. OCT has replaced fundus photos as the primary imaging tool for glaucoma clinical care due to its superiority in measuring RNFL damage [48, 12]. There have been a number of recent works using OCT scans to automatically screen glaucoma with machine learning [15, 60, 5, 35, 24, 80, 49]. Apart from glaucoma detection, a more clinically significant task is progression forecasting using OCT scans [52, 16, 32, 28], which could inform clinical

icians if they should treat the patient conservatively with eye drops or aggressively with invasive surgeries.

For both the glaucoma screening and progression forecasting tasks, a significant challenge is limited labeled patients with the SOTA OCT measurements as the fundus photo is already outdated in clinical practice [15, 60, 5, 35, 24, 80, 49]. This problem with the lack of labeled patients is even more severe for progression forecasting, which requires longitudinal VF measurements to determine progression outcomes [52, 16, 32]. Therefore, there is a strong need for developing novel semi-supervised learning models to maximally utilize unlabeled patient data in clinics to improve glaucoma detection and progression forecasting. However, apart from several papers using standard semi-supervised learning methods for glaucoma detection [20, 2, 23, 79, 84, 9], there has been no dedicated algorithm development for semi-supervised learning with its application in glaucoma. SOTA semi-supervised learning (SSL) approaches in computer vision are generally based on consistency learning [11, 68, 10, 73] or pseudo-labeling [40, 61]. Unlike consistency learning, the pseudo-labeling approaches do not need domain knowledge to design the image augmentation strategy. While many prior works have endeavored to optimally select pseudo-labeled samples to improve the classifier training, the tasks of glaucoma detection and progression forecasting still have some unique challenges that might not be well tackled by existing pseudo-labeling SSL methods [61, 63, 13, 40, 56]. Figure 1 shows examples of RNFLT maps with high similarities but from different label groups. With such obscured group differences, existing pseudo-labeling SSL methods might not be able to provide sufficient guidance to select correct pseudo-labeled samples to improve the classifier training. Furthermore, having high-quality public datasets with OCT scans for glaucoma detection and progression forecasting is equally important to enable more computer vision researchers to study this topic to push forward the field,

In this study, we make two contributions: (1) developing a generalization-reinforced semi-supervised learning model called pseudo supervisor to improve glaucoma detection and progression forecasting. Specifically, our pseudo supervisor model will select pseudo-labeled samples via a policy net that optimizes generalization error with a validation subset to update the classifier. (2) releasing a pilot multitask and multimodal glaucoma dataset, called Harvard Glaucoma Detection and Progression (Harvard-GDP) Dataset, for computer vision researchers to study this topic. In Table 1, we have listed major public glaucoma datasets. Most of them are fundus photo datasets, while there are two major OCT datasets previously published [24, 80]. In comparison, our dataset has the *largest* patient numbers among all datasets with patient numbers available. In addition, our dataset is the *first* and the only one with progression forecasting task data, which is a way more clinically significant

Table 1: Public Glaucoma Datasets.

Study	Imaging Modality	Sample Size	Glaucoma Detection	Progression Forecasting	Multimodal	Label Source	Accessibility
LAG [38]	Fundus	NA (4,854)	✓	✗	✗	Clinician Assessment	Upon Request
REFUGE [54]	Fundus	NA (1,200)	✓	✗	✗	Clinician Assessment	Web Download
GH020 [6]	Fundus	432 (1,020)	✓	✗	✗	Clinician Assessment	Web Download
ACRIMA [21]	Fundus	NA (705)	✓	✗	✗	Clinician Assessment	Web Download
RIM-ONE [7]	Fundus	169 (485)	✓	✗	✗	Clinician Assessment	Web Download
ORIGA [84]	Fundus	NA (650)	✓	✗	✗	Clinician Assessment	Upon Request
PAPILA [36]	Fundus	244 (488)	✓	✗	✗	Clinician Assessment	Web Download
GOALS [24]	OCT	158 (300)	✓	✗	✗	Clinician Assessment	Web Download
GAMMA [80]	OCT	300 (300)	✓	✗	✓	Clinician Assessment	Web Download
Ours	OCT	1,000 (1,000)	✓	✓	✓	Functional Test	Web Download

task than glaucoma detection itself. Compared with the other two OCT datasets, our dataset is superior in patient numbers (1,000 versus up to 300), task versatility (glaucoma detection and progression forecasting versus glaucoma detection only), and availability of objective visual function tests (i.e. VF test). Furthermore, since our dataset is from the US population, the racial representation in our data is more diverse than the two OCT datasets [24, 80] and the large LAG dataset [38], which consists of mainly Asian patients. Therefore, our dataset can be potentially used for medical fairness learning studies, especially since we know Black people have more than doubled glaucoma prevalence than other races, which is a significant disparity [62, 25].

2. Related Work

Glaucoma Detection. Vast majority of glaucoma detection works using deep learning are based on the fundus photo, which is a 2D snapshot of the patient’s retina [37, 57, 33, 14, 42, 1, 64, 15, 76, 75]. However, the 3D imaging modality OCT scan has replaced fundus photos to become the de facto standard for structural damage assessment in glaucoma. Compared with the deep learning glaucoma detection using fundus photos, deep learning glaucoma detection works using OCT scans generally have much smaller sample sizes, which are up to several thousand [15, 60, 5, 35, 24, 80, 49]. As deep learning models are typically data-hungry, there is a clear need to develop new approaches to deal with the limited labeled data issue. It is known that semi-supervised learning can improve prediction accuracy by leveraging unlabeled data in model training. So far, apart from several glaucoma detection papers using standard semi-supervised learning models [20, 2, 23, 79, 84, 9], no novel semi-supervised learning model has been developed with an application focus on glaucoma.

Progression Forecasting. Compared with the glaucoma detection task, progression forecasting is a much more clinically significant task for patients. If we could have an accurate progression forecasting model to inform clinicians of which patient will progress or not, then the clinicians can accordingly decide if a patient should be treated more aggressively with invasive surgeries, which may have a strong side effect, or more conservatively with eye drops. So far, due to the data scarcity issue, we have only found several recent papers using retinal imaging data (whether fundus

or OCT) to forecast progression [52, 16, 32]. The sample sizes of these works are only up to several hundred. Imaginably, we have not found any prior studies applying semi-supervised learning for progression forecasting.

Semi-Supervised Learning. Semi-Supervised Learning (SSL) aims to optimize a model with both labeled and unlabeled data. The current SSL methods can be divided into pseudo-labeling-based methods [61, 63, 13, 40] and/or consistency-based methods [73, 68, 11, 10, 41, 43]. Consistency-based SSL methods optimize the model with standard classification losses (e.g., cross-entropy) on labeled images and regularizes/minimizes the prediction outputs of different weakly and/or strongly augmented views of unlabeled images, where these views are constructed using different types of image augmentations, such as flip-and-shift, color sharpness, and contrast, etc. Even though consistency-based SSL methods achieve SOTA results in many computer vision benchmarks [73, 68, 11, 10], these methods significantly rely on the orderly design of augmentation functions that requires domain knowledge to design proper image augmentation strategies, which are even more challenging for multimodal modeling. Prior pseudo-labeling SSL methods [61, 63, 13, 40, 56] typically utilize a model that trained with the labeled subset to estimate the pseudo labels of the confident unlabeled samples and take these pseudo-labeled data to re-train the model. As mentioned previously, pseudo-labeling-based SSL approaches is a general SSL learning framework for different data modalities naturally. More importantly, as mentioned in Section 1, since the structural differences between glaucoma and non-glaucoma (similarly for progression versus non-progression) are obscured as shown in Figure 1, existing pseudo-labeling SSL methods [61, 63, 13, 40, 56] might not be able to provide sufficient guidance to select correct pseudo labeled samples to benefit the classifier training. In this aspect, we propose a generalization-reinforced SSL model by selecting optimal pseudo-labeled samples for classifier training that explicitly optimize generalization error with a policy net stochastically generates pseudo labels.

Reinforcement Learning. Reinforcement learning studies an agent interacts with the environment and learns an optimal policy to tackle decision-making problems effectively [70, 8, 26, 65, 66]. Reinforcement learning has been applied to semi-supervised learning setting for real-world applications, e.g. domain adaptation [39]. However, how to utilize reinforcement learning techniques to determine pseudo labels for improving the generalization of classification models remains under exploration.

3. Multimodal and Multitask Dataset

3.1. Data Collection and Quality Control

This work was approved by the institutional review board (IRB) and adhered to the tenets of the Declaration of

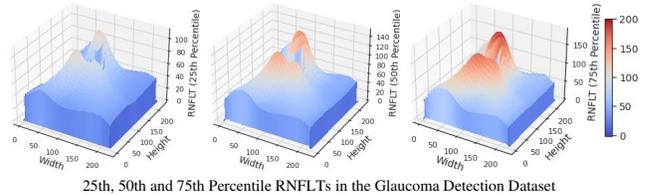


Figure 2: Pointwise RNFLT distribution at different percentiles.

Helsinki. Because of the retrospective nature of the study, the IRB waived the need for informed consent of patients. The glaucoma patients are from Mass Eye and Ear of Harvard Medical School, which were tested between 2010 and 2021. The RNFLT (0 – 350 microns) is from the OCT report provided by the OCT device manufacturer (Cirrus, Carl Zeiss Meditec, Jena, Germany) calculated as the vertical distance between the internal limiting membrane boundary and nerve fiber layer boundary. The RNFLT measurement is the de facto structural measurement for clinicians to diagnose glaucoma. Each two-dimensional 224×224 RNFLT map is generated within the $6 \times 6 \text{ mm}^2$ peripapillary area around the optic nerve head. The VF testing is performed using the Humphrey Field Analyzer with the 24-2 protocol, which measures the patient’s eye’s visual spatial sensitivity with a radius of 24 degrees from the fixation. Data quality control based on standard clinical practice guideline following the manufacturer’s recommendation was applied to exclude unreliable data: (1) OCT scans with signal strength less than 6 (10 represents the best imaging quality) were excluded; (2) VF tests with fixation losses $> 33\%$, false positive rates $> 20\%$ or false negative rates $> 20\%$ were excluded. Each VF measurement is represented as a vector of 52 total deviation (TD) values between -38 dB and 26 dB . In this work, we refer RNFLT map and VF test as two data modalities which constitute a multimodal dataset. Note that, the modifier word “Harvard” from the Harvard GDP Dataset only indicates that our dataset is from the Department of Ophthalmology of Harvard Medical School and does not imply an endorsement, sponsorship, or assumption of responsibility by either Harvard University or Harvard Medical School as a legal identity.

3.2. Data Analysis

The glaucoma detection dataset contains 15,725 training samples, 1,047 validation samples, and 4,287 samples from 5,640, 376, and 1,505 glaucoma patients, respectively. Half of the samples are unlabeled. Among the half samples with labels, 45.2% of the samples are glaucoma, while the rest are non-glaucoma. The progression forecasting dataset includes 5,889 training samples, 431 validation samples, and 1,547 test samples from 1,292, 86, and 345 glaucoma patients, respectively. Among the half samples with progression outcome labels, 24.0% and 2.6% of the samples are TD Progression and MD Fast Progression, respectively. The

two clinically accepted criteria to define glaucoma progression: (1) TD Progression: eyes with at least three locations with TD slope ≤ -1 dB; (2) MD Fast Progression: eyes with MD slope ≤ -1 dB [78]. The MD Fast Progression is more detrimental for the patients than TD progression [3].

Figure 2 presents the pointwise distributions of 25th, 50th, and 75th percentile RNFLT maps in the train (validation set included) and test sets, which demonstrate two obvious superotemporal and inferotemporal RNFL bundles with their damages are known closely related to the visual function loss. As in the two representative examples shown in Figure 1, RNFLT maps with similar appearances may have different labels. This suggests the class boundary may be obscure to determine, which naturally prompts us to adopt semi-supervised solutions to address the glaucoma detection and progression forecasting problems which fully leverage the abundant information in unlabeled samples to mitigate the challenges.

The released dataset including 1,000 patients is a random subset from the entire dataset used to build the glaucoma detection and progression forecasting models in this paper, it naturally inherits the original data characteristics. The released glaucoma detection dataset contains 1,000 samples from 1,000 patients, where each sample includes two data modalities (i.e. RNFLT map and VF test) and respective glaucoma labels. The progression forecasting dataset is a subset from the released glaucoma detection dataset, which entails 500 patients with progression labels and patient demographics including age, sex, and race. For all glaucoma patients, the age distribution of 25th, 50th, and 75th percentiles are 54.3, 64.5, and 72.8 years, respectively. The male and female account for 56% and 44% of the 1,000 patients, respectively. Among the 1,000 patients, the White, Black, and Asian groups account for 76.4%, 14.9%, and 8.7%, respectively. To the best of our knowledge, this is the first public glaucoma dataset with the patient demographic information available, which carries great potential to enable interested researchers to conduct and evaluate fairness learning studies.

4. Problem Formalization & Background

In this section, we formally introduce the problems of glaucoma detection and progression forecasting.

Denote $\mathcal{D}^l = \{(x_i, y_i) | 1 \leq i \leq N^l\}$ be a labeled dataset, where $x \in \mathbb{R}^n$ are the RNFLT maps and $y \in \mathcal{Y}$ are corresponding glaucoma/progression labels. Also, we have a stack of unlabeled RNFLT maps, i.e. $\mathcal{D}^{ul} = \{x_j | 1 \leq j \leq N^{ul}\}$. We assume that we have a classification model $f : \mathbb{R}^n \xrightarrow{\theta} \mathcal{Y}$ (f_θ for short), where θ are the parameters. To update the classification model with the labeled data, a predefined loss function $\ell : \mathcal{Y} \rightarrow \mathcal{Y}$ is used to gauge the discrepancy between the predictions and the ground-truth labels such that the classification model can be updated along

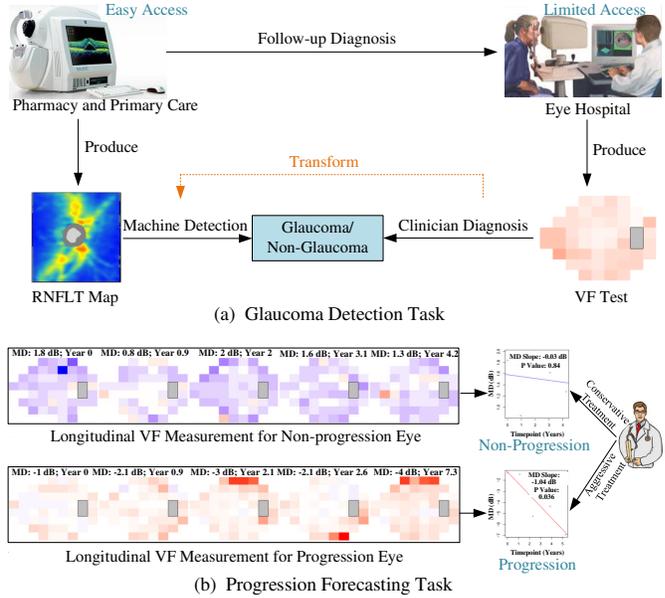


Figure 3: The definitions of (a) glaucoma detection task; and (b) progression forecasting task. In the VF test, blue indicates normal visual function while red indicates abnormal visual function.

the direction of minimizing the loss, i.e. $\min_{\theta} \ell(f_{\theta}(x), y)$.

As shown in Figure 3 (a), the practical utility scenario of glaucoma detection with deep learning is to use OCT scans that can be easily measured in the setting of primary care or pharmacy to predict glaucoma diagnosis defined by visual field tests that typically only can be measured in the setting of an ophthalmology clinic or eye hospital with specialized technicians. The clinical relevance of progression forecasting as shown in Figure 3 (b) with deep learning is that clinicians can better decide if a patient should be treated more conservatively with eye drops or aggressively with invasive surgeries. As suggested by conventional semi-supervised learning methods [20, 2, 40, 68, 11], pseudo labels \tilde{y} are predicted by a supervisor model $\pi : \mathbb{R}^n \xrightarrow{\omega} \mathcal{Y}$ with the unlabeled data \tilde{x} , i.e. $\tilde{y} = \pi(\tilde{x}; \omega)$. Similar to the supervised learning setting, the unlabeled data with the pseudo labels can participate in the supervised learning process, namely $\min_{\theta} \ell(f_{\theta}(\tilde{x}), \tilde{y})$.

However, as shown in Figure 4, the natural images associated with different visual concepts are relatively easy to be distinguished. In contrast, the differences in RNFLT maps between glaucoma versus non-glaucoma (similarly for progression versus non-progression) can be very subtle due to large retinal physiological inter-subject variations. There might not be clear and adequate guidance in the learning process to determine the underlying labels of the unlabeled RNFLT maps. Instead, a more sensible way to better determine the pseudo labels is to apply a policy to obtain pseudo labels. Then, we observe how these pseudo labels affect

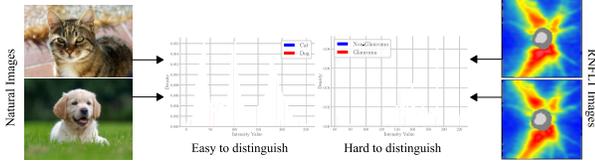


Figure 4: A challenge in glaucoma detection and progression forecasting. Unlike the classification with natural images, where the images with different labels are likely visually distinguishable, RNFLT images with different labels (e.g. glaucoma vs. non-glaucoma and progression vs non-progression) might not be easily distinguishable.

the training process of the classifier and adjust the policy to predict the pseudo labels such that the generalizability of the classifier learned with the pseudo labels can be maximized. In short, the pseudo labels predicted by the supervisor model should be aligned with the improvement of the generalizability of the classifier on a set of unseen data \mathcal{D}^{unseen} . This objective can be formulated as

$$\min_{\{\tilde{y}\}} \frac{1}{|\mathcal{D}^{unseen}|} \sum_{(x,y) \in \mathcal{D}^{unseen}} \ell(f_{\theta^*}(x), y) \quad (1)$$

where $\theta^* \in \arg \min_{\theta} \frac{1}{|\mathcal{D}^{l+ul}|} \sum_{(x,y) \in \mathcal{D}^{l+ul}} \ell(f_{\theta}(x), y)$,

\mathcal{D}^{l+ul} is the union of labeled dataset \mathcal{D} and unlabeled dataset with pseudo labels \mathcal{D}^{ul} .

5. Generalization-Reinforced Pseudo Supervisor

Different from prior-work semi-supervised learning methods [68, 11, 10, 53], the proposed method explicitly leverages the empirical generalizability of the classifier that is quantified with unseen RNFLT maps (e.g. the ones from the validation set) to optimize the policy of predicting the pseudo labels. The empirical generalizability can be viewed as a type of informative reward as the goal of training a classifier is to find one that can generalize well on the unseen test data drawn from the underlying distribution.

Figure 5 shows how the proposed generalization-reinforced pseudo supervisor predicts pseudo labels and is trained according to the observed sets of states (maps), actions (pseudo labels), and rewards (the empirical generalizability with a mini-batch of validation samples). The proposed pseudo supervisor is based on the policy gradient [70], which is one type of generic reinforcement learning methods [8, 26, 65].

Here we formally introduce the concepts related to the proposed pseudo supervisor. An RNFLT map x_t at time step t can be viewed as a state $s_t \in \mathcal{S}$, while the predicted pseudo label \tilde{y}_t can be viewed as an action $a_t \in \mathcal{A}$. A reward function $r_t : \mathcal{S} \times \mathcal{A} \rightarrow \mathbf{R}$ is defined on the state-action space. The proposed pseudo supervisor that employs

a policy to determine pseudo labels is denoted as $\pi_{\omega} : \mathcal{S} \rightarrow \mathcal{P}(\mathcal{A})$, where $\mathcal{P}(\mathcal{A})$ are the probability measures on \mathcal{A} and ω are the parameters of the proposed pseudo supervisor.

In the first place, the classifier takes labeled samples for training. The unlabeled samples will be fed into the pseudo supervisor for predicting pseudo labels. Before updating the weights of the classifier, the empirical generalizability of classifier f_{θ_t} can be quantified with validation samples $(x^{(val)}, y^{(val)})$, that is, $\ell(f_{\theta_t}(x^{(val)}), y^{(val)})$ ($\ell_t^{(val)}$). Specifically, ℓ is the widely-used cross-entropy loss for classification. After updating the weights of the classifier with unlabeled samples and corresponding pseudo labels, given the same mini-batch of validation samples, the empirical generalizability of the updated classifier can be re-quantified again, i.e. $\ell(f_{\theta_{t+1}}(x^{(val)}), y^{(val)})$ ($\ell_{t+1}^{(val)}$). Then, the discrepancy between $\ell_t^{(val)}$ and $\ell_{t+1}^{(val)}$ is informative to measure the effect of updating the classifier with the unlabeled samples and the pseudo labels. In this sense, it is desired to use in the reward function, i.e.

$$r_t = \max(\exp(\ell_t^{(val)} - \ell_{t+1}^{(val)}) - 1, 0) \quad (2)$$

In Eq. (2), if $\ell_{t+1}^{(val)}$ is far less than $\ell_t^{(val)}$, it implies that the unlabeled samples and pseudo labels lead to a positive effect on the learning process such that the reward increases adaptively.

A trajectory of states, actions, and rewards within a time window β will be saved. The expected return of policy optimization can be written as

$$J(\pi_{\omega}) = \mathbb{E}_{\pi_{\omega}} \left(\sum_{k=0}^{\beta} \gamma^k r_k \right) \quad (3)$$

where $\gamma \in [0, 1]$ is the discount rate for the numerical stability of reinforcement learning, and β is the time window for updating π_{ω} periodically. According to the policy gradient theorem [70], the expectation of the sample gradient is equal to the actual gradient

$$\nabla_{\pi_{\omega}} J(\pi_{\omega}) = \mathbb{E}_{\pi_{\omega}} \left[\sum_{k=0}^{\beta} \gamma^k r_{t+k+1} \nabla_{\pi_{\omega}} \log \pi_{\omega}(a_t | s_t) \right] \quad (4)$$

Then the policy parameterization can be updated according to the gradient ascend rule

$$\pi_{\omega} \leftarrow \pi_{\omega} + \eta_{\pi} \nabla_{\pi_{\omega}} J(\pi_{\omega}) \quad (5)$$

where η_{π} is the learning rate of π_{ω} .

6. Experiments

6.1. Set-Up & Implementation Detail

To optimize our models, we employ the AdamW optimizer [46] and train all models for 10 epochs in all experiments conducted in this study. For supervised classification

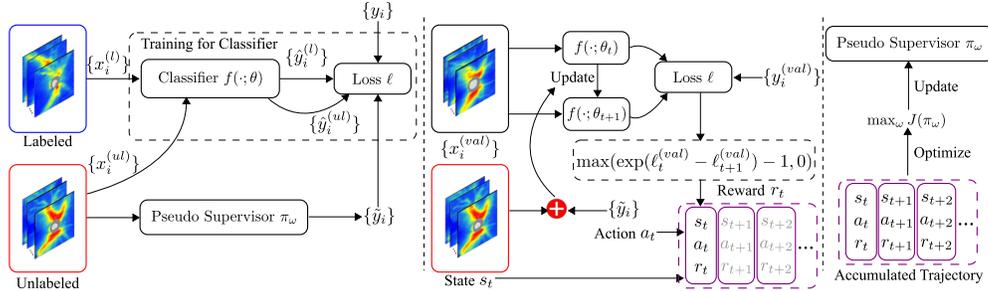


Figure 5: An overview of the proposed generalization-reinforced pseudo supervisor in the semi-supervised learning framework. The learning framework consists of the process of training the classifier, the process of measuring the empirical generalizability of the classifier to be used as rewards, and the process of updating the proposed pseudo supervisor to predict pseudo labels that can maximize the classifier’s empirical generalizability.

models for glaucoma detection, we use a learning rate of 4×10^{-5} and weight decay of 0, and for progression forecasting, we use a learning rate of 2×10^{-5} and weight decay of 0. These hyperparameters are determined to achieve the best baseline performance based on our experiments. The proposed method is trained using a learning rate of 4×10^{-5} and weight decay of 0. The batch sizes used in the supervised model, semi-supervised models, and the proposed method vary due to the different semi-supervised learning contexts [73, 11, 10, 68, 53]. We adopt the batch sizes that result in the best performance for each respective context. To ensure consistency with previous work, we apply data augmentation techniques to the samples in the semi-supervised learning models, including MeanTeacher [73], MixMatch [11], ReMixMatch [10], FixMatch [68], and DASO [53]. Specifically, we utilize random crop and resize, as well as random horizontal flip. However, we do not apply these data augmentation techniques to the supervised baseline or the proposed method.

For the glaucoma detection task, we train all models with 7,860 labeled samples and 7,865 unlabeled samples for training. We use 1,047 and 4,287 labeled samples for validation and testing, respectively. For the progression forecasting task, all the models are trained with 2,944 labeled samples and 2,945 unlabeled samples for training. We use 431 and 1,547 labeled samples for validation and testing, respectively. For multi-modality modeling, we simply concatenate the RNFLT images and VF matrices for joint learning using the same architecture as our single-modality model and apply this model to the progression forecasting task. The model is trained and evaluated using the same train-val-test split as our single-modality experiments.

We use the accuracy, F1 score, and AUC on the test sets as our evaluation matrices. The code is written in PyTorch [55] and we use one RTX A6000 GPU for all experiments.

6.2. Baseline Methods

Our Supervised baseline is a standard ImageNet pre-trained EfficientNet classifier [72] learning with cross-

entropy using only labeled images. For Semi-Supervised baselines, we choose MeanTeacher [73], MixMatch [11], ReMixMatch [10], FixMatch [68], and DASO [53] due to their previous powerful performances. MeanTeacher [73] is one of the initial methods to enforce models to predict consistent pseudo predictions for labeled and unlabeled data, which construct a target-generating teacher model by averaging the model weights. MixMatch [11], ReMixMatch [10], and FixMatch [68] are consistency-based SSL methods that are based on the data augmentations (i.e., weak and strong augmentation) to make consistent classification predictions between different views of unlabeled data. DASO [53] suggests that those methods fail for imbalanced datasets and proposes a pseudo-labeling framework based on different consistency SSL models by adaptively mixing the linear and semantic pseudo-labels to mitigate the overall bias between majority and minority classes. We adopt the FixMatch version of DASO, following the default setting reported in the paper [53]. We implement all baseline methods using Pytorch [55] from DASO’s official codes and conducted experiments under the same codebase and experimental protocols for fair comparison. All models are equipped with the same EfficientNetV2-S [72] backbone and their hyper-parameters are specifically tuned to suit our Glaucoma detection and progression forecasting tasks.

6.3. Results for Glaucoma Detection

In Table 2, we evaluate our Pseudo-Sup model against multiple SOTA baselines on the glaucoma detection task. Our model achieves the best performance in terms of accuracy, F1 score, and AUC score. We compute the mean performances over 5 different runs. Notably, our model achieves such results without any specific design of image perturbations/augmentations, guaranteeing its simplicity and applicability to different tasks. Our method surpasses the previous SOTA method DASO by 6.8% of accuracy, 11.1% of F1 score, and 3.7% of AUC, respectively. The best-performing consistency method FixMatch still worse by 0.4% of accuracy, 1.2% of F1 score, and 1.5%

Table 2: Performance on the cross-sectional data for the **glaucoma detection** task. The baseline method is the supervised model trained with labeled samples only, while pseudo-sup stands for the proposed *pseudo supervisor* method.

Model	Acc \uparrow	F1 \uparrow	AUC \uparrow
Baseline	0.8131 \pm 0.0039	0.8057 \pm 0.0033	0.8693 \pm 0.0022
MeanTeacher	0.8148 \pm 0.0053	0.8053 \pm 0.0072	0.8541 \pm 0.0040
MixMatch	0.6392 \pm 0.0161	0.6273 \pm 0.0059	0.8451 \pm 0.0062
ReMixMatch	0.6741 \pm 0.0019	0.6716 \pm 0.0018	0.8442 \pm 0.0032
DASO	0.7444 \pm 0.0115	0.6929 \pm 0.0050	0.8371 \pm 0.0118
FixMatch	0.8077 \pm 0.0032	0.7912 \pm 0.0021	0.8574 \pm 0.0019
Pseudo-Sup	0.8124 \pm 0.0025	0.8038 \pm 0.0046	0.8727 \pm 0.0006
Pseudo-Sup+Aug	0.8224 \pm 0.0101	0.8128 \pm 0.0122	0.8818 \pm 0.0020

of AUC than our Pseudo-Sup model, indicating the effectiveness of our approach and the challenges of designing appropriate augmentations for RNFLT maps from previous approaches. We also apply **weak** augmentations (similar to FixMatch) to our Pseudo-Sup (Pseudo-Sup+Aug) leading to improved performances. Please note that consistency methods (e.g., DASO, FixMatch, etc.) employ **both** weak and strong augmentations.

6.4. Results for Progression Forecasting

Unimodal Results. We further evaluate our model on progression forecasting, a more clinically significant task than glaucoma detection, as shown in Table 3. Accurate forecasting of glaucoma progression plays a vital role in prescribing appropriate treatment plans to efficiently prevent further vision losses and is important for assisting the development of new pharmaceutical medicines to treat glaucoma. Adding weak augmentations further improves the results for both progression tasks. Please note that similar to glaucoma detection experiments, we also compute the mean performances over 5 different runs.

For TD Progression, our model achieves the best 74.47% of mean AUC, respectively. When compared with consistency-based approaches, our method improves between 0.5% - 2.1% in AUC. For MD Fast Progression, our approach consistently achieves the best AUC performance, illustrating the effectiveness and accuracy of our method regardless of any glaucoma progression criteria. Given that most of the glaucoma patients are non-progressive, our progression dataset for TD Progression and MD Fast Progression are both highly imbalanced (i.e., most of the samples are labeled as non-progression). Therefore, it is worth noting that our method also outperforms the previous imbalanced SSL SOTA DASO by a large margin in all measures under such imbalanced scenarios.

Multimodal Results. One of the major contributions of this work is to introduce a multi-modal dataset for glaucoma progression forecasting. To the best of our knowledge, this is the *first* work to explore and benchmark SSL approaches using multimodal model inputs. In Table 5, we show the

Table 3: Performance on the longitudinal data with **single modality** RNFLT for the **progression forecasting** task. The baseline method is the supervised model trained with labeled samples only, while Pseudo-Sup represents the proposed *pseudo supervisor* method.

Models	AUC \uparrow (TD)	AUC \uparrow (MD Fast)
Baseline	0.7367 \pm 0.0126	0.6677 \pm 0.0017
MeanTeacher	0.7367 \pm 0.0126	0.6677 \pm 0.0017
MixMatch	0.7238 \pm 0.0098	0.6518 \pm 0.0069
ReMixMatch	0.7390 \pm 0.0124	0.6595 \pm 0.0083
DASO	0.7231 \pm 0.0188	0.6787 \pm 0.0148
FixMatch	0.7395 \pm 0.0152	0.6747 \pm 0.0054
Pseudo-Sup	0.7447 \pm 0.0084	0.7004 \pm 0.0218
Pseudo-Sup+Aug	0.7477 \pm 0.0074	0.7215 \pm 0.0162

progression forecasting results using two modalities (RNFLT and VF). Our model achieves the best performance on all evaluation measures. To keep its simplicity and adaptability, we concatenate the RNFLT and the up-scaled VF together as the inputs of our model. Please note that consistency SSL approaches that are based on image augmentations are not included in this table because of the lack of design of augmentation techniques in multimodal medical data. The SOTA results with multimodal glaucoma data enable more accurate and applicable potential computer-aided diagnosis (CAD) systems in real clinical settings.

6.5. Supervised Benchmarks on Released Data

In Table 6, we show the supervised classification results for the progression forecasting tasks with multiple SOTA supervised CNN and transformer baseline methods, including VGG [67], ResNet [27], ResNext [82], WideResNet [83], EfficientNet [71], ConvNext [45], ViT [22], and Swin Transformers [44]. This benchmark is conducted on our future data release with 500 patients upon acceptance, of which 400 patients are used for training and the remaining are used for testing. To the best of our knowledge, this is the first supervised glaucoma progression forecasting benchmark, which aims to provide broader impacts for the computer vision and medical imaging communities to build clinically effective computer-aided diagnosis systems. Moreover, the results are reported with both single RNFLT modality and multi-modality (RNFLT and VF) samples. It is worth noting that the accuracy and F1 are the same on MD fast progression because the testing set only contains 100 samples under highly imbalanced class distributions. For more experimental details regarding this supervised progression forecasting benchmark and the supervised results on future release data with 1,000 patients for the glaucoma detection task, please refer to the Supplemental Material.

Table 4: Performance comparisons for **glaucoma detection** task on **LAG dataset**.

Model	Acc \uparrow	F1 \uparrow	AUC \uparrow
MeanTeacher	0.8280	0.8279	0.9082
MixMatch	0.5020	0.3378	0.9075
ReMixMatch	0.6980	0.4785	0.8369
DASO	0.7971	0.7906	0.9164
FixMatch	0.8366	0.8297	0.9198
Pseudo-Sup	0.9160	0.9158	0.9858

Table 5: Performance on the longitudinal data with **two modalities** (RNFLT and VF) for the **progression forecasting** task. The baseline method is the supervised model trained with labeled samples only, while pseudo-sup stands for the proposed *pseudo supervisor* method. Please note that most methods in Table 3 are based on weak/strong augmentations [68, 11, 10] requiring RGB channels, which are not adaptable for multi-modal learning.

Models	AUC (TD)	AUC (MD Fast)
Baseline	0.7343 \pm 0.0098	0.6814 \pm 0.0081
MeanTeacher	0.7436 \pm 0.0133	0.6697 \pm 0.0035
Pseudo-Sup	0.7517 \pm 0.0041	0.7265 \pm 0.0095

6.6. Performances on LAG

In Table 4, we show the performance comparisons on LAG dataset [37]. LAG is a large-scale fundus imaging dataset for glaucoma detection. For this experiment, we use 3,854 images for training with half with labels and the remaining without labels, and 1,000 images for testing, with the size of 224×224 . For fair comparisons, we use the same hyper-parameters and experimental setups as our other benchmarks. It is worth noting that our approach surpasses the SOTA methods by a large margin and achieves the highest 98.58% AUC with only half-labeled data, illustrating the effectiveness and applicability of our proposed Pseudo-Sup in different data and tasks.

6.7. Ablation Study

According to Section 5, the time window β and discount rate γ are the most important hyperparameters affecting the proposed pseudo supervisor. The ablation studies of the two hyperparameters are shown in Figure 6. We can see that the proposed pseudo supervisor behaviors stably with various β on AUC. Specifically, when $\beta = 50$ leads to better performance. The consistent pattern can be also observed in the analysis of γ . The discount factor is used to achieve a trade-off between the rewards in the distant future and the rewards in the near future. If it is 0, reinforcement learning agents would only look for an optimal policy for the first action. As shown in Figure 6 (b), we find the optimal discount rate at 0.9, which implies that the states, actions, and rewards in the progression task are coherent to a certain degree.

Table 6: Performance of different **supervised methods** on the longitudinal data with single/multi-modality for the **progression forecasting tasks** on the released 500 glaucoma progression data.

Model	TD Progression			MD Fast Progression		
	Acc \uparrow	F1 \uparrow	AUC \uparrow	Acc \uparrow	F1 \uparrow	AUC \uparrow
VGG [67]	0.78/0.80	0.75/0.74	0.84/0.79	0.98/0.98	0.49/0.49	0.78/0.68
ResNet [27]	0.75/0.74	0.64/0.67	0.74/0.75	0.98/0.98	0.49/0.49	0.71/0.59
ResNext [82]	0.78/0.77	0.65/0.68	0.78/0.76	0.98/0.98	0.49/0.49	0.61/0.80
WideResNet [83]	0.77/0.79	0.71/0.75	0.77/0.80	0.98/0.98	0.49/0.49	0.71/0.64
EfficientNet [71]	0.73/0.78	0.66/0.73	0.76/0.79	0.98/0.98	0.49/0.49	0.55/0.66
ConvNext [45]	0.74/0.81	0.62/0.77	0.78/0.81	0.98/0.98	0.49/0.49	0.84/0.62
ViT [22]	0.73/0.77	0.54/0.71	0.68/0.79	0.98/0.98	0.49/0.49	0.92/0.69
Swin [44]	0.71/0.77	0.49/0.72	0.68/0.77	0.98/0.98	0.49/0.49	0.87/0.57

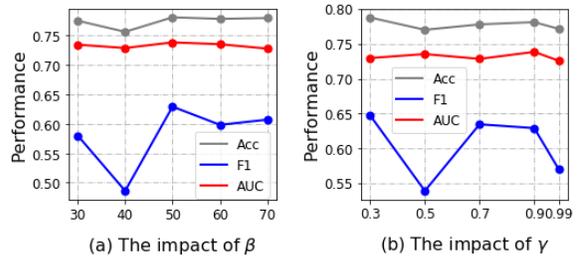


Figure 6: Ablation study of hyperparameters' impact on TD Progression forecasting.

6.8. Experiment on Racial Groups

The patients in our dataset are from three racial groups (White (0.88 AUC), Black (0.83 AUC), and Asian (0.85 AUC)). In our experiments, we observe that black patients have the lowest AUC yet the highest glaucoma prevalence [62, 25]. Therefore, our dataset is well-suited for future fairness studies.

7. Conclusion

In this paper, we have made two contributions. First, we have developed a generalization-reinforced semi-supervised learning model termed pseudo supervisor for glaucoma detection and progression forecasting. Specifically, the pseudo supervisor predicts pseudo labels following a policy that optimizes the generalization error in training the classifier. Our pseudo supervisor model overall demonstrates superior performance over a number of SOTA methods. Second, we release a multimodal and multi-task dataset for glaucoma detection and progression forecasting with SOTA 3D OCT imaging data. Our dataset is the *largest* by far for glaucoma detection with 3D OCT data and the *first* for progression forecasting task. More importantly, our dataset entails demographic information including gender and race, which is generally unavailable in existing public glaucoma datasets. Given the diverse racial representation in our dataset as detailed in Section 3, our dataset is well equipped for potential fairness learning studies, especially considering the significant glaucoma prevalence disparity between Black and other races [62, 25].

Acknowledgment

This work was supported by NIH R00 EY028631, Research to Prevent Blindness International Research Collaborators Award, Alcon Young Investigator Grant, NIH R01 EY030575, and NIH P30 EY003790.

References

- [1] Jin Mo Ahn, Sangsoo Kim, Kwang-Sung Ahn, Sung-Hoon Cho, Kwan Bok Lee, and Ungsoo Samuel Kim. A deep learning model for the detection of both advanced and early glaucoma using fundus photography. *PloS one*, 13(11):e0207982, 2018. 1, 2
- [2] Manal Al Ghamdi, Mingqi Li, Mohamed Abdel-Mottaleb, and Mohamed Abou Shousha. Semi-supervised transfer learning for convolutional neural networks for glaucoma detection. In *ICASSP 2019-2019 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, pages 3812–3816. IEEE, 2019. 2, 4
- [3] Andrew J Anderson, Abadh K Chaurasia, Ajay Sharma, Amisha Gupta, Shikha Gupta, Anurag Khanna, and Viney Gupta. Comparison of rates of fast and catastrophic visual field loss in three glaucoma subtypes. *Investigative ophthalmology & visual science*, 60(1):161–167, 2019. 4
- [4] Paul H Artes and Balwantray C Chauhan. Longitudinal changes in the visual field and optic disc in glaucoma. *Progress in retinal and eye research*, 24(3):333–354, 2005. 1
- [5] Ryo Asaoka, Hiroshi Murata, Kazunori Hirasawa, Yuri Fujino, Masato Matsuura, Atsuya Miki, Takashi Kanamoto, Yoko Ikeda, Kazuhiko Mori, Aiko Iwase, et al. Using deep learning and transfer learning to accurately diagnose early-onset glaucoma from macular optical coherence tomography images. *American journal of ophthalmology*, 198:136–145, 2019. 1, 2
- [6] Muhammad Naseer Bajwa, Gur Amrit Pal Singh, Wolfgang Neumeier, Muhammad Imran Malik, Andreas Dengel, and Sheraz Ahmed. G1020: A benchmark retinal fundus image dataset for computer-aided glaucoma detection. In *2020 International Joint Conference on Neural Networks (IJCNN)*, pages 1–7. IEEE, 2020. 2
- [7] Francisco José Fumero Batista, Tinguaro Diaz-Aleman, Jose Sigut, Silvia Alayon, Rafael Arnay, and Denisse Angel-Pereira. Rim-one dl: A unified retinal image database for assessing glaucoma using deep learning. *Image Analysis & Stereology*, 39(3):161–167, 2020. 2
- [8] Jonathan Baxter and Peter L Bartlett. Infinite-horizon policy-gradient estimation. *Journal of Artificial Intelligence Research*, 15:319–350, 2001. 3, 5
- [9] Mohammed El Amine Bechar, Nesma Settouti, Vincent Barra, and Mohamed Amine Chikh. Semi-supervised super-pixel classification for medical images segmentation: application to detection of glaucoma disease. *Multidimensional Systems and Signal Processing*, 29(3):979–998, 2018. 2
- [10] David Berthelot, Nicholas Carlini, Ekin D Cubuk, Alex Kurakin, Kihyuk Sohn, Han Zhang, and Colin Raffel. Mixmatch: Semi-supervised learning with distribution matching and augmentation anchoring. In *International Conference on Learning Representations*, 2019. 2, 3, 5, 6, 8
- [11] David Berthelot, Nicholas Carlini, Ian Goodfellow, Nicolas Papernot, Avital Oliver, and Colin A Raffel. Mixmatch: A holistic approach to semi-supervised learning. *Advances in Neural Information Processing Systems*, 32:5049–5059, 2019. 2, 3, 4, 5, 6, 8
- [12] Igor I Bussel, Gadi Wollstein, and Joel S Schuman. Oct for glaucoma diagnosis, screening and detection of glaucoma progression. *British Journal of Ophthalmology*, 98(Suppl 2):ii15–ii19, 2014. 1
- [13] Paola Cascante-Bonilla et al. Curriculum labeling: Revisiting pseudo-labeling for semi-supervised learning. *arXiv preprint arXiv:2001.06001*, 2020. 2, 3
- [14] Xiangyu Chen, Yanwu Xu, Shuicheng Yan, Damon Wing Kee Wong, Tien Yin Wong, and Jiang Liu. Automatic feature learning for glaucoma detection based on deep learning. In *International conference on medical image computing and computer-assisted intervention*, pages 669–677. Springer, 2015. 1, 2
- [15] Mark Christopher, Akram Belghith, Christopher Bowd, James A Proudfoot, Michael H Goldbaum, Robert N Weinreb, Christopher A Girkin, Jeffrey M Liebmann, and Linda M Zangwill. Performance of deep learning architectures and transfer learning for detecting glaucomatous optic neuropathy in fundus photographs. *Scientific reports*, 8(1):1–13, 2018. 1, 2
- [16] Mark Christopher, Akram Belghith, Robert N Weinreb, Christopher Bowd, Michael H Goldbaum, Luke J Saunders, Felipe A Medeiros, and Linda M Zangwill. Retinal nerve fiber layer features identified by unsupervised machine learning on optical coherence tomography scans predict glaucoma progression. *Investigative ophthalmology & visual science*, 59(7):2748–2756, 2018. 1, 2, 3
- [17] Anne Louise Coleman. Advances in glaucoma treatment and management: surgery. *Investigative ophthalmology & visual science*, 53(5):2491–2494, 2012. 1
- [18] Ronan Conlon, Hady Saheb, and Iqbal Ike K Ahmed. Glaucoma treatment trends: a review. *Canadian Journal of Ophthalmology*, 52(1):114–124, 2017. 1
- [19] David P Crabb, Nicholas D Smith, Fiona C Glen, Robyn Burton, and David F Garway-Heath. How does glaucoma look?: patient perception of visual field loss. *Ophthalmology*, 120(6):1120–1126, 2013. 1
- [20] Andres Diaz-Pinto, Adrián Colomer, Valery Naranjo, Sandra Morales, Yanwu Xu, and Alejandro F Frangi. Retinal image synthesis and semi-supervised learning for glaucoma assessment. *IEEE transactions on medical imaging*, 38(9):2211–2218, 2019. 2, 4
- [21] Andres Diaz-Pinto, Sandra Morales, Valery Naranjo, Thomas Köhler, Jose M Mossi, and Amparo Navea. Cnns for automatic glaucoma assessment using fundus images: an extensive validation. *Biomedical engineering online*, 18(1):1–19, 2019. 2
- [22] Alexey Dosovitskiy, Lucas Beyer, Alexander Kolesnikov, Dirk Weissenborn, Xiaohua Zhai, Thomas Unterthiner, Mostafa Dehghani, Matthias Minderer, Georg Heigold, Sylvain Gelly, et al. An image is worth 16x16 words: Trans-

- formers for image recognition at scale. *arXiv preprint arXiv:2010.11929*, 2020. 7, 8
- [23] Rui Fan, Christopher Bowd, Nicole Brye, Mark Christopher, Robert N Weinreb, David Kriegman, and Linda Zangwill. One-vote veto: Semi-supervised learning for low-shot glaucoma diagnosis. *arXiv preprint arXiv:2012.04841*, 2020. 2
- [24] Huihui Fang, Fei Li, Huazhu Fu, Junde Wu, Xiulan Zhang, and Yanwu Xu. Dataset and evaluation algorithm design for goals challenge. In *International Workshop on Ophthalmic Medical Image Analysis*, pages 135–142. Springer, 2022. 1, 2
- [25] David S Friedman, Henry D Jampel, Beatriz Munoz, and Sheila K West. The prevalence of open-angle glaucoma among blacks and whites 73 years and older: the salisbury eye evaluation glaucoma study. *Archives of ophthalmology*, 124(11):1625–1630, 2006. 2, 8
- [26] Mohammad Ghavamzadeh and Sridhar Mahadevan. Hierarchical policy gradient algorithms. *Computer Science Department Faculty Publication Series*, page 173, 2003. 3, 5
- [27] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 770–778, 2016. 7, 8
- [28] Patrick Herbert, Kaihua Hou, Chris Bradley, Michael V Boland, Pradeep Y Ramulu, Mathias Unberath, and Jithin Yohannan. Forecasting risk of future rapid glaucoma worsening using early visual field, optical coherence tomography and clinical data. *Investigative Ophthalmology & Visual Science*, 63(7):2294–2294, 2022. 1
- [29] Donald C Hood and Randy H Kardon. A framework for comparing structural and functional measures of glaucomatous damage. *Progress in retinal and eye research*, 26(6):688–710, 2007. 1
- [30] Sen Hu, Nicholas D Smith, Luke J Saunders, and David P Crabb. Patterns of binocular visual field loss derived from large-scale patient data from glaucoma clinics. *Ophthalmology*, 122(12):2399–2406, 2015. 1
- [31] David Huang, Eric A Swanson, Charles P Lin, Joel S Schuman, William G Stinson, Warren Chang, Michael R Hee, Thomas Flotte, Kenton Gregory, Carmen A Puliafito, et al. Optical coherence tomography. *science*, 254(5035):1178–1181, 1991. 1
- [32] Alireza Kamalipour, Sasan Moghimi, Pooya Khosravi, Vahid Mohammadzadeh, Takashi Nishida, Eleonora Micheletti, Jo-Hsuan Wu, Golnoush Mahmoudinezhad, Elizabeth HF Li, Mark Christopher, et al. Combining oct and oct-angiography longitudinal data for the detection of visual field progression in glaucoma: Longitudinal oct/octa data to detect glaucoma progression. *American Journal of Ophthalmology*, 2022. 1, 2, 3
- [33] Stuart Keel, Jinrong Wu, Pei Ying Lee, Jane Scheetz, and Mingguang He. Visualizing deep learning models for the detection of referable diabetic retinopathy and glaucoma. *JAMA ophthalmology*, 137(3):288–292, 2019. 1, 2
- [34] Philip J Kellman and Thomas F Shipley. A theory of visual interpolation in object perception. *Cognitive psychology*, 23(2):141–221, 1991. 1
- [35] Yuka Kihara, Giovanni Montesano, Andrew Chen, Nishani Amerasinghe, Chrysostomos Dimitriou, Aby Jacob, Almira Chabi, David P Crabb, and Aaron Y Lee. Policy-driven, multimodal deep learning for predicting visual fields from the optic disc and oct imaging. *Ophthalmology*, 2022. 1, 2
- [36] Oleksandr Kovalyk, Juan Morales-Sánchez, Rafael Verdú-Monedero, Inmaculada Sellés-Navarro, Ana Palazón-Cabanes, and José-Luis Sancho-Gómez. Papila: Dataset with fundus images and clinical data of both eyes of the same patient for glaucoma assessment. *Scientific Data*, 9(1):1–12, 2022. 2
- [37] Liu Li, Mai Xu, Xiaofei Wang, Lai Jiang, and Hanruo Liu. Attention based glaucoma detection: a large-scale database and cnn model. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 10571–10580, 2019. 1, 2, 8
- [38] Liu Li, Mai Xu, Xiaofei Wang, Lai Jiang, and Hanruo Liu. Attention based glaucoma detection: A large-scale database and cnn model. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, June 2019. 2
- [39] Bingyu Liu, Yuhong Guo, Jieping Ye, and Weihong Deng. Selective pseudo-labeling with reinforcement learning for semi-supervised domain adaptation. *arXiv preprint arXiv:2012.03438*, 2020. 3
- [40] Fengbei Liu, Yu Tian, Yuanhong Chen, Yuyuan Liu, Vasileios Belagiannis, and Gustavo Carneiro. Acpl: Anticurriculum pseudo-labelling for semi-supervised medical image classification. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 20697–20706, 2022. 2, 3, 4
- [41] Fengbei Liu, Yu Tian, Filipe R Cordeiro, Vasileios Belagiannis, Ian Reid, and Gustavo Carneiro. Self-supervised mean teacher for semi-supervised chest x-ray classification. In *International Workshop on Machine Learning in Medical Imaging*, pages 426–436. Springer, 2021. 3
- [42] Hanruo Liu, Liu Li, I Michael Wormstone, Chunyan Qiao, Chun Zhang, Ping Liu, Shuning Li, Huaizhou Wang, Dapeng Mou, Ruiqi Pang, et al. Development and validation of a deep learning system to detect glaucomatous optic neuropathy using fundus photographs. *JAMA ophthalmology*, 137(12):1353–1360, 2019. 1, 2
- [43] Yuyuan Liu, Yu Tian, Yuanhong Chen, Fengbei Liu, Vasileios Belagiannis, and Gustavo Carneiro. Perturbed and strict mean teachers for semi-supervised semantic segmentation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 4258–4267, 2022. 3
- [44] Ze Liu, Yutong Lin, Yue Cao, Han Hu, Yixuan Wei, Zheng Zhang, Stephen Lin, and Baining Guo. Swin transformer: Hierarchical vision transformer using shifted windows. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 10012–10022, 2021. 7, 8
- [45] Zhuang Liu, Hanzi Mao, Chao-Yuan Wu, Christoph Feichtenhofer, Trevor Darrell, and Saining Xie. A convnet for the 2020s. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 11976–11986, 2022. 7, 8
- [46] Ilya Loshchilov and Frank Hutter. Decoupled weight decay regularization. In *International Conference on Learning Representations*, 2019. 5

- [47] Rizwan Malik, William H Swanson, and David F Garway-Heath. ‘structure–function relationship’ in glaucoma: past thinking and current concepts. *Clinical & experimental ophthalmology*, 40(4):369–380, 2012. 1
- [48] Felipe A Medeiros, Alessandro A Jammal, and Atalie C Thompson. From machine to machine: an oct-trained deep learning algorithm for objective quantification of glaucomatous damage in fundus photographs. *Ophthalmology*, 126(4):513–521, 2019. 1
- [49] Parmita Mehta, Christine A Petersen, Joanne C Wen, Michael R Banitt, Philip P Chen, Karine D Bojkian, Catherine Egan, Su-In Lee, Magdalena Balazinska, Aaron Y Lee, et al. Automated detection of glaucoma with interpretable machine learning using clinical data and multimodal retinal images. *American Journal of Ophthalmology*, 231:154–169, 2021. 1, 2
- [50] MJ Morgan and RJ Watt. Mechanisms of interpolation in human spatial vision. *Nature*, 299(5883):553–555, 1982. 1
- [51] Jacqueline M Nelson-Quigg, Kimberly Cello, and Chris A Johnson. Predicting binocular visual field sensitivity from monocular visual field results. *Investigative ophthalmology & visual science*, 41(8):2212–2221, 2000. 1
- [52] Kouros Nouri-Mahdavi, Vahid Mohammadzadeh, Alessandro Rabiolo, Kiumars Edalati, Joseph Caprioli, and Siamak Yousefi. Prediction of visual field progression from oct structural measures in moderate to advanced glaucoma. *American Journal of Ophthalmology*, 226:172–181, 2021. 1, 2, 3
- [53] Youngtaek Oh, Dong-Jin Kim, and In So Kweon. Distribution-aware semantics-oriented pseudo-label for imbalanced semi-supervised learning. *arXiv preprint arXiv:2106.05682*, 2021. 5, 6
- [54] José Ignacio Orlando, Huazhu Fu, João Barbosa Breda, Karel van Keer, Deepti R Bathula, Andrés Diaz-Pinto, Ruogu Fang, Pheng-Ann Heng, Jeyoung Kim, JoonHo Lee, et al. Refuge challenge: A unified framework for evaluating automated methods for glaucoma assessment from fundus photographs. *Medical image analysis*, 59:101570, 2020. 2
- [55] Adam Paszke, Sam Gross, Francisco Massa, Adam Lerer, James Bradbury, Gregory Chanan, Trevor Killeen, Zeming Lin, Natalia Gimelshein, Luca Antiga, Alban Desmaison, Andreas Kopf, Edward Yang, Zachary DeVito, Martin Raison, Alykhan Tejani, Sasank Chilamkurthy, Benoit Steiner, Lu Fang, Junjie Bai, and Soumith Chintala. Pytorch: An imperative style, high-performance deep learning library. In H. Wallach, H. Larochelle, A. Beygelzimer, F. d’Alché-Buc, E. Fox, and R. Garnett, editors, *Advances in Neural Information Processing Systems 32*, pages 8024–8035. Curran Associates, Inc., 2019. 6
- [56] Hieu Pham, Zihang Dai, Qizhe Xie, and Quoc V Le. Meta pseudo labels. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 11557–11568, 2021. 2, 3
- [57] Sonia Phene, R Carter Dunn, Naama Hammel, Yun Liu, Jonathan Krause, Naho Kitade, Mike Schaeckermann, Rory Sayres, Derek J Wu, Ashish Bora, et al. Deep learning and glaucoma specialists: the relative importance of optic disc features to predict glaucoma referral in fundus photographs. *Ophthalmology*, 126(12):1627–1639, 2019. 1, 2
- [58] Harry A Quigley. Number of people with glaucoma worldwide. *British journal of ophthalmology*, 80(5):389–393, 1996. 1
- [59] Harry A Quigley and Aimee T Broman. The number of people with glaucoma worldwide in 2010 and 2020. *British journal of ophthalmology*, 90(3):262–267, 2006. 1
- [60] An Ran Ran, Carol Y Cheung, Xi Wang, Hao Chen, Lu-yang Luo, Poemen P Chan, Mandy OM Wong, Robert T Chang, Suria S Mannil, Alvin L Young, et al. Detection of glaucomatous optic neuropathy with spectral-domain optical coherence tomography: a retrospective training and validation deep-learning analysis. *The Lancet Digital Health*, 1(4):e172–e182, 2019. 1, 2
- [61] Mamshad Nayeem Rizve et al. In defense of pseudo-labeling: An uncertainty-aware pseudo-label selection framework for semi-supervised learning. In *International Conference on Learning Representations*, 2020. 2, 3
- [62] Alicja R Rudnicka, Shahrul Mt-Isa, Christopher G Owen, Derek G Cook, and Deborah Ashby. Variations in primary open-angle glaucoma prevalence by age, gender, and race: a bayesian meta-analysis. *Investigative ophthalmology & visual science*, 47(10):4254–4261, 2006. 2, 8
- [63] Weiwei Shi, Yihong Gong, Chris Ding, Zhiheng MaXiaoyu Tao, and Nanning Zheng. Transductive semi-supervised deep learning using min-max features. In *Proceedings of the European Conference on Computer Vision (ECCV)*, pages 299–315, 2018. 2, 3
- [64] Naoto Shibata, Masaki Tanito, Keita Mitsuhashi, Yuri Fujino, Masato Matsuura, Hiroshi Murata, and Ryo Asaoka. Development of a deep residual learning algorithm to screen for glaucoma from fundus photography. *Scientific reports*, 8(1):1–9, 2018. 1, 2
- [65] David Silver, Guy Lever, Nicolas Heess, Thomas Degris, Daan Wierstra, and Martin Riedmiller. Deterministic policy gradient algorithms. In *International conference on machine learning*, pages 387–395. PMLR, 2014. 3, 5
- [66] David Silver, Julian Schrittwieser, Karen Simonyan, Ioannis Antonoglou, Aja Huang, Arthur Guez, Thomas Hubert, Lucas Baker, Matthew Lai, Adrian Bolton, Yutian Chen, Timothy P. Lillicrap, Fan Hui, Laurent Sifre, George van den Driessche, Thore Graepel, and Demis Hassabis. Mastering the game of go without human knowledge. *Nat.*, 550(7676):354–359, 2017. 3
- [67] Karen Simonyan and Andrew Zisserman. Very deep convolutional networks for large-scale image recognition. *arXiv preprint arXiv:1409.1556*, 2014. 7, 8
- [68] Kihyuk Sohn, David Berthelot, Nicholas Carlini, Zizhao Zhang, Han Zhang, Colin A Raffel, Ekin Dogus Cubuk, Alexey Kurakin, and Chun-Liang Li. Fixmatch: Simplifying semi-supervised learning with consistency and confidence. *Advances in Neural Information Processing Systems*, 33, 2020. 2, 3, 4, 5, 6, 8
- [69] Yi Sun, Aiming Chen, Minjie Zou, Yichi Zhang, Ling Jin, Yi Li, Danying Zheng, Guangming Jin, and Nathan Congdon. Time trends, associations and prevalence of blindness and vision loss due to glaucoma: an analysis of observational data from the global burden of disease study 2017. *BMJ open*, 12(1):e053805, 2022. 1

- [70] Richard S. Sutton, David A. McAllester, Satinder Singh, and Yishay Mansour. Policy gradient methods for reinforcement learning with function approximation. In *Advances in Neural Information Processing Systems 12, [NIPS Conference, Denver, Colorado, USA, November 29 - December 4, 1999]*, pages 1057–1063, 1999. 3, 5
- [71] Mingxing Tan and Quoc Le. Efficientnet: Rethinking model scaling for convolutional neural networks. In *International conference on machine learning*, pages 6105–6114. PMLR, 2019. 7, 8
- [72] Mingxing Tan and Quoc V. Le. Efficientnetv2: Smaller models and faster training. In *Proceedings of the 38th International Conference on Machine Learning, ICML 2021, 18-24 July 2021, Virtual Event*, volume 139 of *Proceedings of Machine Learning Research*, pages 10096–10106. PMLR, 2021. 6
- [73] Antti Tarvainen and Harri Valpola. Mean teachers are better role models: Weight-averaged consistency targets improve semi-supervised deep learning results. *Advances in neural information processing systems*, 30, 2017. 2, 3, 6
- [74] Yih-Chung Tham, Xiang Li, Tien Y Wong, Harry A Quigley, Tin Aung, and Ching-Yu Cheng. Global prevalence of glaucoma and projections of glaucoma burden through 2040: a systematic review and meta-analysis. *Ophthalmology*, 121(11):2081–2090, 2014. 1
- [75] Yu Tian, Fengbei Liu, Guansong Pang, Yuanhong Chen, Yuyuan Liu, Johan W VERJANS, and Rajvinder Singh. Self-supervised multi-class pre-training for unsupervised anomaly detection and segmentation in medical images. 2021. 2
- [76] Yu Tian, Guansong Pang, Fengbei Liu, Yuanhong Chen, Seon Ho Shin, Johan W Verjans, Rajvinder Singh, and Gustavo Carneiro. Constrained contrastive distribution learning for unsupervised anomaly detection and localisation in medical images. In *International Conference on Medical Image Computing and Computer-Assisted Intervention*, pages 128–140. Springer, 2021. 2
- [77] Rohit Varma, Paul P Lee, Ivan Goldberg, and Sameer Kotak. An assessment of the health and economic burdens of glaucoma. *American journal of ophthalmology*, 152(4):515–522, 2011. 1
- [78] Eija Vesti, Chris A Johnson, and Balwantray C Chauhan. Comparison of different methods for detecting glaucomatous visual field progression. *Investigative ophthalmology & visual science*, 44(9):3873–3879, 2003. 4
- [79] Xi Wang, Hao Chen, An-Ran Ran, Luyang Luo, Poemen P Chan, Clement C Tham, Robert T Chang, Suria S Mannil, Carol Y Cheung, and Pheng-Ann Heng. Towards multi-center glaucoma oct image screening with semi-supervised joint structure and function multi-task learning. *Medical Image Analysis*, 63:101695, 2020. 2
- [80] Junde Wu, Huihui Fang, Fei Li, Huazhu Fu, Fengbin Lin, Jiongcheng Li, Lexing Huang, Qinji Yu, Sifan Song, Xingxing Xu, et al. Gamma challenge: glaucoma grading from multi-modality images. *arXiv preprint arXiv:2202.06511*, 2022. 1, 2
- [81] Jo-Hsuan Wu, Takashi Nishida, Robert N Weinreb, and Jou-Wei Lin. Performances of machine learning in detecting glaucoma using fundus and retinal optical coherence tomography images: A meta-analysis. *American Journal of Ophthalmology*, 2021. 1
- [82] Saining Xie, Ross Girshick, Piotr Dollár, Zhuowen Tu, and Kaiming He. Aggregated residual transformations for deep neural networks. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 1492–1500, 2017. 7, 8
- [83] Sergey Zagoruyko and Nikos Komodakis. Wide residual networks. *arXiv preprint arXiv:1605.07146*, 2016. 7, 8
- [84] Zhuo Zhang, Feng Shou Yin, Jiang Liu, Wing Kee Wong, Ngan Meng Tan, Beng Hai Lee, Jun Cheng, and Tien Yin Wong. Origa-light: An online retinal fundus image database for glaucoma analysis and research. In *2010 Annual international conference of the IEEE engineering in medicine and biology*, pages 3065–3068. IEEE, 2010. 2