

Self-Supervised Domain Adaptation for Diabetic Retinopathy Grading using Vessel Image Reconstruction

Duy M. H. Nguyen¹, Truong T. N. Mai², Ngoc T. T. Than³, Alexander Prange¹,
and Daniel Sonntag^{1,4}

¹ German Research Center for Artificial Intelligence (DFKI)
Saarland Informatics Campus Saarbrücken, Germany

² Department of Multimedia Engineering, Dongguk University, South Korea

³ Byers Eye Institute, Stanford University, United States

⁴ Oldenburg University, Germany

Abstract. This paper investigates the problem of domain adaptation for diabetic retinopathy (DR) grading. We learn invariant target-domain features by defining a novel self-supervised task based on retinal vessel image reconstructions, inspired by medical domain knowledge. Then, a benchmark of current state-of-the-art unsupervised domain adaptation methods on the DR problem is provided. It can be shown that our approach outperforms existing domain adaptation strategies. Furthermore, when utilizing entire training data in the target domain, we are able to compete with several state-of-the-art approaches in final classification accuracy just by applying standard network architectures and using image-level labels.

Keywords: Domain Adaption, Diabetic Retinopathy, Self-Supervised Learning, Deep Learning, Interactive Machine Learning

1 Introduction

Diabetic retinopathy (DR) is a type of ocular disease that can cause blindness due to damaged blood vessels in the back of the eye. The causes of DR are high blood pressure and high blood sugar concentration, which are very common in modern lifestyles [40]. People with diabetes usually have higher risks of developing DR. In fact, one-third of diabetes patients show the symptoms of diabetic retinopathy according to recent studies [42]. Therefore, early detection of DR is critical to ensure successful treatment. Unfortunately, detecting and grading diabetic retinopathy in practice is a laborious task, and DR is difficult to diagnose at an early stage even for professional ophthalmologists. As a result, developing a precise automatic DR diagnostic device is both necessary and advantageous.

Automated DR diagnosis systems take retinal images (fundus images) and yield DR grades. In the common retinal imaging dataset of DR, the grades of DR can be categorized into five stages [6]: 0 - no DR, 1 - mild DR, 2 - moderate

DR, 3 - severe DR, and 4 - proliferative DR. Specifically, the severity of DR is determined by taking the numbers, sizes, and appearances of lesions into account. For instance, figure 1 provides an illustration of five DR grades in the Kaggle DR dataset [12]. As can be seen, the characteristics of DR grades are complex in both structure and texture aspects. Therefore, automated diagnosis systems are required to be capable of extracting meaningful visual features from retinal images for precise DR grading.




Fig. 1. Illustration of different DR grades.

With the success of deep learning, several CNN-based methods for DR grading of retinal images have been proposed. The paper from 2016 [6] introduces the development and validation of a deep learning algorithm for detection of diabetic retinopathy—with high sensitivity and specificity when compared with manual grading by ophthalmologists for identifying diabetic retinopathy. Jiang et al. [11] also propose an ensemble of conventional deep learning methods to increase the predictive performance of automated DR grading. Lin et al. [14] in other direction introduce a joint model for lesion detection as well as DR identification, in which the DR is inferred from the fusion of original images and lesion information predicted by an attention-based network. Similarly, Zhou et al. in [41] apply a two-step strategy: first produce a multi-lesion mask by using a semantic segmentation component, then the severity of DR is graded by exploiting the lesion mask. Recently, Wu et al. [35] address the problem in a similar way, the classification is performed by employing pixel-level segmentation maps.

While recent works have demonstrated its effectiveness when trained and tested on a single dataset, they often suffer from the domain adaptation problem in practice. In particular, medical images in clinical applications are acquired from devices of different manufactures that vary in many aspects, including imaging modes, image processing algorithms, and hardware components. Therefore, the performance of a trained network from a particular source domain can dramatically decrease when applied to a different target domain. One possible way to overcome this barrier is to collect and label new samples in the target domain, which is necessary for fine-tune trained networks. Nevertheless, this task is laborious and expensive especially with medical images, as the data are limited and labeling requires extreme caution. As a result, it is highly desirable to develop an algorithm that can adapt well in the new domain without additional

labeled data for training. Such an approach is known as unsupervised domain adaptation.

In this paper, we propose a self-supervised method to reduce domain shift in the fundus images’ distribution by learning the invariant feature representations. To this end, feature extraction layers are trained by using both labeled data from the source domain and a self-supervised task on a target domain by defining image reconstruction tasks around retinal vessel positions. Moreover, we also incorporate additional restricted loss functions throughout the training phase to encourage the acquired features to be consistent with the main objective.

At a glance, we make three main contributions. First, we address the domain adaptation problem for DR grading on fundus images using a novel self-supervised approach motivated by medical domain knowledge. Second, we provide a benchmark of current state-of-the-art unsupervised domain adaptation methods on the DR problem. Finally, we show that our approach when using fully training data in the target domain obtains competitive performance just by employing standard network architectures and using image-level labels.




Fig. 2. Illustration of our vessel segmentation reconstruction-based SSL. (a) input image x^t , (b) its vessel segmentation y_v^t , (c) binary masks B^t (inside green rectangles) sampled along edges of y_v^t in (b) (red rectangles). The image regions inside B^t are removed to define \hat{y}^t and asking the encoder-decoder network to reconstruct them given the remaining pixels in \hat{x}^t .

2 Related Work

Over the last decade, research in domain adaptation has achieved remarkable results. Tzeng et al. [32] propose a deep domain confusion technique to minimize the maximum mean discrepancy, a non-parametric metric of distribution divergence proposed by Gretton et al. [5], so that the divergence between two distributions is reduced. The algorithm developed by Sun et al. [28] is an extension of their previous work [27], in which CNNs are employed to learn a nonlinear transformation for correlation alignment. Recently, Wang et al. [33] have presented a domain adaptation algorithm for screening normal and abnormal retinopathy

in optical coherence tomography (OCT) images. The system consists of several complex components guided by the Wasserstein distance [23] to extract invariant representations across different domains.

In other directions, researchers have employed generative adversarial networks (GANs) to learn better invariant features. Tzeng et al. [31] combine discriminative modeling with untied weight sharing and a GAN-based loss to create an adversarial discriminative domain algorithm. Shen et al.’s algorithm [23] extracts domain invariant feature representations by optimizing the feature extractor network, which minimizes the Wasserstein distance trained in an adversarial manner between the source and target domains. In a different way, Long et al. [15] design a conditional domain adversarial network by exploiting two strategies, namely multilinear conditioning, to capture the cross-domain covariance, and entropy conditioning, to ensure the transferability.

Our method in this paper follows the self-supervised learning (SSL) approach [13], which is recently an active research direction due to its effectiveness in learning feature representations. In particular, SSL refers to a representation learning method where a supervised task is defined over unlabelled data to reduce the data labeling cost and leverage the available unlabelled data. Until now, several algorithms based on SSL have been introduced. The method presented by Xu et al. [38] is a generic network with several kinds of learning tasks in SSL that can adapt to various datasets and diverse applications. In medical image analysis, authors in [1] introduce a SSL pretext task based on context restoration, thereby two isolated small regions are selected randomly and swap their positions. A deep network is then trained to recover original orders in input images. Unfortunately, these prior works are mostly designed in the same domain. Recently, Xiao et al. [36] have pioneered to apply the SSL method for domain adaptation problems. Specifically, target-domain-aware features are learned from unlabeled data for image classification through an image rotation-based pretext task trained by a unified encoder for both source and target domains.

Difference w.r.t. Previous Work: Our method follows Xiao et al. [36]; however, we make the following modifications for our setting. First, rather than a rotation task like [36], we study medical domain knowledge to create a novel SSL prediction task, i.e., vessel segmentation reconstruction that has a solid connection to the severity of diabetic retinopathy [6, ?]. Second, a two-player procedure is integrated through a discriminate network to ensure mission regions generated in SSL tasks look realistic and consistent with the image context. As a results, our objective function has more constraints on learned features when compared to [36].

3 Method

3.1 Overview

Our proposed method aims at learning invariant features across different domains through encoder layers shared to optimize several relevant tasks. In specific,

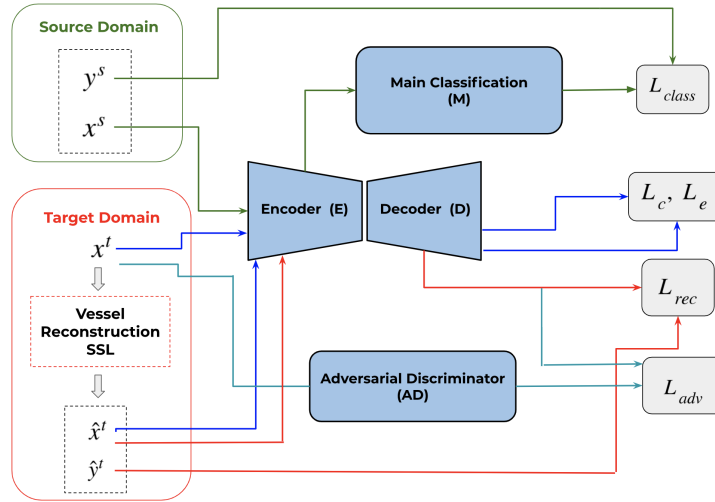


Fig. 3. Overview our proposed unsupervised domain adaptation with vessel reconstruction-based self-supervised learning.

we define labeled images in the source domain $X_s = \{(x_i^s, y_i^s)\}_{i=0}^{N_s}$ with y_i^s is the corresponding label (DR grades) of image x_i^s and N_s is the total of images. In the target domain, we assume that only a set of unlabeled images denoted by $X_t = \{x_i^t\}_{i=0}^{N_t}$ with N_t samples is available. Our framework, which uses labeled X_s and unlabeled X_t for domain adaptation, consists of four distinct blocks: an encoder network E , a decoder network D , an adversarial discriminator AD , and a main classifier M . These blocks are parameterized by $\theta_e, \theta_d, \theta_{ad}$ and θ_m respectively. For each image $x_i^t \in X_t$, we transform it through the self-supervised learning task based on vessel image reconstruction to define a new set $\hat{X}_t = \{(\hat{x}_i^t, \hat{y}_i^t)\}_{i=1}^{N_t}$, which are used to train E and D blocks for predicting removed sub-patch images. To encourage that the reconstructed regions look authentic, the adversarial discriminator AD is integrated through the two-player game learning procedure for distinguishing generated and ground-truth samples. Finally, the block M is built on top of the encoder layer E and acts as the main classification task. We describe below each aforementioned architecture in detail.

3.2 Retinal Vessel Reconstruction-based SSL

According to medical protocol [6, 7], the severity of DR can be predicted by observing the number and size of related lesion appearances and complications. While their positions tend to cluster near vessel positions, we use this attribute to create a new SSL task that forces learnt feature representation to capture such lesions.

Given a sample $x^t \in X_t$, we extract its vessel segmentation image $y_v^t = f(x_i^t)$ with $f(\cdot)$ is a trained deep network (Figure 2a and 2b). In this work, we use

$f(\cdot)$ as a proposed architecture in [29]. Let B^t is a binary mask corresponding to the dropped image region in x^t , with a value of 1 if a pixel was dropped and 0 for input pixels. Unlike related works [1, 3], we generate region masks in B^t by randomly sampling sub-patch images along vessel positions in y_v^t as indicated in Figure 2c. We then define a new pair of samples:

$$\hat{x}^t = (1 - B^t) \odot x^t, \hat{y}^t = B^t \odot x^t \quad (1)$$

where \odot is the element-wise product operation.

Reconstruction Loss We train a context encoder F formed from the encoder E and the decoder D to reconstruct target regions \hat{y}^t given the input \hat{x}^t . A normalized $L2$ distance is employed as our reconstruction objective function:

$$L_{rec} = \min_{\theta_e, \theta_d} \mathbb{E}_{x^t \in X^t} \|B^t \odot F(\hat{x}^t) - \hat{y}^t\|_2^2 \quad (2)$$

Adversarial Loss The objective function L_{rec} takes into account the overall construction of the missing region and agreement with its context, but tends to average together the multiple forms in predictions. We thus adapt the adversarial discriminator AD as [21, 34] to make the predictions of the context encoder F look real through selecting similar instances from the target distribution. The joint min-max objective function of discriminator AD and generator F is:

$$L_{adv} = \min_{\theta_e, \theta_d} \max_{\theta_{ad}} \mathbb{E}_{x^t \in X^t} [\log(D(x^t)) + \log(1 - D(F(\hat{x}^t)))] \quad (3)$$

By jointly optimizing L_{rec} and L_{adv} , we encourage the output of the context encoder F to look realistic on the entire prediction, not just the missing regions as in L_{rec} .

3.3 Relevant Features from SSL

Main Classification Loss In our framework, the block M takes the feature representation from the encoder E to predict a corresponding label y^t for each image x^t in the target domain. The network M and encoder E are trained with labeled data in the source domain by optimizing the classification problem:

$$L_{class} = \min_{\theta_e, \theta_m} \mathbb{E}_{x^s, y^s \in X^s} [-\log p(y^s | x^s)] \quad (4)$$

where $p(y^s | x^s)$ is a conditional probability distribution of y^s given x^s parameterized by E and M networks.

Constrained Features from SSL While the SSL task is designed to encourage the encoder E to capture invariant features across different domains and pay attention to vessel positions, there is no guarantee of the compatibility between this SSL task and the main classification target. Inspired from prior works in semi-supervised learning [16, 37], we adapt two additional loss constraints on the

feature representation generated by the SSL \hat{x}^t , the input x^t , and the target label y^t :

$$L_c = \min_{\theta_e, \theta_m} \mathbb{E}_{x^t \in X^t} \mathbb{E}_{\hat{x}^t \in \hat{X}^t} [D_{KL}(\hat{p}(y^t|x^t)||p(y^t|\hat{x}^t))] \quad (5)$$

$$L_e = \min_{\theta_e, \theta_m} \mathbb{E}_{x^t \in X^t} \left[- \sum_{y^t} p(y^t|x^t) \log(p(y^t|x^t)) \right] \quad (6)$$

where D_{KL} is the Kullback-Leibler consistency loss [16, 37], $\hat{p}(y^t|x^t)$ is a fixed copy of the current $p(y^t|x^t)$ with parameters θ_e, θ_m , it means that $\hat{p}(y^t|x^t)$ is only used for each inference step and the gradient is not propagated through them.

Intuitively, the consistency objective function L_c forces the feature representation in E to be insensitive to data augmentation in defined SSL task while the objective L_e penalizes uncertain predictions, leading to more discriminative representations. However, the equations L_c, L_e require labels y^t in the target domain to optimize, which are assumed to be not available in our unsupervised domain adaption. We address this challenge by integrating pseudo-labels y^t generated by predictions using E and M blocks and updating it progressively after each training step.

Overall Objective Function In summary, our overall objective function is:

$$L = L_{class} + \lambda_{rec}L_{rec} + \lambda_{adv}L_{adv} + \lambda_cL_c + \lambda_eL_e \quad (7)$$

where $\lambda_{rec}, \lambda_{adv}, \lambda_c, \lambda_e$ are coefficients of corresponding objective functions. Due to the generative adversarial function in L_{adv} , L is the min-max objective problem. We adapt the alternative optimization strategy to first update parameters $\theta_e, \theta_d, \theta_m$, second update θ_{ad} and repeating this process until convergence. In our experiment, we use feature extraction layers from ResNet-50 [8] for both the encoder E , decoder D and adversarial discriminator AD . These layers are shaped in certain architectural constraints as in [22]. For the main classification M , we adapt a simple average pooling followed by a fully connected layer.

4 Experiments and Results

4.1 Evaluation Method

We assess our method, denoted as VesRec-SSL, in two DR grading scenarios: unsupervised domain adaption (UDA) and conventional classification problems. In the first case, all UDA methods are trained using both supervised samples in the source domain and unlabeled samples in the target domain. The performance is then evaluated using the target domain’s testing set. In the second case, we train and test in the same domain, i.e., the training set’s labeled images are utilized in the training step, and trained networks are measured on the remaining data. For the latter case, our method may be viewed as a pre-training phase [1, 18]; thereby, obtained weights after training VesRec-SSL will be used in the fine-tuning step using partially or completely supervised training samples.

4.2 Dataset and Metrics

We employ two DR-graded retinal image datasets, Kaggle EyePACS [12] and FGADR [42], for training and testing with DR gradings from 0-4 (Figure 1). We follow the splitting standard in EyePACS with 35126 training images and 53576 testing images. With the FGADR dataset, we can only access 1842 images (SegSet) out of a total of 2842 images at the moment due to data privacy. Because there is no specific train/test on the SegSet, we apply 3-fold cross-validation to compute the final performance. For quantitative metrics, we use classification accuracy and Quadratic Weighted Kappa (Q. W. Kappa) [42].

4.3 Performance of Unsupervised Adaption Methods

In this task, we choose one dataset as the source domain and the other as the target domain. We provide a benchmark of three different methods in literature: Xiao et al.’s Rotation-based SSL [36], Long et al.’s CDAN and CDAN-E [15]. For fairly comparison, we choose ResNet50 as the backbone network for all methods. The quantitative evaluation is shown in Table 1 where “EyePACS \rightarrow FGADR (SegSet)” indicates the source domain is EyePACS the target domain is FGADR restricted on SegSet with 1842 images, and similarly for “FGADR (SegSet) \rightarrow EyePACS”. In practice, we found that training baselines directly in our setting is not straightforward due to the imbalance among grading types and the complexity of distinguishing distinct diseases. Therefore, we applied the following training methods:

- First, we only activate the main classification loss using fully supervised samples in the source domain in the initial phase and training until the model converges. Next, auxiliary loss functions will be activated, and the network is continued to train in the latter phase.
- Second, we apply the progressive resizing technique introduced in the fast.ai¹, and the DAWNbench challenge [2] in which the network is trained with smaller images at the beginning, and obtained weights are utilized for training another model with larger images. We use two different resolutions in our setting: 256×256 and 512×512 .
- Finally, the optimal learning rate is automatically chosen by the Cyclical Learning method [25] with the SGD optimizer [4], which sets the learning rate to cyclically change between reasonable boundary values.

As shown in Table 1, our VesRec-SSL outperforms competitors by a remarkable margin in all settings and metrics. For instance, we achieve 2 – 3% more for FGADR and 1 – 3% more for EyePACS, compared to the second competitor CDAN-E. In addition, we can observe that the performance in “FGADR (SegSet) \rightarrow EyePACS” is lower than that in “EyePACS \rightarrow FGADR (SegSet)” in most of the cases. We argue this happens due to the number of training instances in the source domain of “FGADR”, which is much lower than that of “EyePACS”.

¹ <https://course.fast.ai/>

Table 1. Performance of Unsupervised Domain Adaption Methods.

Method	EyePACS \rightarrow FGADR (SegSet)		FGADR (SegSet) \rightarrow EyePACS	
	Acc.	Q.W. Kappa	Acc.	Q.W. Kappa
Rotation-based SSL [36]	0.728	0.672	0.681	0.660
CDAN [15]	0.741	0.685	0.697	0.685
CDAN-E [15]	0.755	0.706	0.702	0.691
VesRec-SSL (Our)	0.782	0.725	0.736	0.702

Table 2. Performance of competitor methods on the DR grading prediction. Red, blue, black, and orange represent the top four best results.

Method	EyePACS	
	Acc.	Q.W. Kappa
VGG-16 [24]	0.836	0.820
ResNet-50 [8]	0.846	0.824
Inception v3 [30]	0.840	0.811
DenseNet-121 [10]	0.854	0.835
Lin et al., [14]	0.867	0.857
Zhou et al., [41]	0.895	0.885
Wu et al., [35]	0.886	0.877
VesRec-SSL (ResNet-50) + 0%	0.736	0.702
VesRec-SSL (ResNet-50) + 50%	0.798	0.774
VesRec-SSL (ResNet-50) + 100%	0.864	0.852
VesRec-SSL (DenseNet-121) + 0%	0.744	0.711
VesRec-SSL (DenseNet-121) + 50%	0.815	0.793
VesRec-SSL (DenseNet-121) + 100%	0.871	0.862
VesRec-SSL (ResNet-50 + DenseNet-121) + 100%	0.891	0.879

4.4 Performance of Baseline Methods on DR Grading Prediction

In this task, we compare our algorithm to the most recent state-of-the-art method reported in [42]. Due to the data privacy on the FGADR dataset, we can only benchmark baselines on the EyePACS dataset. For ablation studies, we also fine-tune our VesRec-SSL with additional 0%, 50%, and 100% labeled data pairs from the target domain. The evaluation results are shown in Table 2. Besides the default backbone with ResNet-50, we consider a variation with DenseNet-121 network for fairly evaluation with two top methods in [41, 35]. Moreover, we also utilize feature extraction layers as average pooling of feature maps obtained from ResNet-50 and DenseNet-121 at the last row and train this network with 100% training data.

The results indicate that without labeled data from the target domain, our two settings perform considerably worse than all baselines trained with fully supervised images. However, by progressively increasing the amount of labeled data from 50% to 100%, we can significantly increase performance. For example, the ResNet-50 with 50% data outperforms the 0% case with approximately 6/7% in Acc/Q.W.Kappa. DenseNet-121 follows a similar pattern, improving 7/8% (50% data), and even with 100% training data, our VesRec-SSL can achieve

the fourth rank in total. Finally, we observe that utilizing both the ResNet-50 and DenseNet-121 backbones can result in a second-rank overall without modifying the network architecture or adding extra pixel-level segmentation maps for relative lesion characteristics as in [35, 41]. In summary, we argue that our method with vessel reconstruction-based SSL has proven effective for domain adaptation under DR grading applications, especially as partial or complete annotations are available.

5 Conclusion

Domain shift is a big obstacle of deep learning-based medical analysis, especially as images are collected by using various devices. In this work, we showed that the unsupervised domain adaptation for diabetic retinopathy grading can benefit from our novel self-supervised learning (SSL) based on the medical vessel image reconstruction tasks. Furthermore, when fully integrating annotation data and simply using standard network architectures, our technique achieves comparable performance to cutting-edge benchmarks. In future work, we consider to extend the SSL task to include related lesion appearances such as microaneurysms (MAs), Hard exudates, and Soft exudates [42] to acquire improved invariant feature representation guided by medical domain knowledge. Moreover, making our network’s predictions understandable and explainable to clinicians is also a crucial question for further investigation based on our recent medical application projects [17, 19, 20, 26]. We also aim to investigate in the direction of information fusion and explainable AI by incorporating multimodal embeddings with Graph Neural Networks [9, 39].

Acknowledgement

This research has been supported by the Ophthalmology-AI project (BMBF, 16SV8639), the Ki-Para-Mi project (BMBF, 01IS19038B), the pAItient project (BMG, 2520DAT0P2), and the Endowed Chair of Applied Artificial Intelligence, Oldenburg University. We would like to thank all student assistants that contributed to the development of the platform, see iml.dfki.de.

References

1. Chen, L., Bentley, P., Mori, K., Misawa, K., Fujiwara, M., Rueckert, D.: Self-supervised learning for medical image analysis using image context restoration. *Medical image analysis* **58**, 101539 (2019)
2. Coleman, C., Narayanan, D., Kang, D., Zhao, T., Zhang, J., Nardi, L., Bailis, P., Olukotun, K., Ré, C., Zaharia, M.: Dawnbench: An end-to-end deep learning benchmark and competition. *Training* **100**(101), 102 (2017)
3. Doersch, C., Gupta, A., Efros, A.A.: Unsupervised visual representation learning by context prediction. In: *Proceedings of the IEEE international conference on computer vision*. pp. 1422–1430 (2015)

4. Goodfellow, I., Bengio, Y., Courville, A.: Deep learning. MIT press (2016)
5. Gretton, A., Borgwardt, K.M., Rasch, M.J., Schölkopf, B., Smola, A.: A kernel two-sample test. *The Journal of Machine Learning Research* **13**(1), 723–773 (2012)
6. Gulshan, V., Peng, L., Coram, M., Stumpe, M.C., Wu, D., Narayanaswamy, A., Venugopalan, S., Widner, K., Madams, T., Cuadros, J., et al.: Development and validation of a deep learning algorithm for detection of diabetic retinopathy in retinal fundus photographs. *Jama* **316**(22), 2402–2410 (2016)
7. Haneda, S., Yamashita, H.: International clinical diabetic retinopathy disease severity scale. *Nihon rinsho. Japanese journal of clinical medicine* **68**, 228–235 (2010)
8. He, K., Zhang, X., Ren, S., Sun, J.: Deep residual learning for image recognition. In: *Proceedings of the IEEE conference on computer vision and pattern recognition*. pp. 770–778 (2016)
9. Holzinger, A., Malle, B., Saranti, A., Pfeifer, B.: Towards multi-modal causability with graph neural networks enabling information fusion for explainable AI. *Information Fusion* **71**, 28–37 (2021)
10. Huang, G., Liu, Z., Van Der Maaten, L., Weinberger, K.Q.: Densely connected convolutional networks. In: *Proceedings of the IEEE conference on computer vision and pattern recognition*. pp. 4700–4708 (2017)
11. Jiang, H., Yang, K., Gao, M., Zhang, D., Ma, H., Qian, W.: An interpretable ensemble deep learning model for diabetic retinopathy disease classification. In: *2019 41st Annual International Conference of the IEEE Engineering in Medicine and Biology Society (EMBC)*. pp. 2045–2048. IEEE (2019)
12. Kaggle: Diabetic retinopathy detection (2015), <https://www.kaggle.com/c/diabetic-retinopathy-detection/data>
13. Kolesnikov, A., Zhai, X., Beyer, L.: Revisiting self-supervised visual representation learning. In: *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*. pp. 1920–1929 (2019)
14. Lin, Z., Guo, R., Wang, Y., Wu, B., Chen, T., Wang, W., Chen, D.Z., Wu, J.: A framework for identifying diabetic retinopathy based on anti-noise detection and attention-based fusion. In: *International Conference on Medical Image Computing and Computer-Assisted Intervention*. pp. 74–82. Springer (2018)
15. Long, M., Cao, Z., Wang, J., Jordan, M.I.: Conditional adversarial domain adaptation. In: *Advances in Neural Information Processing Systems*. pp. 1640–1650 (2018)
16. Miyato, T., Maeda, S.i., Koyama, M., Ishii, S.: Virtual adversarial training: a regularization method for supervised and semi-supervised learning. *IEEE transactions on pattern analysis and machine intelligence* **41**(8), 1979–1993 (2018)
17. Nguyen, D.M., Nguyen, D.M., Vu, H., Nguyen, B.T., Nunnari, F., Sonntag, D.: An attention mechanism using multiple knowledge sources for covid-19 detection from ct images. In: *The Thirty-Fifth AAAI Conference on Artificial Intelligence (AAAI-21), Workshop: Trustworthy AI for Healthcare* (2021)
18. Nguyen, D.M., Nguyen, T.T., Vu, H., Pham, Q., Nguyen, M.D., Nguyen, B.T., Sonntag, D.: TATL: Task agnostic transfer learning for skin attributes detection. *arXiv preprint arXiv:2104.01641* (2021)
19. Nguyen, D.M.H., Ezema, A., Nunnari, F., Sonntag, D.: A visually explainable learning system for skin lesion detection using multiscale input with attention u-net. In: *German Conference on Artificial Intelligence (Künstliche Intelligenz)*. pp. 313–319. Springer (2020)

20. Nunnari, F., Sonntag, D.: A software toolbox for deploying deep learning decision support systems with XAI capabilities. In: Companion of the 2021 ACM SIGCHI Symposium on Engineering Interactive Computing Systems. pp. 44–49 (2021)
21. Pathak, D., Krahenbuhl, P., Donahue, J., Darrell, T., Efros, A.A.: Context encoders: Feature learning by inpainting. In: Proceedings of the IEEE conference on computer vision and pattern recognition. pp. 2536–2544 (2016)
22. Radford, A., Metz, L., Chintala, S.: Unsupervised representation learning with deep convolutional generative adversarial networks. arXiv preprint arXiv:1511.06434 (2015)
23. Shen, J., Qu, Y., Zhang, W., Yu, Y.: Wasserstein distance guided representation learning for domain adaptation. In: Proceedings of the AAAI Conference on Artificial Intelligence. vol. 32 (2018)
24. Simonyan, K., Zisserman, A.: Very deep convolutional networks for large-scale image recognition. arXiv preprint arXiv:1409.1556 (2014)
25. Smith, L.N.: Cyclical learning rates for training neural networks. In: 2017 IEEE winter conference on applications of computer vision (WACV). pp. 464–472. IEEE (2017)
26. Sonntag, D., Nunnari, F., Profitlich, H.J.: The skincare project, an interactive deep learning system for differential diagnosis of malignant skin lesions. technical report. arXiv preprint arXiv:2005.09448 (2020)
27. Sun, B., Feng, J., Saenko, K.: Return of frustratingly easy domain adaptation. In: Proceedings of the AAAI Conference on Artificial Intelligence. vol. 30 (2016)
28. Sun, B., Saenko, K.: Deep coral: Correlation alignment for deep domain adaptation. In: European conference on computer vision. pp. 443–450. Springer (2016)
29. Sun, X., Cao, X., Yang, Y., Wang, L., Xu, Y.: Robust retinal vessel segmentation from a data augmentation perspective. arXiv preprint arXiv:2007.15883 (2020)
30. Szegedy, C., Vanhoucke, V., Ioffe, S., Shlens, J., Wojna, Z.: Rethinking the inception architecture for computer vision. In: Proceedings of the IEEE conference on computer vision and pattern recognition. pp. 2818–2826 (2016)
31. Tzeng, E., Hoffman, J., Saenko, K., Darrell, T.: Adversarial discriminative domain adaptation. In: Proceedings of the IEEE conference on computer vision and pattern recognition. pp. 7167–7176 (2017)
32. Tzeng, E., Hoffman, J., Zhang, N., Saenko, K., Darrell, T.: Deep domain confusion: Maximizing for domain invariance. arXiv preprint arXiv:1412.3474 (2014)
33. Wang, J., Chen, Y., Li, W., Kong, W., He, Y., Jiang, C., Shi, G.: Domain adaptation model for retinopathy detection from cross-domain oct images. In: Medical Imaging with Deep Learning. pp. 795–810. PMLR (2020)
34. Wang, Y., Chen, Y.C., Zhang, X., Sun, J., Jia, J.: Attentive normalization for conditional image generation. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 5094–5103 (2020)
35. Wu, Y.H., Gao, S.H., Mei, J., Xu, J., Fan, D.P., Zhang, R.G., Cheng, M.M.: JCS: An explainable covid-19 diagnosis system by joint classification and segmentation. *IEEE Transactions on Image Processing* **30**, 3113–3126 (2021)
36. Xiao, L., Xu, J., Zhao, D., Wang, Z., Wang, L., Nie, Y., Dai, B.: Self-supervised domain adaptation with consistency training. In: 2020 25th International Conference on Pattern Recognition (ICPR). pp. 6874–6880. IEEE (2021)
37. Xie, Q., Dai, Z., Hovy, E., Luong, T., Le, Q.: Unsupervised data augmentation for consistency training. In: Larochelle, H., Ranzato, M., Hadsell, R., Balcan, M.F., Lin, H. (eds.) *Advances in Neural Information Processing Systems*. vol. 33, pp. 6256–6268. Curran Associates, Inc. (2020)

38. Xu, J., Xiao, L., López, A.M.: Self-supervised domain adaptation for computer vision tasks. *IEEE Access* **7**, 156694–156706 (2019)
39. Yuan, H., Yu, H., Gui, S., Ji, S.: Explainability in graph neural networks: A taxonomic survey. *arXiv preprint arXiv:2012.15445* (2020)
40. Yun, W.L., Acharya, U.R., Venkatesh, Y.V., Chee, C., Min, L.C., Ng, E.Y.K.: Identification of different stages of diabetic retinopathy using retinal optical images. *Information sciences* **178**(1), 106–121 (2008)
41. Zhou, Y., He, X., Huang, L., Liu, L., Zhu, F., Cui, S., Shao, L.: Collaborative learning of semi-supervised segmentation and classification for medical images. In: *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*. pp. 2079–2088 (2019)
42. Zhou, Y., Wang, B., Huang, L., Cui, S., Shao, L.: A benchmark for studying diabetic retinopathy: Segmentation, grading, and transferability. *IEEE Transactions on Medical Imaging* (2020)