Synthesizing New Retinal Symptom Images by Multiple Generative Models

Yi-Chieh Liu¹, Hao-Hsiang Yang¹, C.-H. Huck Yang^{2,3}, Jia-Hong Huang^{1,2}, Meng Tian⁴, Hiromasa Morikawa², Yi-Chang James Tsai³, and Jesper Tegnèr^{2,5}

¹ National Taiwan University, Taiwan {b01310048,r05921014}@ntu.edu.tw
² Biological and Environmental Sciences and Engineering Division, Computer, Electrical and Mathematical Sciences and Engineering Division, King Abdullah University of Science and Technology, Thuwal, Saudi Arabia {chao-han.yang,

jiahong.huang, hiromasa.morikawa, jesper.tegner}@kaust.edu.sa

³ Georgia Institute of Technology, GA, USA

james.tsai@ce.gatech.edu

⁴ Department of Ophthalmology, Bern University Hospital, Bern, Switzerland tianmeng1231@gmail.com

⁵ Unit of Computational Medicine, Center for Molecular Medicine, Department of Medicine, Karolinska Institutet, Stockholm, Sweden

Abstract. Age-Related Macular Degeneration (AMD) is an asymptomatic retinal disease which may result in loss of vision. There is limited access to high-quality relevant retinal images and poor understanding of the features defining sub-classes of this disease. Motivated by recent advances in machine learning we specifically explore the potential of generative modeling, using Generative Adversarial Networks (GANs) and style transferring, to facilitate clinical diagnosis and disease understanding by feature extraction. We design an analytic pipeline which first generates synthetic retinal images from clinical images; a subsequent verification step is applied. In the synthesizing step we merge GANs (DCGANs and WGANs architectures) and style transferring for the image generation, whereas the verified step controls the accuracy of the generated images. We find that the generated images contain sufficient pathological details to facilitate ophthalmologists' task of disease classification and in discovery of disease relevant features. In particular, our system predicts the drusen and geographic atrophy sub-classes of AMD. Furthermore, the performance using CFP images for GANs outperforms the classification based on using only the original clinical dataset. Our results are evaluated using existing classifier of retinal diseases and class activated maps, supporting the predictive power of the synthetic images and their utility for feature extraction.

1 Introduction

As a rule, it is challenging to automatically diagnose retinal diseases from images, partly because of the difficulty of acquiring public data with a sufficient number of annotated images due to concerns of personal privacy. Meanwhile, different ophthalmologists may provide conflicting judgments about identical images; therefore, it can be arduous to reach consensus about a diagnosis. Thus, it is clear that a larger number of retinal images collected from a system with provided unbiased feature detection would be beneficial for ophthalmologists clinical practice.

Generative models, such as Generative Adversarial Networks [1] (GANs), and style transferring [2] techniques, have achieved impressive results for generating sharp and realistic images. Therefore, these two methods are used to synthesize the disease images from healthy retinal images and diseased ones. Synthesized images not only impose high-level symptom features to the original ones but help ophthalmologists build the understanding of related diseases. The definition of image synthesis in [3] is seen as an image reconstruction process coupled with feature transformation. The synthesized part is responsible for inverting features back to the color space and the feature transformation matches certain statistics of a original image to a generated image [4].

We consider images with Age-Related Macular Degeneration (AMD) as an asymptomatic retinal disease and the leading cause of irreversible visual loss among the aged population. Despite the advances of therapeutics, there is still no satisfactory treatment. It raises the issue that diagnosing AMD from its early stage and having proper managing it properly are more important than ever. The development of AMD is classified as several stages that can be discerned by two explicit symptoms, drusen and Geographic Atrophy(GA). Drusen are one of the earliest clinical indications of AMD, which appears as focal, with yellow excrescences deep in the retina with extra-cellular deposits located beneath the retinal pigment, epithelium, and Bruchs membrane; the number, size and distribution of these deposits is highly variable. GA, symptomatic of a more advanced stage of AMD, is described as a well-demarcated area of decreased retinal thickness. Such areas have relative changes in color compared to surroundings allowing an increased visualization of the underlying choroidal vessels. The phenomenon is that less intense and more diffuse hyperfluorescence in which pigment clumping sometimes forms a microreticular pattern, is demonstrated [5] [6] [7]. To sum up, two symptoms (drusen and GA) are established clinical hallmarks of AMD. Drusen size and confluency have been historically associated with the progression of AMD, which also contributes to the development of GA. Our chief objective is to generate images equipped with a sufficient number of pathological features to capture the two different stages of AMD.

The contribution of this paper consists of two parts. First, style transferring, WGANs and DCGANs are used to build a new artificial neural network as the framework for the generation of synthetic pathologically relevant but detailed images. Second, after new images are obtained and diagnosed by ophthalmologists, we use Class Activation Maps (CAMs) [8] to locate the advanced features within the generated images. Finally, the EyeNet [9] is used to classify generated images according to the established labelling of diseases.

The paper is organized into five sections. Following this introduction, in the second section, we survey related work. In the third section, we present our analytic pipeline, including an account of how we fuse DCGANs, WGANs, style transferring, EyeNet, and CAMs. In the fourth section, we present and discuss computational experimental results. In the last section, we summarize conclusions and outline prospects the future work.

2 Related Work

Below, we survey previous work on GANs, from which we benefit, synthetic image generation, and computational retinal disease methods.

2.1 GANs

Since the pioneering formulation of GANs [1], there have been numerous studies of how to formulate the optimization problem of balancing on the one hand the training of a generative network \mathbf{G} producing realistic synthetic samples, and on the other, a discriminator network \mathbf{D} that distinguishes between real and synthetic (generated) data. We adopt an adversarial loss

$$min_G max_D L_{GANS} = E_{x \sim p_{data(x)}} [log D(x)] + E_{x \sim p_{prior(z)}} [log(1 - D(z))] \quad (1)$$

Yet, a major issue has been the stability and convergence of training a GAN. Recent work [10] demonstrated improved stability when using a Kantorovich-Rubinstein metric, which we have adopted in our training of the GAN for retinal images. Rapid advances have demonstrated that GANs generate realistic images, with a rich number of features. For example, GANs have been successfully applied for face generation [11], indoor scene reconstruction [12] and person reidentification [13]. Here we benefit from recent progress with GANs to generate new synthetic retinal disease images using both Deep Convolutional Generative Adversarial Networks (DCGANs) [14] and Wasserstein GANs (WGANs) [10]. These architectures utilize a convolutional decoder, and DCGAN enables the employment of large GANs using Convolutional Neural Networks (CNNs), resulting in stable training across various datasets. Finally, our use of WGANs improves the stability of learning, thereby avoiding known challenges such as model collapse [10].

2.2 Generation of Synthetic Images

Recently, researchers have used convolutional neural networks to generate images [15] with different given style. The method makes use of a pre-trained network to optimize the image and its features. However, this method operates as a global optimization; therefore, generated image exhibit distortions and detailed parts cannot be presented on the transferred images. Meeting this challenge, recent work [2] accomplished realistic image generation and style transferring. On the other hand, in [14], authors combined a convolutional neural network and GANs to generate new images, thus mitigating the impact of a limited number of features of pathological relevance in original images. While their work clearly improved the state-of-the-art methods, the technique may generate poor image 4 Y.-C. Liu et al.

samples or fails to converge. To ensure convergence and quality of generated images, we deploy a closed form solution for style transferring [16].

2.3 Computational Retinal Disease Methods

There is a challenge to build large high-quality medical databases despite massive investment in, for example, data collection, labeling, and data augmentation. Exceptions include the recently released ChestXray14 [17] dataset which contains 112,120 frontal-view chest radiographs with up to 14 thoracic pathological labels. Yet, in contrast, for retinal research, the DRIVE dataset [18], which contains only 40 retina images, has long been a standard. However, recently the Retina Image Bank (RIB) [19], containing a large number of different kinds of retinal images, is truly an enabler for the kind of work presented in our paper. Despite this, we still need techniques to augment such databases due to various challenges, such as the limited amount of annotation, thus effectively transforming a small dataset with low diversity into one that approximates the underlying data distribution. For example, in [20] and [21], GANs were used to generate a variety of retinal images and targeting control (healthy) images. Using the Retinal Image Bank, we aim to generate new retinal images that have a sufficient number of pathological details so that we have abundant and useful retinal images to train and build robust classifier. In [22], authors propose a method that implements automated segmentation of retina to facilitate the detection of disease. Article [9] uses the whole retinal images in [19] to train the classifier, which can discern multiple diseases with the extraction of visual traits. Our work depends on having a pre-trained network to test the quality of generated images and uses CAMs to present symptoms identified by the classifier.

3 Methodology

In this section, we describe different methods in our proposed pipeline. For generative models, GANs and style transferring based networks are discussed. For verification, we elaborate EyeNet and CAMs.

3.1 Style Transferring

The input contains two images: a content and a style image and pre-trained CNNs; the output is the synthesized image. In our case, the content image means the disease image with pathological details; the style image represents the healthy retinal image. When it comes to the existing style transferring methods, even though the style is changed, the content of the image can be seen in the new image. Thus, we expect generated images with pathological details, so content images are seen as disease images. For each image, the output from the CNNs classifier obtains various level features from many convolutional layers. Generated images preserve the original semantic content from the content image but look like a style image. For the content and the style part, loss functions that are computed from the similarity of images from convolutional layers can be defined; style transferring becomes an optimization problem when the optimal image is obtained with the least loss. The pixels in the image can be computed iteratively by gradient descent.

3.2 DCGANs and WGANs for image generation

Although original GANs provide an intriguing algorithm with surprising results, the instability is what we concern about when it comes to medical applications, which requires precision and detailed images for diagnosis. To improve the quality of generated images, we chose DCGANs and WGANs to establish our generative model. In this part, with a random initialized parameter, we build a generator of retinal diseases while the improving discriminator. For a specific symptom, generated images contain similar optical traits. Furthermore, high dimensional neural networks for computer vision sometimes materialize higher forms of neglected visual features. Therefore, generated retinal images not only become the aid of diagnosis and strategy to explore diseases, but also provide diverse computer training data.

3.3 Class Activation Maps

The class activation maps (CAMs) in [8] provide a method that localizes features on images. From localized features, the performance of the generated image can be evaluated and observed. As discussed in Section 3.1, convolutional layers of CNNs are used to extract the visual feature of images. Through this method, not only the similarity of images is tested with high-level disease features, but a series of pathological details is built.

3.4 EyeNet

Besides CAMs, EyeNet as proposed in [9] is used to evaluate the correctness of generated images. In [9], the authors trained a network that classifies different retinal diseases; 52 kinds of retinal diseases are labeled and classified. Proposed methods [23] include three frameworks: U-net, SVM and ResNet50; predicton by ResNet50 performs best. Therefore, ResNet50 is modified so that the generated images also can be classified to make sure of their correctness.

3.5 Pipeline

We propose a pipeline structure in Fig. 1. Initially, with feature extraction by style transferring and GANs, more images are generated. In order to verify the correctness, CAMs and EyeNet are used to compute the high level visual features and predict the diseases, respectively. Results from the CAMs present pathological details. Moreover, generated images can be applied to feed to other classifier to train the more accurate classifier. All researches benefits not only the newly trained network, but also ophthalmologists. Original retinal images give doctors an initial diagnosis, and the generated images provide them more clues. CAMs help ophthalmologists judge accurately, and they can reach the consensus with EyeNet. To sum up, our pipeline improves the efficiency and accuracy of the medical system and contributes to researchers.

4 Experiments

In this section, we describe the implementation details and experiments we conduct to validate our proposed methods. Initially, the data collection and setup of experiments are emphasized. And generating images by style transferring, GANs are presented. Finally, generated images are diagnosed by doctors and EyeNet



Fig. 1. Proposed pipeline to improve retinal diagnosis efficiency by combining GANs and EyeNet to assist ophthalmologists and doctors. EyeNet is used to check the performance. Finally, observation of similarities among some diseases are analyzed and described.

is used to check the performance. Furthermore, observation of similarity among some diseases are analyzed and described.

4.1 Dataset Collection

Experimental images come from the Retina Image Bank (RIB) [19]. Retinal image collection contains three types of photography that are fluorescein angiography(FA), optical coherence tomography (OCT) and color fundus photography (CFP). FA are gray-scale images and CFP are colorful images. CFP and FA imaging are reliable for whole fundus, and used as our dataset.

4.2 Setup

As discussed above, we use images with AMD for experiments. Images contain CFP and FA type, and present the symptom of drusen and GA. All DNNs were implemented in PyTorch, and we modified the publicly available PyTorch code for the neural network algorithm. Details of various methods are described later, respectively. The derivative of all generative models is sped in CUDA for gradient-based optimization.

4.3 Style Transferring Neural Networks

Style transferring neural network in [16] was modified to generate new disease images. This network adopts layers from "conv1_1" to "conv4_1" in pre-trained VGG-19 [24] network for the encoder, whose weights are provided by ImageNetpretrained weights. What's more, multi-level stylization strategy proposed in [16] is applied to optimize the VGG features in different layers. Input images are three CFP images and three FA images as style images shown in Fig. 2 and 5. Six CFP images with three drusen and three GA images in Fig. 3 and Fig. 4. Also, FA images are applied to generate new images in Fig. 6 and Fig. 7. For CFP images, six images are shown in Fig. 3 and in Fig. 4. In Fig. 3, generated images contain round, discrete yellow-white dots, which are the symptom of drusen. In the same way, in Fig. 4, well-demarcated areas appear on the three images. Therefore, style transferring can generate new retinal symptom images.

Furthermore, generated images from FA images are presented in Fig. 6 and Fig. 7. Results in the images are nearly identical to the original images, because



Fig. 2. Three CFP fundus images that are used to generate new images are seen as style images.



Fig. 3. Three CFP fundus images with symptom of drusen and corresponding generated images. (a), (b), (c) Original images. (d), (e), (f) Generated images.

original networks are applied to stylize color images. However, six generated images contain more concise features than the original ones, which helps ophthalmologists make better judgements. Therefore, this style transferring networks can fulfill edge sharpening and enhancement of contrast. No matter which kinds of images are generated, advanced features in new disease images still exist. Fur-





Fig. 4. Three CFP fundus images with symptom of GA and corresponding generated images. (a), (b), (c) Original images. (d), (e), (f) Generated images.



Fig. 5. Three AF fundus images that are used to generate new images are seen as style images.

thermore, analyses of image performance by EyeNet and CAMs for prediction are presented in a later section.

4.4 DCGANs and WGANs

In this section, DCGANs and WGANs are trained with thousands of CFP and FA images that have symptoms of drusen and GA separately; both of the models require four to six hours to train. Generated images have been diagnosed



Fig. 6. Three FA fundus images with symptom of drusen and corresponding generated images. (a), (b), (c) Original images. (d), (e), (f) Generated images.

by ophthalmologists for verification. Images generated by DCGANs, which are shown in Fig. 8, cannot be identified as a valid retinal image with symptom. However, drusen and GA images generated by WGANs can be used by ophthalmologists to diagnose. In Fig. 4, generated drusen images are diagnosed as insignificant of drusen but can be identified by EyeNet. As for generated GA images in Fig. 4, irregularly shaped macular atrophy can be identified by an ophthalmologist. Macular atrophy is a distinguishable trait of GA, which means WGANs indeed learn the symptoms of drusen and GA from specific AMD and generate new images. Thus, WGANs perform better than DCGANs because of resolution. Structure of DCGANs limits the size of generated images to be 64x64, so some pathological details are lost. We choose WGANs for following experiments.

4.5 EyeNet Results for prediction

EyeNet in [9] is trained to predict the accuracy of images to accomplish the pipeline. Though ImageNet and our retinal dataset are much different, using pre-trained weights on ImageNet rather than random ones has boosted testing accuracy of any models with 5 to 15 percent. Besides, pre-trained models tend to converge much faster than random initialized. The training images encompass 52 kinds of fundus images, which are randomly divided into three parts: 70%



Fig. 7. Three FA fundus images with symptom of GA and corresponding generated images. (a), (b), (c) Original images. (d), (e), (f) Generated images.

for training, 10% for validation and 20% for testing. It is noted that synthesized images are not used to train this network. The training lasts 400 epochs. The first 200 epochs take a learning rate of 1e-4 and the second 200 take 1e-5. Besides, we apply random data augmentation during training. In every epoch, 70% probability for a training sample is affinely transformed. After EyeNet is trained, generated images are fed into it, and the average predicted probabilities are shown in Table 1. Compared to drusen, The accuracy declines when it comes to identifying generated geographic atrophy images. The lack of geographic atrophy images in the EyeNet dataset weaken the capability of the classifier to discern traits about geographic atrophy. Despite of the setback, there is a exciting exploration that the predictions are not randomly distributed but focus on particular diseases, which is likely caused by the high-dimensional features mentioned above.

4.6 Image Sample Size Effect

Generative models are data-driven and the performance highly depending on the sample size. The EyeNet dataset we use contains 19496 retinal images with 1448 AMD images. In this section, We choose 338 drusen images as samples to test the size effect of GANs. Experiments show the difficulty of synthesizing high quality images rises along with the increase of the sample number. Fig. 10



Fig. 8. Drusen images generated by DCGANs.(a) CFP image. (b) FA image.



Fig. 9. Drusen and GA images generated by WGANs.(a) Generated drusen CFP image. (b) Generated drusen FA image. (c) Generated drusen CFP image. (d) Generated GA FA image.

shows accuracy of successfully predicting synthesized images, but AMD slightly declines as the sample number increases. In general, the more samples used to train a generative, the harder it is to extract specific visual features for generative model, which requires images with similarity; this is difficult to achieve when it comes to biological traits. On the other hand, prediction error focuses on some

12 Y.-C. Liu et al.

	Drusen-CFP	Drusen-FA	GA-CFP	GA-FA
Real Images	0.442	0.368	0.166	0.128
WGANs	0.594	0.374	0.254	0.291
Style Transferring	0.376	0.657	0.139	0.231

 Table 1. Average accuracy classified by EyeNet.



Fig. 10. Size effect of image samples and probability of predicting other specific diseases.

specific diseases, and the probability of predicting these diseases rises when the sample number falls. The phenomenon implies that high dimensional features in the retinal images exist. Furthermore, with more sample images, we can more likely to detect the symptom. This is a pathological approach to reveal hidden relations among diseases.

4.7 Pathological retinal diseases classification inspired by size effect of GANs With higher quality images and thriving computer vision skills, visible retinal disease symptoms can being detected and represented. Based on traditional classification, symptoms have pathological correlations among retinal diseases for ophthalmologist to use in diagnoses, as shown in Fig. 11. However, according to the discovery above, retinal diseases have hidden relation connected by invisible features. With GANs, we can propose a method to improve current classification. In this case, the classification could modified by the results of GANs as shown in Fig. 12.

4.8 Neural Network Visualization for Retinal Images

Finally, we verified the hypothesis that vessel-based segmentation and contrast enhancement are two coherent features to decide the type of retinal diseases. Using techniques of generating CAMs introduced in [8], we visualized feature maps of the final convolutional layer of ResNet50 in Fig. 13. In our results, generated drusen images are well identified. However, generated GA images are not focused on the exact location of the symptom, but they are close. As discussed above, in the clinical diagnosis process, "vessel patterns" and "fundus structure" are the most crucial features for identifying the symptoms of different diseases. These types of features cover more than 80% of retinal diseases [27, 28].



Fig. 11. Clinical hierarchical relationship of AMD and Choroidal metastases from [25] and [26]. Drusen increase a persons risk of developing AMD; Geographic atrophy (GA) is an advanced form of age-related macular degeneration that can result in the progressive and irreversible loss of retina. Subretinal fluid and equatorial retina are significant features in choroidal metastases from clinical understanding.



Fig. 12. Pathological Classification referring to GANs results. As a visual representation learned by deep network, Drusen are correlated to AMD (58.3%), Choroidal metastasis (5.12%), Presumed-ocular-histoplasmosis-syndrome (6.52%), and Pattern dystrophy (2.57%).

5 Conclusions and Future Work

We have implemented style transferring, DCGANs and WGANs to generate disease images that are detailed to capture different stages of AMD. Symptoms of images are drusen and GA; both FA and CFP images are generated. Images from DCGANs are difficult to be identified due to limit of resolution. However, images from style transferring and WGANs are easier to identify by ophthalmologists, and generated images preserve pathological details. EyeNet is used to predict the disease label, and results of generated drusen images are similar to original images. However, generated GA images are more distant compared to original images, because of the small number of GA images used during training EyeNet. This phenomenon shows that generated new images can be fed into the classifier to improve it. Also, CAMs are useful for extracting label-specific features. In Fig. 13(c),(f) and(i), warmer color parts are located in the welldemarcated areas or spots, which represents disease features that are close to those parts.

In this paper, only a small number of disease images are synthesized and evaluated, so various images can be tested and enhanced further. Furthermore, different kinds of skills like semantic segmentation can be merged into the original



Fig. 13. Generated images and their CAMs. (a) Original drusen image (CFP). (b) Generated image of (a) by WGANs. (c) CAM of (b). (d) Original GA image (CFP). (e) Generated image of (d) by WGANs. (f) CAM of (e). (g) Original GA image (FA). (h) Generated image of (g) by style transferring. (i) CAM of (g).

GANs framework. With better and diverse generated images, classifier can be trained robustly and applied to predict the disease more precisely. Above all, a re-trained network discovers hidden relationships and provides ophthalmologists with useful disease features warranting further investigation.

15

References

- Goodfellow, I., Pouget-Abadie, J., Mirza, M., Xu, B., Warde-Farley, D., Ozair, S., Courville, A., Bengio, Y.: Generative adversarial nets. In: Advances in neural information processing systems. (2014) 2672–2680
- Luan, F., Paris, S., Shechtman, E., Bala, K.: Deep photo style transfer. CoRR, abs/1703.07511 2 (2017)
- Li, Y., Fang, C., Yang, J., Wang, Z., Lu, X., Yang, M.H.: Universal style transfer via feature transforms. In Guyon, I., Luxburg, U.V., Bengio, S., Wallach, H., Fergus, R., Vishwanathan, S., Garnett, R., eds.: Advances in Neural Information Processing Systems 30. Curran Associates, Inc. (2017) 386–396
- Salehinejad, H., Valaee, S., Dowdell, T., Colak, E., Barfett, J.: Generalization of deep neural networks for chest pathology classification in x-rays using generative adversarial networks. In: 2018 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP), IEEE (2018) 990–994
- Klein, R., Klein, B.E., Knudtson, M.D., Meuer, S.M., Swift, M., Gangnon, R.E.: Fifteen-year cumulative incidence of age-related macular degeneration: the beaver dam eye study. Ophthalmology 114 (2007) 253–262
- Green, W.R., McDONNELL, P.J., Yeo, J.H.: Pathologic features of senile macular degeneration. Ophthalmology 92 (1985) 615–627
- Gheorghe, A., Mahdi, L., Musat, O.: Age-related macular degeneration. Romanian journal of ophthalmology 59 (2015) 74–77
- Zhou, B., Khosla, A., A., L., Oliva, A., Torralba, A.: Learning Deep Features for Discriminative Localization. CVPR (2016)
- Yang, C.H.H., Huang, J.H., Liu, F., Chiu, F.Y., Gao, M., Lyu, W., Tegner, J., et al.: A novel hybrid machine learning model for auto-classification of retinal diseases. arXiv preprint arXiv:1806.06423 (2018)
- Arjovsky, M., Chintala, S., Bottou, L.: Wasserstein generative adversarial networks. In: International Conference on Machine Learning. (2017) 214–223
- Pumarola, A., Agudo, A., Martinez, A.M., Sanfeliu, A., Moreno-Noguer, F.: Ganimation: Anatomically-aware facial animation from a single image. In: Proceedings of the European Conference on Computer Vision (ECCV). (2018) 818–833
- 12. Fan, H., Su, H., Guibas, L.J.: A point set generation network for 3d object reconstruction from a single image. In: CVPR. Volume 2. (2017) 6
- Qian, X., Fu, Y., Wang, W., Xiang, T., Wu, Y., Jiang, Y.G., Xue, X.: Pose-normalized image generation for person re-identification. arXiv preprint arXiv:1712.02225 (2017)
- Radford, A., Metz, L., Chintala, S.: Unsupervised representation learning with deep convolutional generative adversarial networks. arXiv preprint arXiv:1511.06434 (2015)
- Gatys, L.A., Ecker, A.S., Bethge, M.: Image style transfer using convolutional neural networks. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition. (2016) 2414–2423
- Li, Y., Liu, M.Y., Li, X., Yang, M.H., Kautz, J.: A closed-form solution to photorealistic image stylization. arXiv preprint arXiv:1802.06474 (2018)
- Wang, X., Peng, Y., Lu, L., Lu, Z., Bagheri, M., Summers, R.M.: Chestx-ray8: Hospital-scale chest x-ray database and benchmarks on weakly-supervised classification and localization of common thorax diseases. In: 2017 IEEE Conference on Computer Vision and Pattern Recognition (CVPR), IEEE (2017) 3462–3471

- 16 Y.-C. Liu et al.
- Staal, J., Abràmoff, M.D., Niemeijer, M., Viergever, M.A., Van Ginneken, B.: Ridge-based vessel segmentation in color images of the retina. TMI 23 (2004) 501–509
- Retina Image Bank: A project from the American Society of Retina Specialists. (http://imagebank.asrs.org/about) Accessed: 2018-06-30.
- Beers, A., Brown, J., Chang, K., Campbell, J.P., Ostmo, S., Chiang, M.F., Kalpathy-Cramer, J.: High-resolution medical image synthesis using progressively grown generative adversarial networks. arXiv preprint arXiv:1805.03144 (2018)
- Guibas, J.T., Virdi, T.S., Li, P.S.: Synthetic medical images from dual generative adversarial networks. arXiv preprint arXiv:1709.01872 (2017)
- Kaur, J., Mittal, D.: Segmentation and measurement of exudates in fundus images of the retina for detection of retinal disease. Journal of Biomedical Engineering and Medical Imaging 2 (2015) 27
- Yang, C.H.H., Liu, F., Huang, J.H., Tian, M., Liu, Y.C., Lin, I., Tegner, J., et al.: Auto-classification of retinal diseases in the limit of sparse data using a two-streams machine learning model. arXiv preprint arXiv:1808.05754 (2018)
- 24. Simonyan, K., Zisserman, A.: Very deep convolutional networks for large-scale image recognition. arXiv:1409.1556 (2014)
- Coleman, H.R., Chan, C.C., Ferris, F.L., Chew, E.Y.: Age-related macular degeneration. The Lancet 372 (2008) 1835 – 1845
- Mewis, L., Young, S.E.: Breast carcinoma metastatic to the choroid: Analysis of g7 patients. Ophthalmology 89 (1982) 147–151
- 27. Crick, R.P., Khaw, P.T.: A textbook of clinical ophthalmology: a practical guide to disorders of the eyes and their management. World Scientific (1998)
- Akram, I., Rubinstein, A.: Common retinal signs. an overview. Optometry Today (2005)