FAKULTÄT FÜR !NFORMATIK
Faculty of Informatics

VCBM 2023
13th EG Workshop on Visual Computing for Biology and Medicine

Segmentation, Characterization, and Visualization of Retinal Blood Vessels

Christian Muth, Mark-Christian Bront, and Renata Georgia Raidou
TU Wien, Austria

Motivation and Problem Description

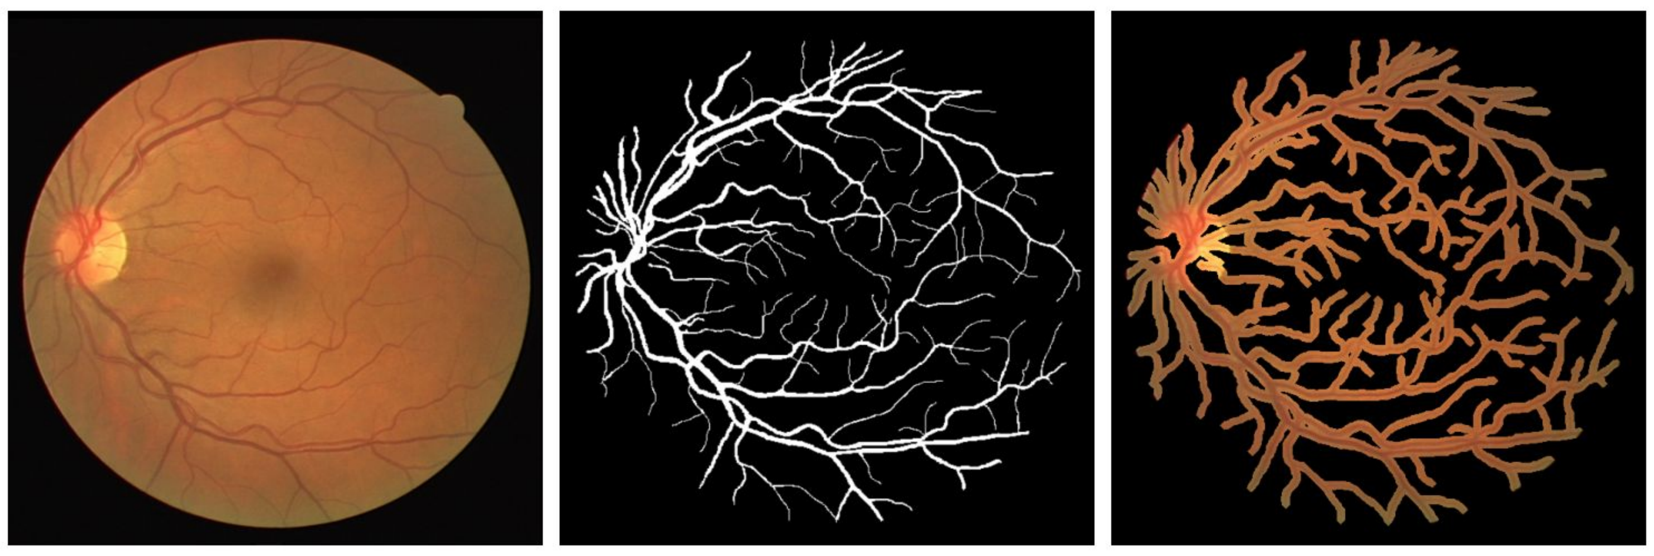

Fundus Segmentation Overlay

Segmentations of retinal vascular structures support ophthalmologists in detecting diseases and predispositions [1]. To streamline the diagnostic process, **automatic segmentation approaches** are being investigated. So far, this has been a **difficult talk, especially for small vasculatures**, as retinal images contain artifacts, such as bright and dark spots.

Generative adversarial networks (GANs) [2] with a U-Net as the deep model have proven useful for the segmentation of retinal images [1]. The training itself can be further augmented by letting the architecture grow incrementally [3].

In this work, we employ the base GAN architecture proposed by Son et al. [1] and **adapt it to a growing GAN architecture for the segmentation of retinal blood vessels.**

Growing Approach

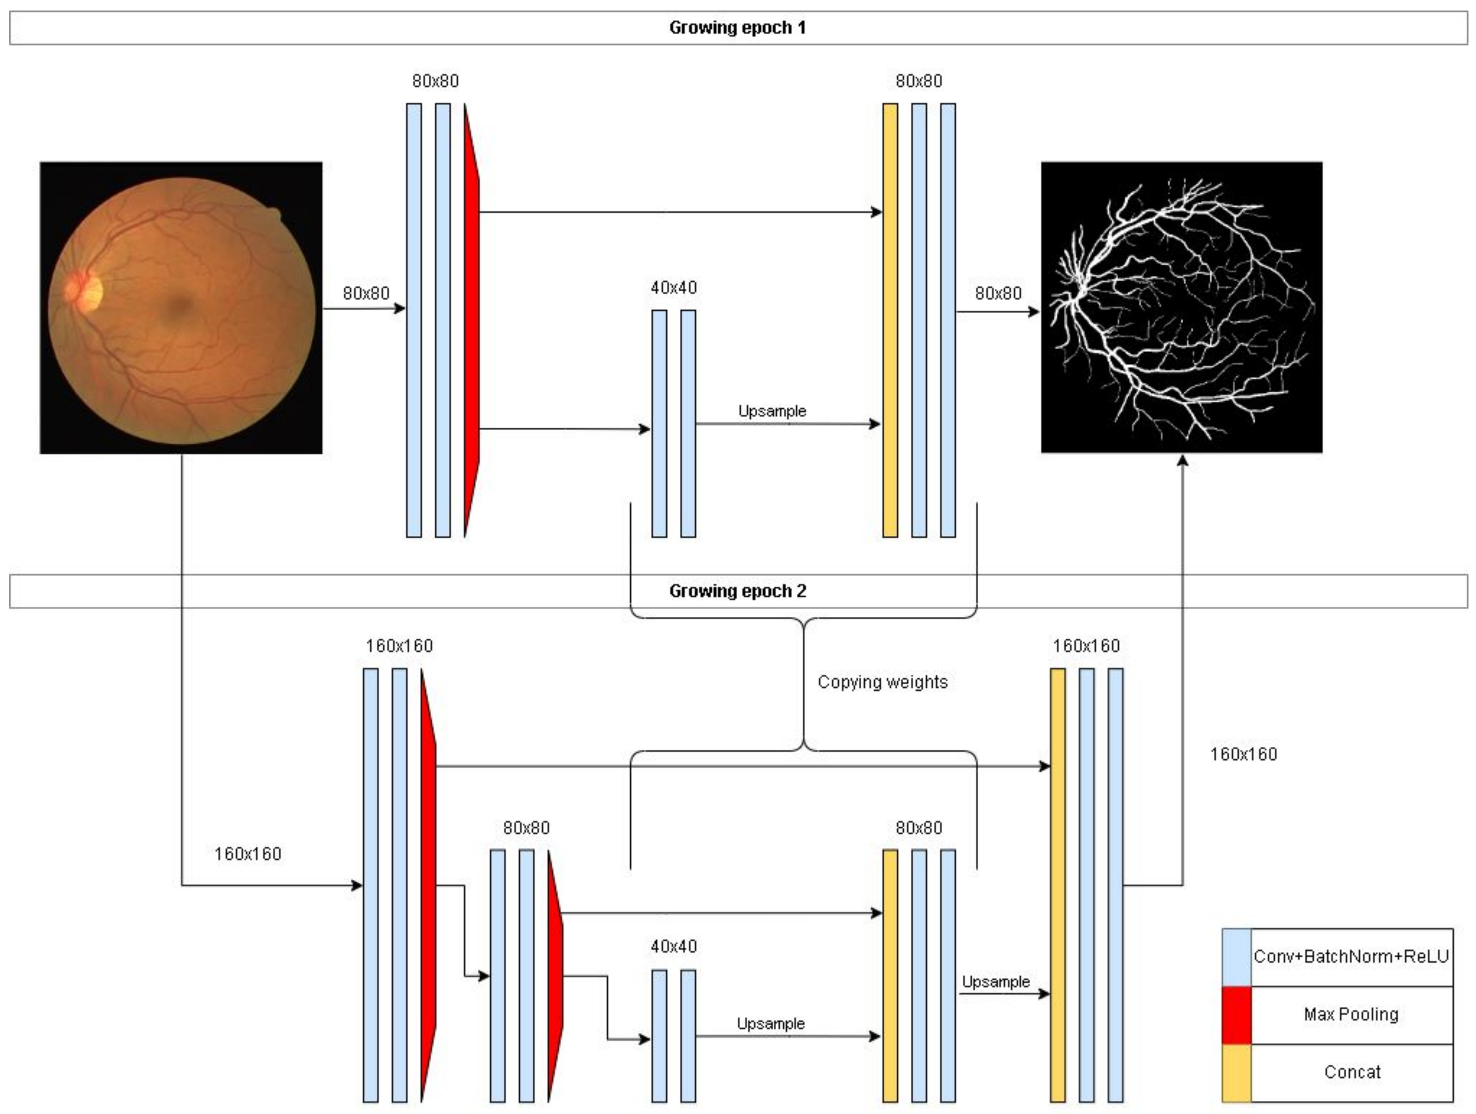

To achieve growth we append **additional layers** at the beginning of both the generator and discriminator, and at the end of the generator. The model is then **trained with the less downsampled images**, and the process is **repeated until the full image size is reached**.

The model is trained at **multiple image sizes** and **different growth strategies**, with varying numbers of rounds, to assess their impact on performance.

Comparison of Results

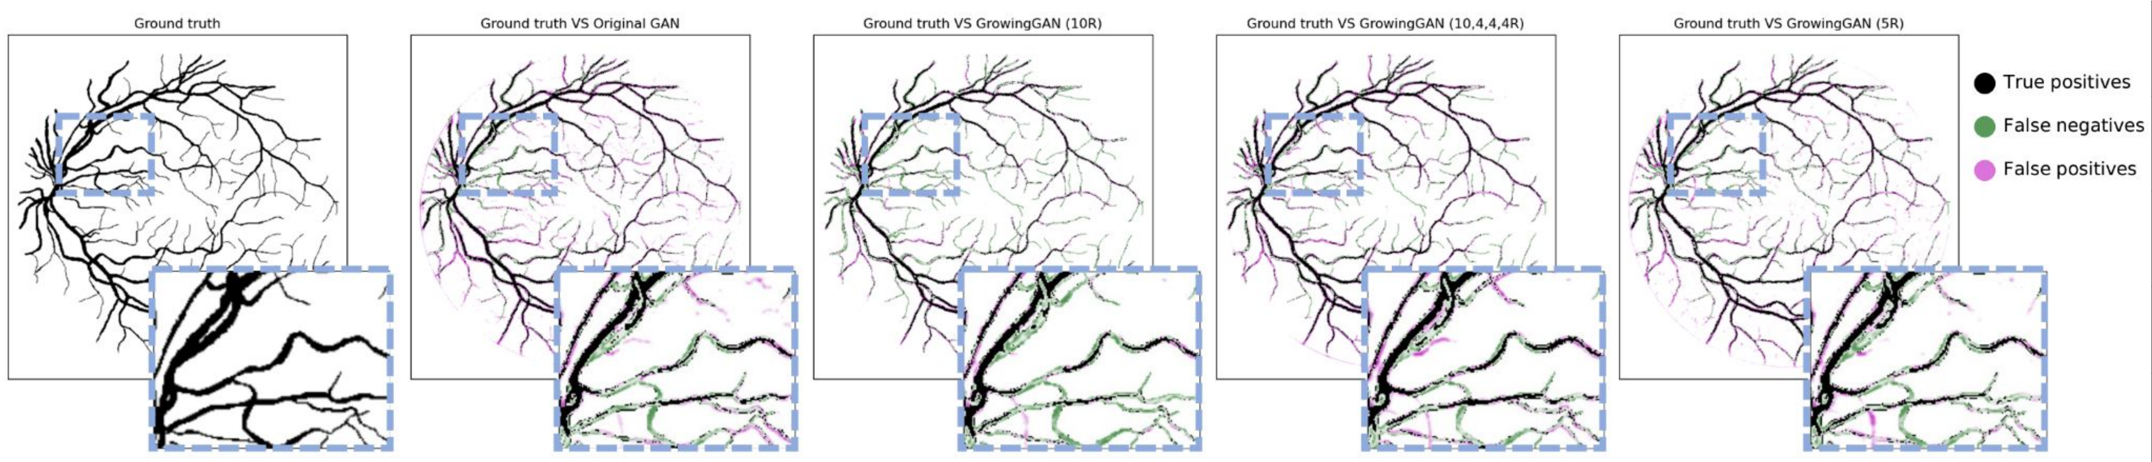

Model	AU-ROC	AU-PR	Dice	Training Time
Original [1]	0.971	0.8892	0.8032	3 hours
Growing GAN (10R)	0.9664	0.8754	0.7604	4.7 hours
Growing GAN (10,4,4,4R)	0.96	0.8748	0.7912	2.3 hours
Growing GAN (5R)	0.967	0.8774	0.7897	2.5 hours

We present hereby a **comparison** of the ground truth, the original GAN by Son et al. [1], and our growing approach with a varying number of rounds per growing step:

To the left, we indicate **qualitatively** the resulting true positives (correct segmentation), false negatives (under-segmentation), and false positives (over-segmentation) measured against the ground truth:
- The approach with **the varying growth generates more false positives** than the other two approaches.
- The growing approach with the **10 rounds yields more fragmented segmentations than the 5 rounds**.
- In areas with thin vasculatures (top right corner of the zoomed-in image), the **original approach generates more false positives than the growing approaches**.

To the right, we show **metrics for the tested models**: area under the receiver operating characteristic (AU-ROC), area under the precision and recall curve (AU-PR), Dice coefficient, and training time.
- The growing GAN model with the **best score was the approach with the 5 rounds.**
- Our growing GAN approach yields **comparable results to the original by Son et al.** [1]

Conclusion

Growing GANs — once more mature — **can enhance the accuracy** of diagnosing retinal vascular diseases, enabling early detection and (tailored) intervention. They can also provide **valuable data for tracking disease progression and treatment efficacy**, leading to better-informed clinical decisions.

Future work might combine aspects of new approaches, using noisy labels, alternative growing procedures, deeper convolutional neural networks, transformer layers, or different loss metrics.

References

[1] SON J., PARK S. J., JUNG K.-H.: Towards accurate segmentation of retinal vessels and the optic disc in fundoscopic images with generative adversarial networks. Journal of Digital Imaging 32, 3 (2019), 499–512

[2] GOODFELLOW I., POUGET-ABADIE J., MIRZA M., XU B., WARDE-FARLEY D., OZAIR S., COURVILLE A., BENGIO Y.: Generative adversarial networks. Communications of the ACM 63, 11 (2020), 139–144.

[3] KARRAS T., AILA T., LAINE S., LEHTINEN J.: Progressive growing of GANs for improved quality, stability, and variation. arXiv preprint arXiv:1710.10196 (2017).

GitHub