OCTA-based AMD Stage Grading Enhancement via Class-Conditioned Style Transfer

Haochen Zhang, Anna Heinke, Krzysztof Broniarek, Carlo Galang, Daniel Deussen, Katarzyna Michalska-Małecka, Dirk-Uwe Bartsch, William Freeman, Truong Nguyen, Cheolhong An

the 46th Annual International Conference of the IEEE Engineering in Medicine and Biology Society | July 16, 2024

Contents

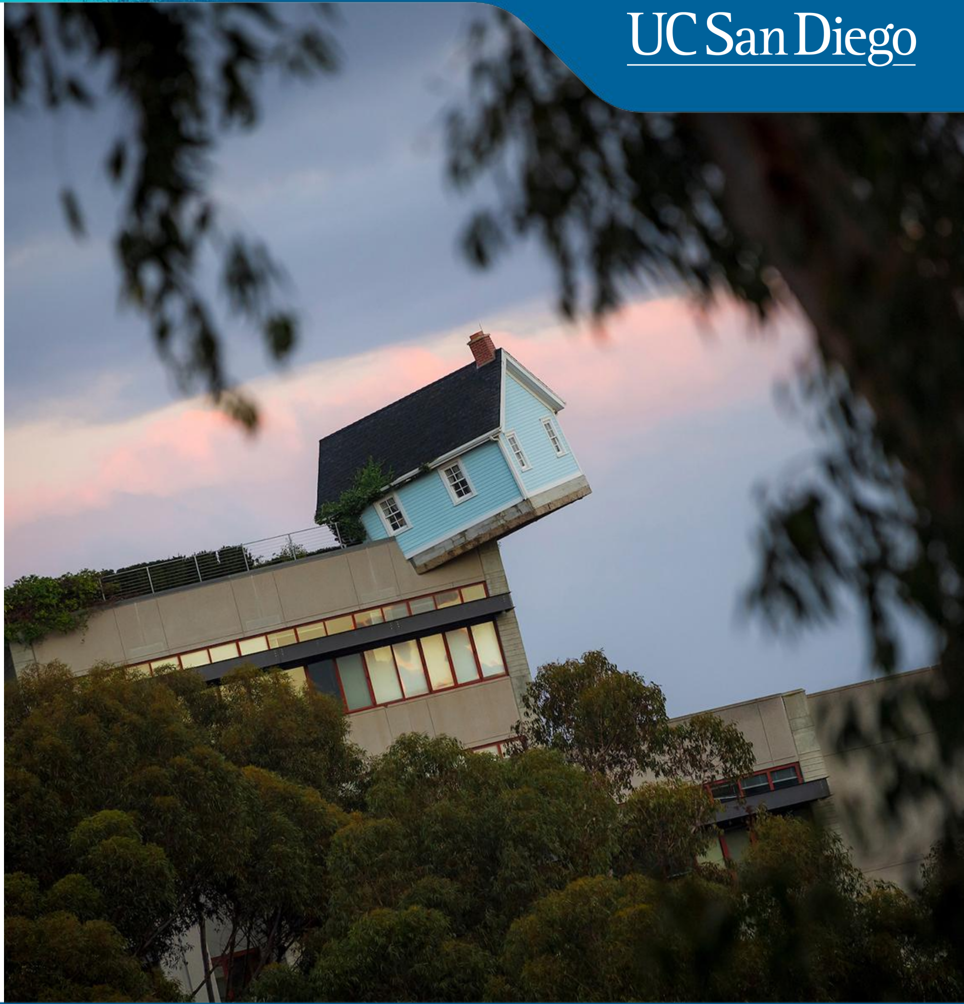

UC San Diego

Background

Age-related Macular Degeneration (AMD)

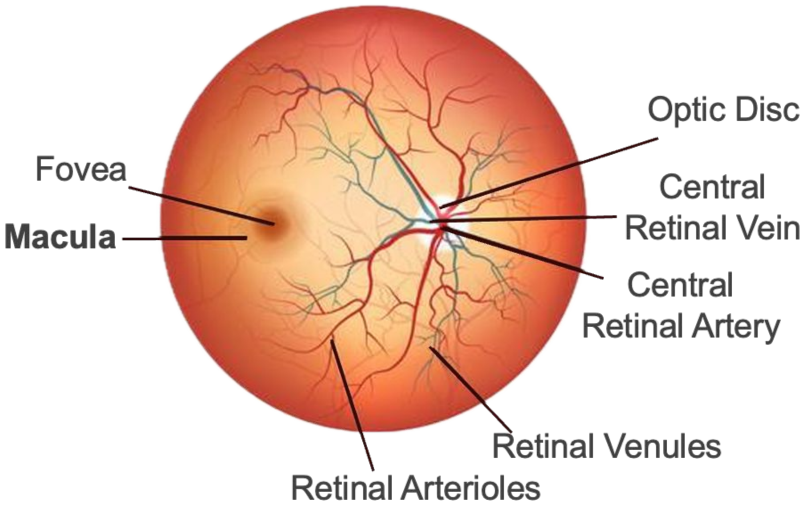

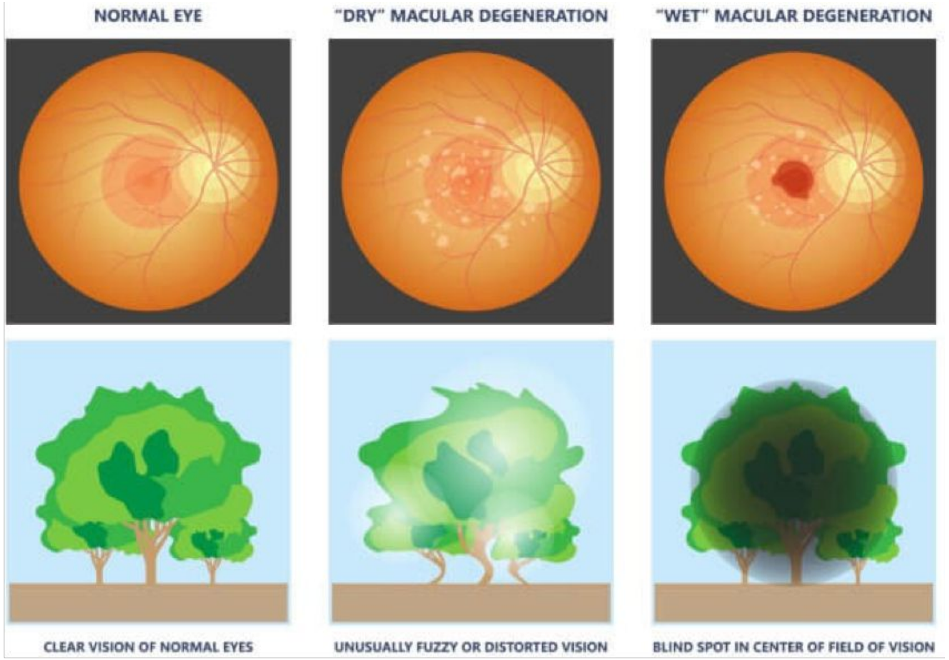

Normal Dry AMD Wet AMD

https://www.centralfloridaretina.com/retinal-vascular-disease-orlando/
https://www.allaboutvision.com/conditions/amd.htm

Background

Retina Imaging

Fundus

- Invasive (inject dye)
- Show vessels as a whole
- Diagnose wet AMD based on a macular hemorrhage (late symptom)

Optical Coherence Tomography (OCT)

- Non-invasive (no dye)
- Cross sectional view of the retina (layers)

OCT Angiography (OCTA)

- Non-invasive
- The microvasculature of the retina
- Visualize vessels in different retinal layers
    - Enable an earlier detection,
    - Monitor the clinical response to treatment.

Popular, Gold standard

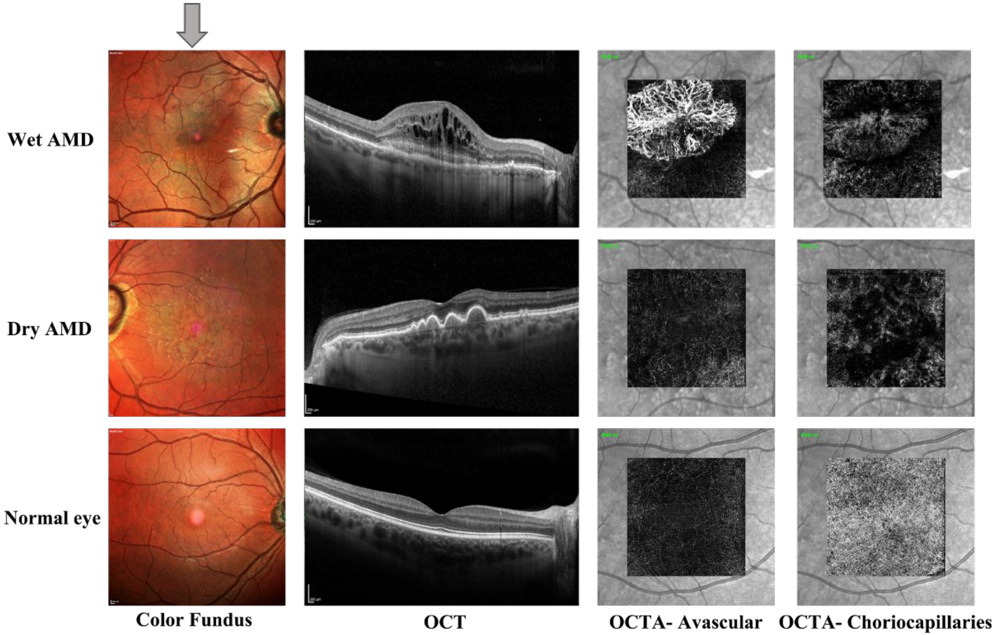

Background

Target Categories

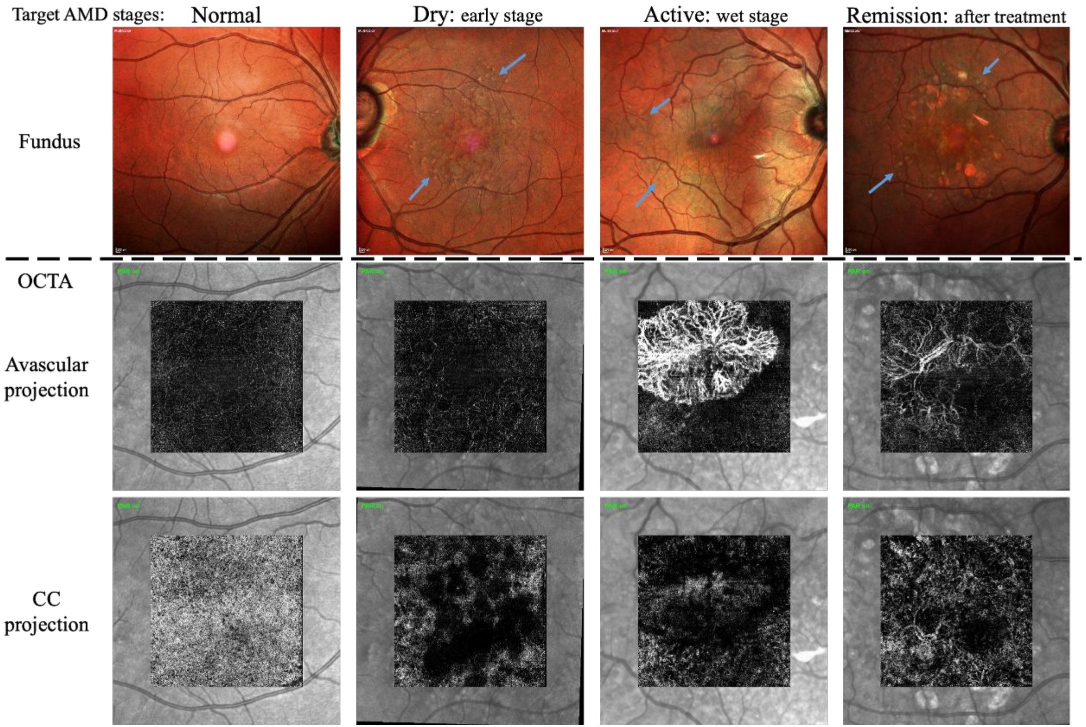

- All AMD stages show drusens
- Hard to tell each stage based on the pattern of drusens

- Dry vs Normal ✓
- Dry vs Active/Remission ✓
- Active/Remission ✕

Motivation

UC San Diego

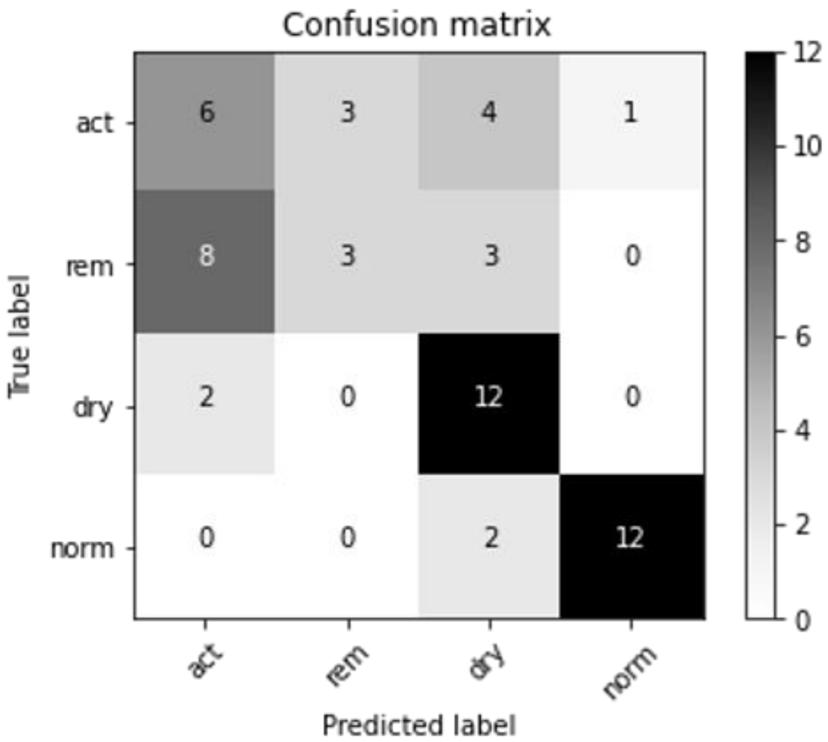

Human Expert performance

- There is no well established rules to distinguish each category
- Accuracy of ophthalmologists is low, 58.92%
- It is more efficient for computer to handle 3D data or multiple projections than human
- Could deep learning achieve a better accuracy?

UC San Diego

Motivation - Existing Datasets

OCTA dataset	Task	2D/3D	# of samples	# of AMD
PREVENT[3]	Segmentation	2D only	55	Unknown
ROSE [5]	Segmentation	2D only	229	Unknown
OCTAGON [3]	Segmentation	2D only	213	0
FOCTAIR [3]	Registration	2D only	86	Unknown
OCTA-500 [4]	Segmentation	3D and 2D	500	49
Ours	Classification	3D and 2D	889	749

Discussion

- No public OCTA dataset available for AMD stage grading
- OCTA is an emerging modality —-> Collecting cross-instrument data (Heidelberg and Optovue)

3 https://datashare.ed.ac.uk/handle/10283/3528

Style Transfer

CycleGAN

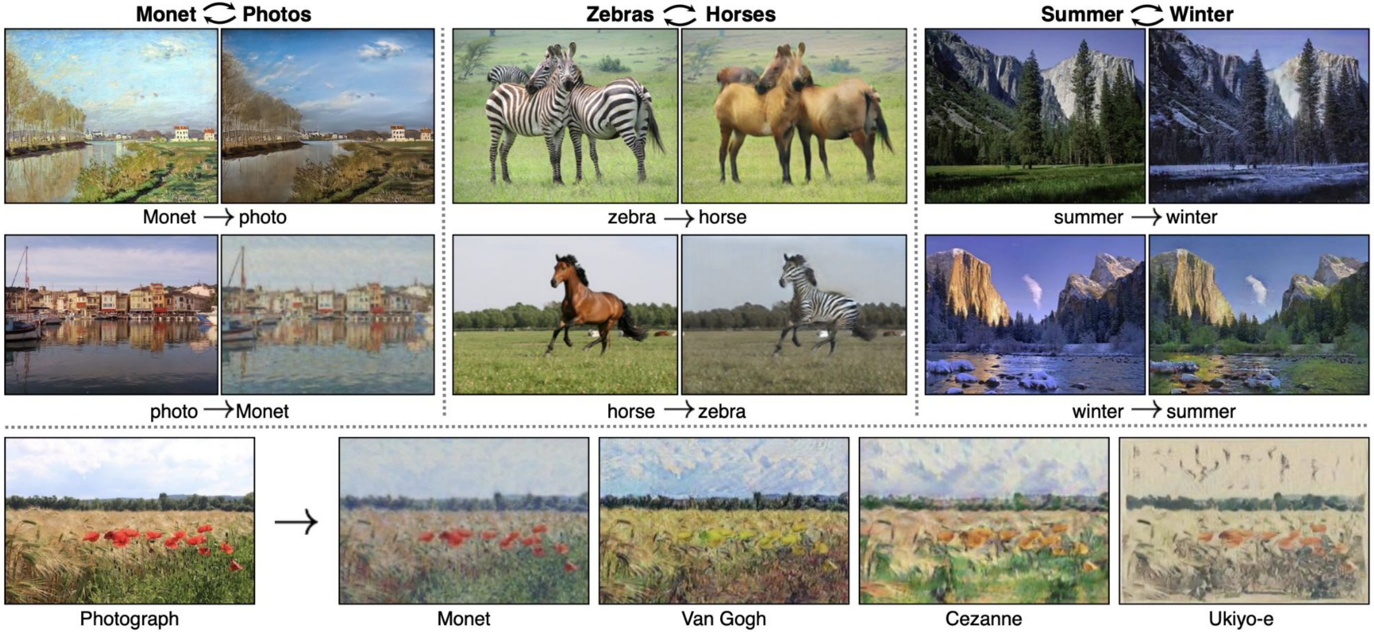

Figure 1: Given any two unordered image collections X and Y, our algorithm learns to automatically "translate" an image from one into the other and vice versa: *(left)* Monet paintings and landscape photos from Flickr; *(center)* zebras and horses from ImageNet; *(right)* summer and winter Yosemite photos from Flickr. Example application *(bottom)*: using a collection of paintings of famous artists, our method learns to render natural photographs into the respective styles.

Zhu, J.Y., Park, T., Isola, P. and Efros, A.A., 2017. Unpaired image-to-image translation using cycle-consistent adversarial networks. In Proceedings of the IEEE international conference on computer vision (pp.2223-2232).

Contribution

UC San Diego

- **CycleGAN for Unpaired Style Transfer**: Utilizing a CycleGAN translator to convert samples between Heidelberg and Optovue domains, introducing additional content and enhancing sample diversity.

- **Explicit Class Constraint**: Additional constraints are introduced to optimize machine recognition quality. Our proposed class constraint ensures that transferred images belong to the same class in fully-supervised manner or maintains similarity in unsupervised manner.

- Experimental results demonstrate that the CycleGAN translator effectively enhances classification performance in each domain, and our proposed class-conditioned approach further boosts the accuracy.

Contents

CycleGAN Framework

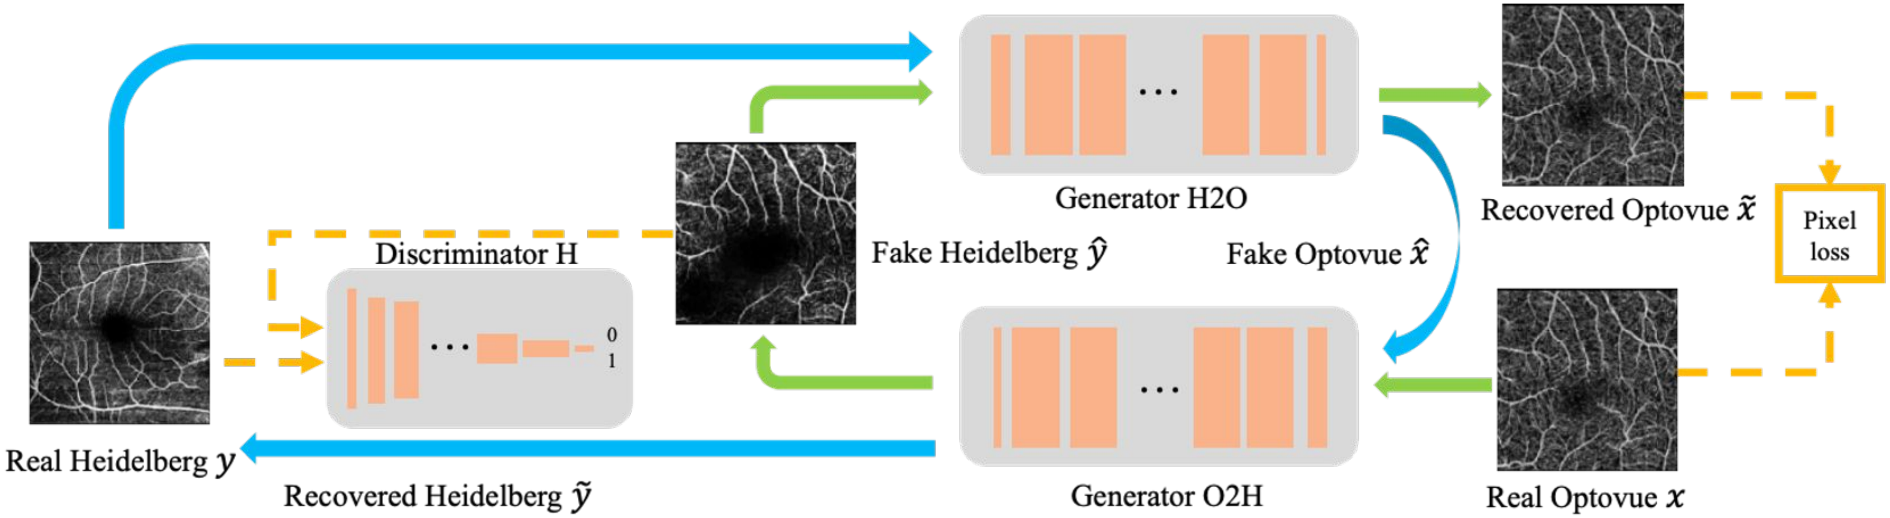

- No need of paired training data
- Only one channel of OCTA projection is displayed
- Identity loss is omitted

CycleGAN loss

$$\begin{aligned} l_{cyclegan}(x,y) = \; & l_{cyc}(x,\tilde{x}) + l_{cyc}(y,\tilde{y}) \\ & + \alpha\big(l_{GAN}(x,\hat{x}) + l_{GAN}(y,\hat{y})\big) \\ & + \beta\left(l_{idt}\big(x, G_{H2O}(x)\big) + l_{idt}\big(y, G_{O2H}(y)\big)\right) \end{aligned}$$

Explicit Class Condition

Supervised manner

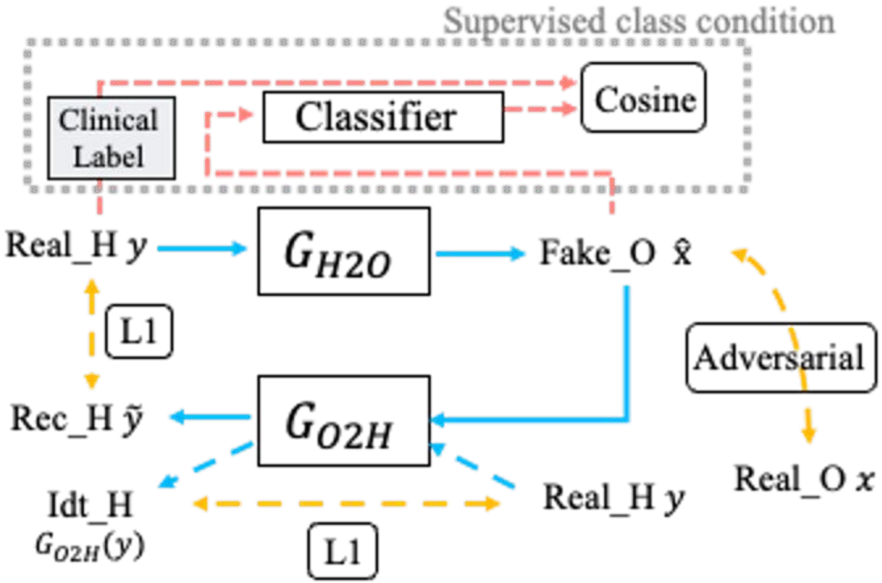

Training loss

$$l_{total} = l_{cyclegan}(x, y) + \gamma\left(l_{cos}(label(x), cls(\hat{y})) + l_{cos}(label(y), cls(\hat{x}))\right)$$

- Performing style transfer
- Preserve category information
- Utilize Cosine loss instead of cross entropy loss

Explicit Class Condition

Unsupervised manner

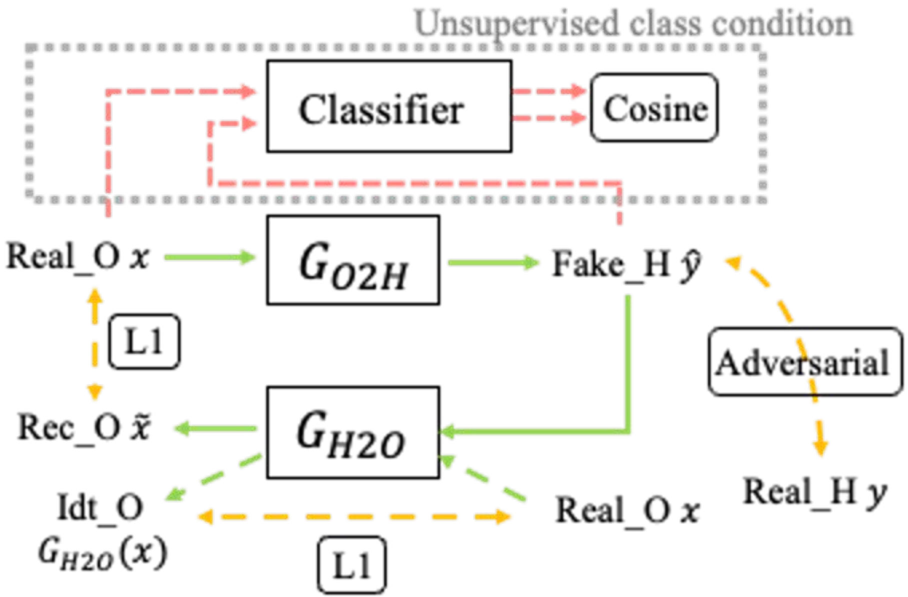

Training loss

$$l_{total} = l_{cyclegan}(x, y) + \gamma\left(l_{cos}(cls(x), cls(\hat{y})) + l_{cos}(cls(y), cls(\hat{x}))\right)$$

- Classifier generates pseudo labels
- ~Soft labels utilized in teacher-student distillation learning
- Utilize Cosine loss instead of cross entropy loss

Classifier training with CycleGAN

UC San Diego

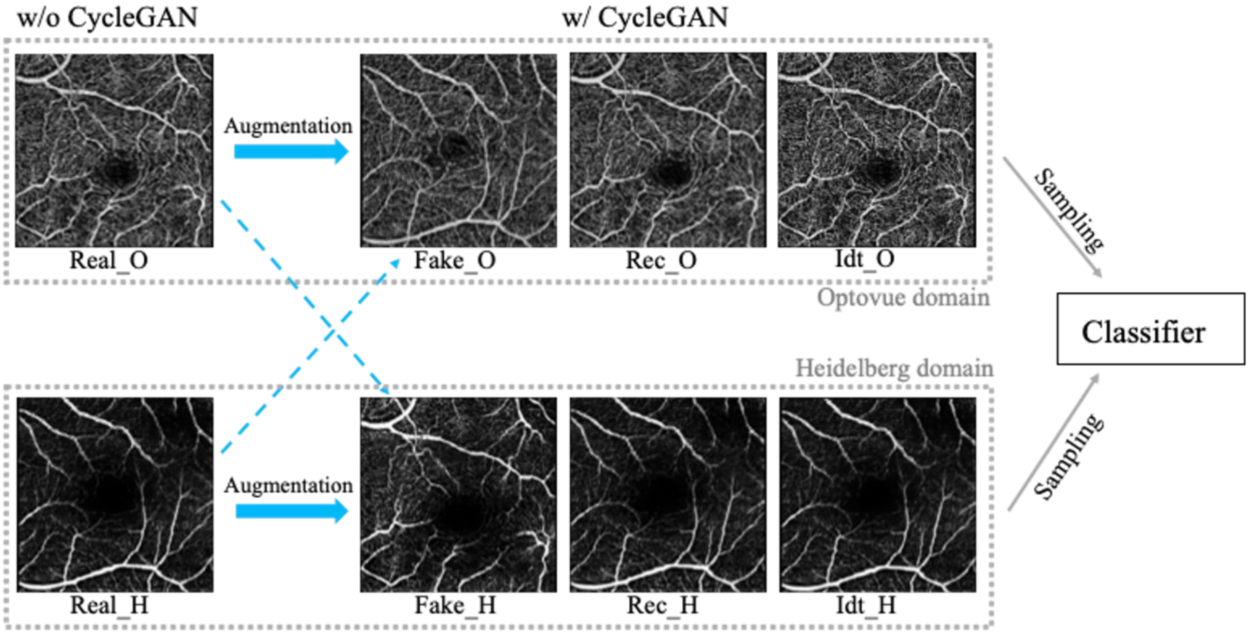

Previously unseen content in each domain(‘Fake O’ and ‘Fake H’)

Intra-domain sample augmentation (‘Rec O’ and ‘Idt O’)

Predetermined sampling strategy

Classifier Structure

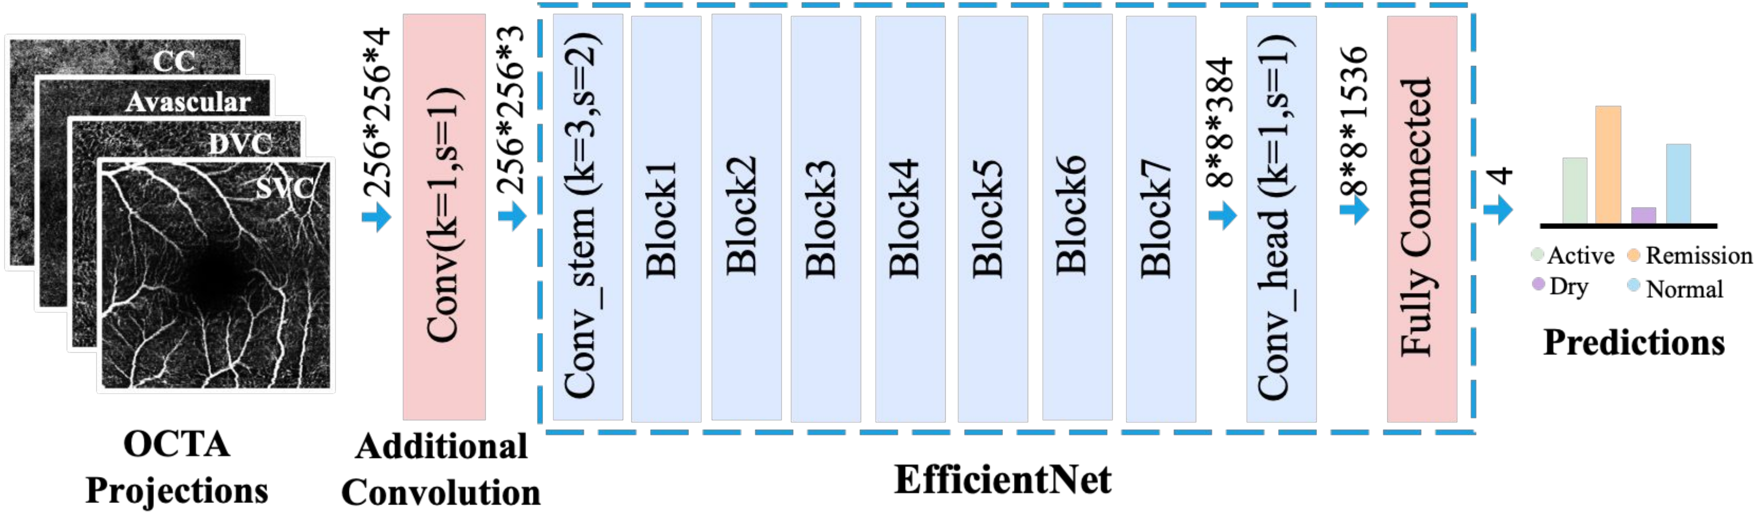

- EfficientNet pretrained on ImageNet + Additional conv
- Warm-Up strategy
  - Fix blue and finetune red only for 600 epoch
  - Finetune all together with smaller learning rate

Tan, M. and Le, Q., 2019, May. Efficientnet: Rethinking model scaling for convolutional neural networks. In International conference on machine learning (pp. 6105-6114).

Contents

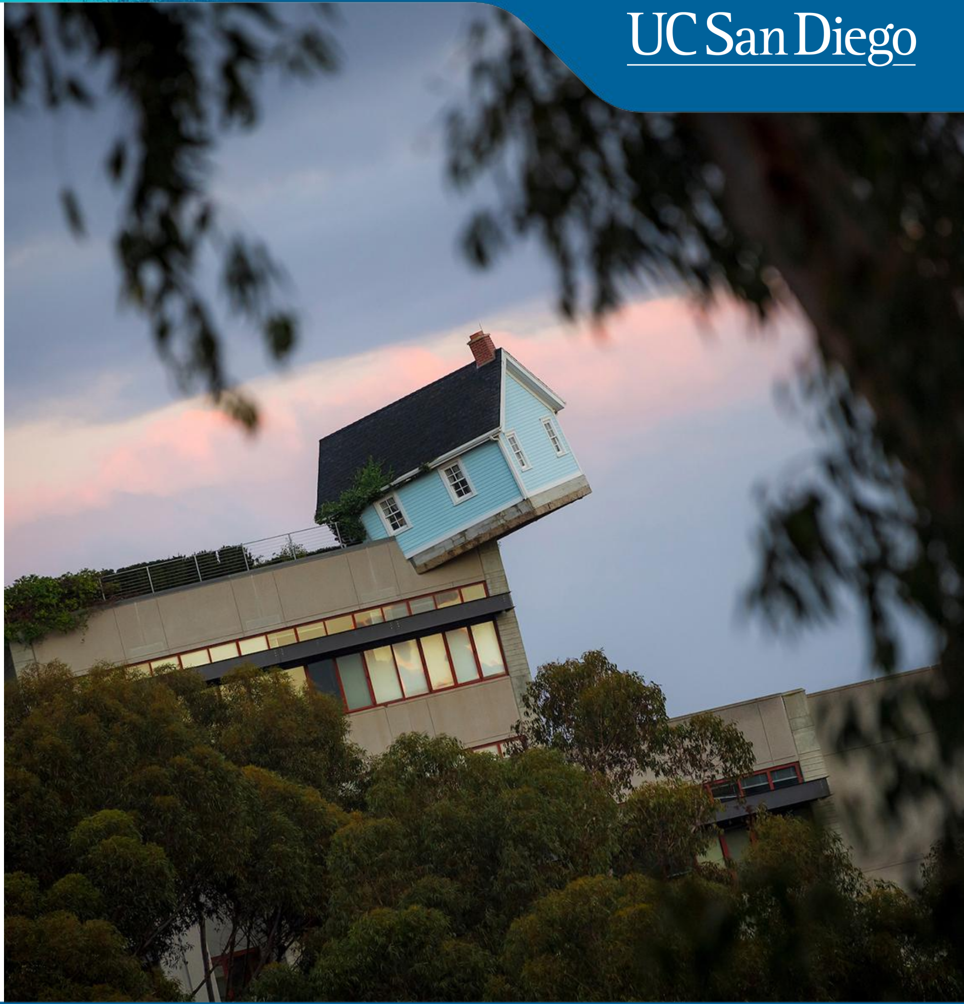

Baseline method

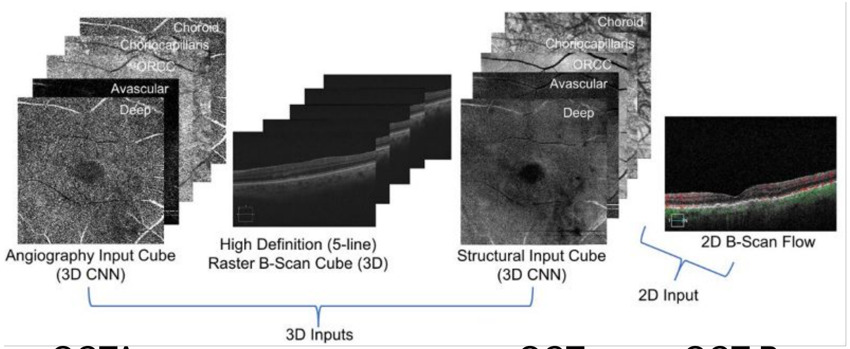

OCTA projections

OCT B-scans

OCT projections

OCT B-scan with OCTA dots

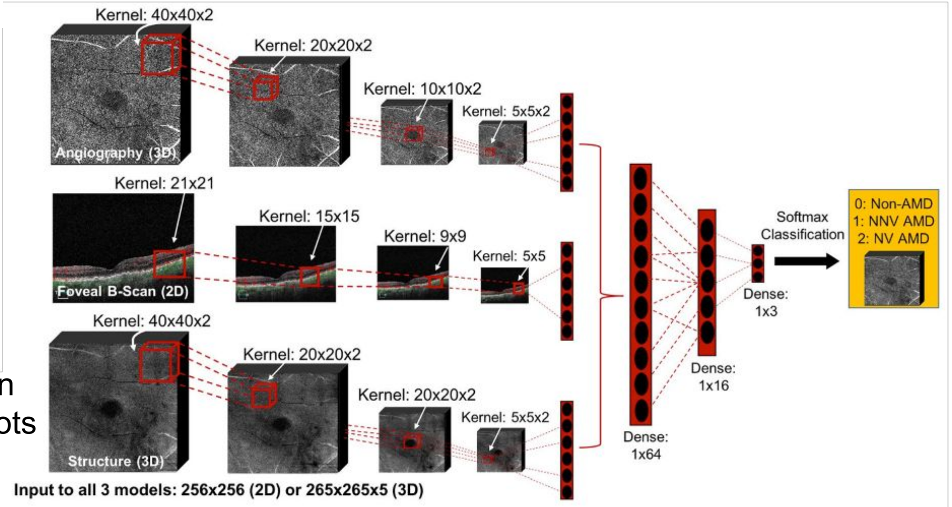

- 3D convolution layer *4
- Fully connected layer *3
- No pretraining

Thakoor, K.A., Yao, J., Bordbar, D., Moussa, O., Lin, W., Sajda, P. and Chen, R.W., 2022. A multimodal deep learning system to distinguish late stages of AMD and to compare expert vs. AI ocular biomarkers. Scientific reports, 12(1), p.2585.

Experiment Results

UC San Diego

Backbone	Training sample	Target domain	Test Acc (Heid/Opto)
CustomCNN [25]	Real_O	Opto only	0.31 / 0.45
	Real_H	Heid only	0.56 / 0.3
EfficientNet [8]	Real_O	Opto only	0.33 / **0.7333**
	Real_H	Heid only	**0.72** / 0.5167
	Real_{H,O}	Heid+Opto	0.67 / 0.7
	Real_H 3D	Heid only	*0.8* / -

Discussion

- Trained within a single domain, the classifier shows optimal performance on the corresponding test set but significantly underperforms on the other.
- Simply combining the two domains helps mitigate imbalanced performance but achieves suboptimal outcomes on both test sets

Experiment Results

UC San Diego

Training sample	Target domain	Class condition	Test Acc (Heid/Opto)
Real_O, Rec_O Fake_O, Idt_O	Opto only	No	0.51 / 0.7167
	Opto only	Supervised	0.55 / **0.8**
	Opto only	Unspv	0.64 / 0.7833
Real_H, Rec_H Fake_H, Idt_H	Heid only	No	0.73 / 0.65
	Heid only	Supervised	**0.78** / 0.6
	Heid only	Unspv	0.77 / 0.6833
Real_{H,O}, Rec_{H,O} Fake_{H,O}, Idt_{H,O}	Heid+Opto	No	0.68 / 0.7333
	Heid+Opto	Supervised	**0.76** / **0.7667**
	Heid+Opto	Unspv	0.74 / 0.75

Discussion

- Two-domain mixture training improves performance on both test sets
- Domain-specific training, the CycleGAN further enhances accuracy.
- The performance improvement is marginal without our proposed class constraint.

Visual inspection

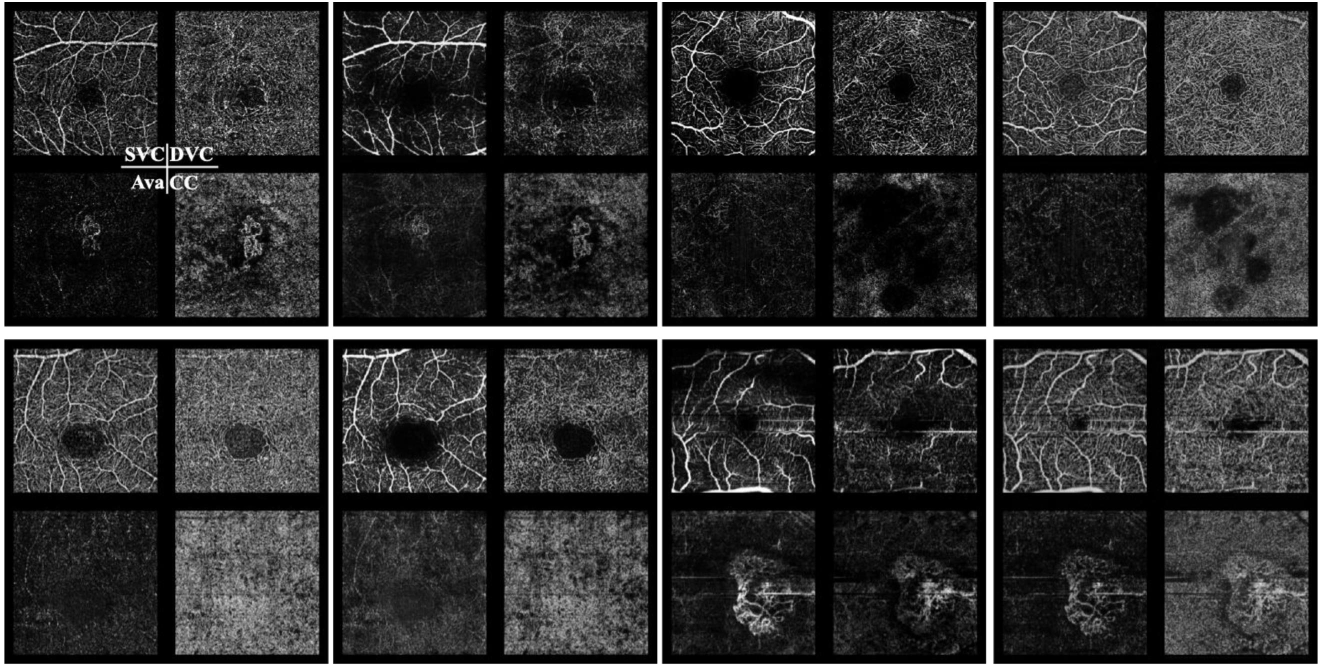

Real Optovue | Fake Heidelberg | Real Heidelberg | Fake Optovue

Real Optovue sample shares vessel patterns with the fake Heidelberg sample.

Same for real Heidelberg and fake Optovue samples

Real and fake Heidelberg samples share the same style.

Same for real and fake Optovue samples

Contents

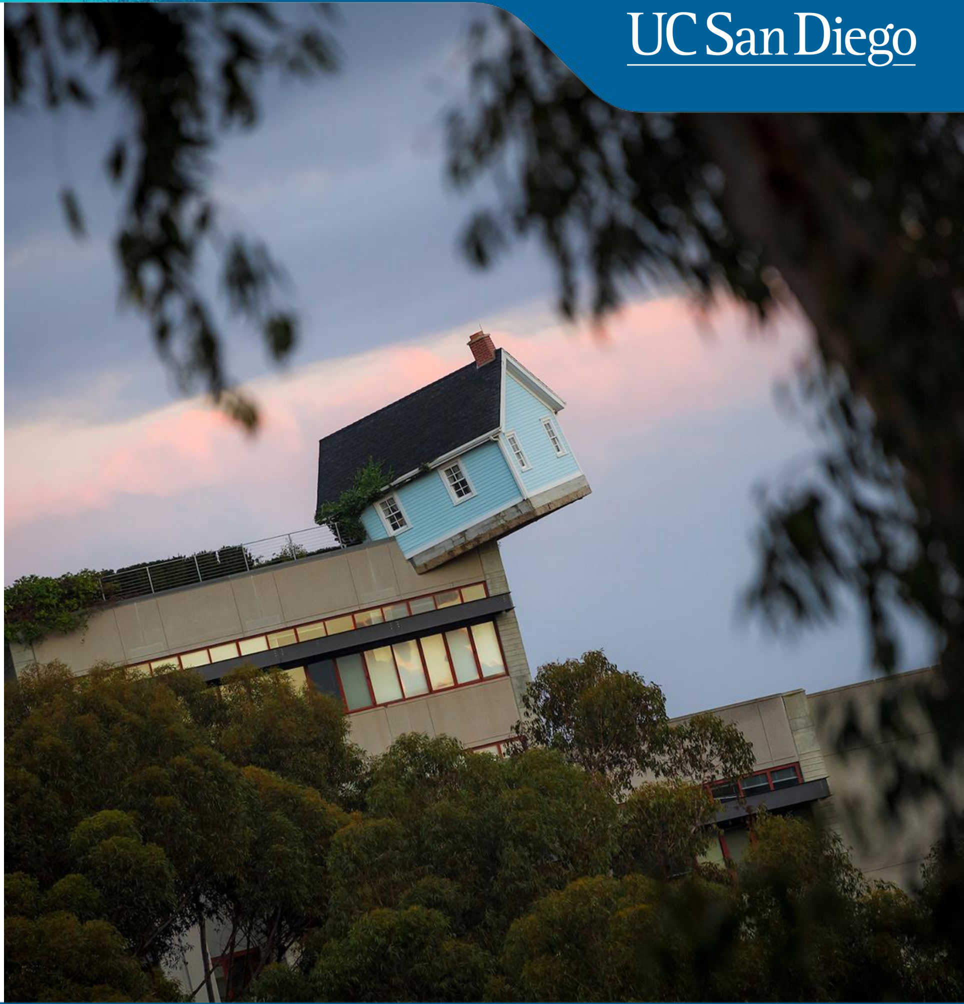

Conclusion

- We proposed CycleGAN for cross-instrument data augmentation, enhancing AMD stage grading accuracy using en-face OCTA images. It increases diversity within each domain and generates unseen data from the other domain.

- We proposed explicit class constraint in both supervised and unsupervised manner to optimize machine recognition quality of the transferred images.

- Experiments show the CycleGAN translator is particularly effective when trained with class-related constraints.

- We hope this study draws more attention to the significance of semantic-related loss design in image processing algorithms aimed at improving classification as downstream task.

References

[3] Macarena Diaz, Jorge Novo, Manuel G Penedo, and Marcos Ortega. Automatic extraction of vascularity measurements using OCT-A images. Procedia Computer Science, 126:273281, 2018.

[4] Mingchao Li, Yuhan Zhang, Zexuan Ji, Keren Xie, Songtao Yuan, Qinghuai Liu, and Qiang Chen. IPN-v2 and OCTA500: Methodology and dataset for retinal image segmentation. arXiv preprint arXiv:2012.07261, 2020.

[5] Yuhui Ma, Huaying Hao, Jianyang Xie, Huazhu Fu, Jiong Zhang, Jianlong Yang, Zhen Wang, Jiang Liu, Yalin Zheng, and Yitian Zhao. ROSE: A retinal oct-angiography vessel segmentation dataset and new model. IEEE transactions on medical imaging, 40(3):928–939, 2020.

[8] H. Zhang, A. Heinke, C. M. B. Galang, D. N. Deussen, B. Wen, D.- U. G. Bartsch, W. R. Freeman, T. Q. Nguyen, and C. An, "Robust amd stage grading with exclusively octa modality leveraging 3D volume," in ICCV Workshop, 2023, pp. 2411–2420.

[25] K. A. Thakoor, J. Yao, D. Bordbar, O. Moussa, W. Lin, P. Sajda, and R. W. Chen, "A multimodal deep learning system to distinguish late stages of AMD and to compare expert vs. AI ocular biomarkers," Scientific reports, vol. 12, no. 1, pp. 1–11, 2022.

UC San Diego

Thank you!

Cross-instrument Data

Example

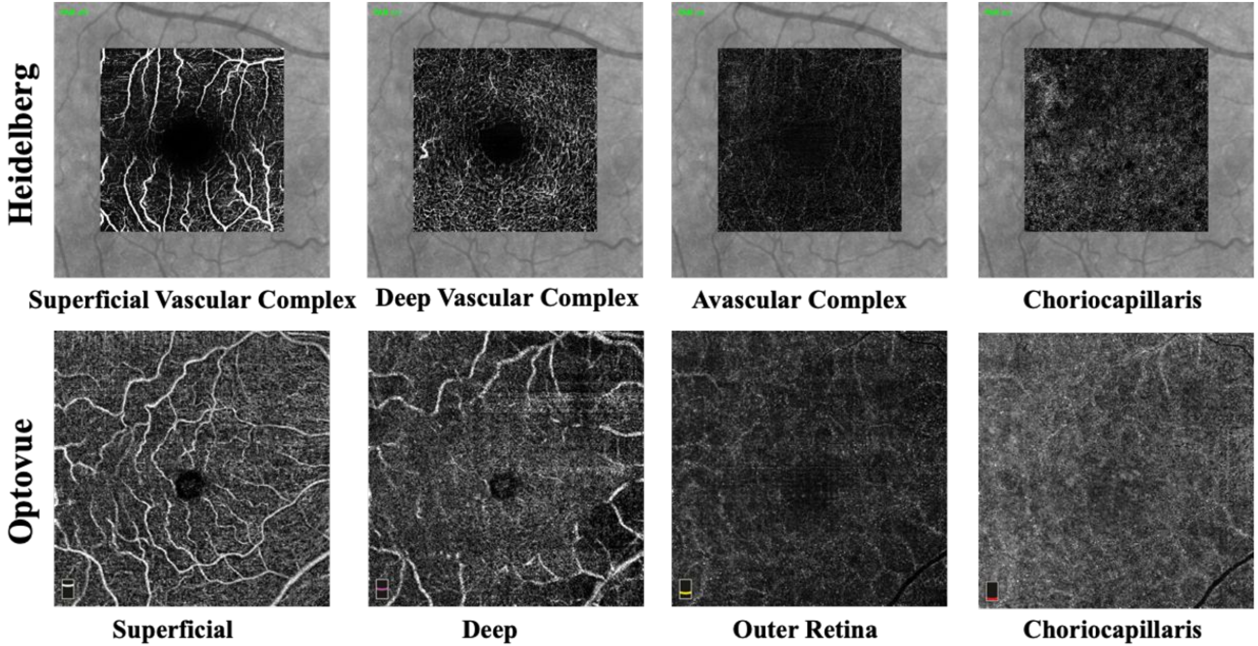

- Crop OCTA area
- Resize to 256
- 256 resolution, 15 FoV

- Central crop 256 area
- 256 resolution, 15 FoV

Dataset

	Active	Remission	Dry	Normal
Heidelberg	324	187	242	140
Optovue	258	83	132	107

- 100 of Heidelberg sample used for inference (25 each category)
- 60 of Heidelberg sample used for inference (15 each category)
- Remaining for training and validation (5 fold validation experiments)
- Oversampling within each category to get a balance training set