Synthesis of Diabetic Retina Fundus Images Using Semantic Label Generation

Joon-Ho Son, Amir Alansary, Daniel Rueckert, Bernhard Kainz, Benjamin Hou
Biomedical Image Analysis, Imperial College London

Abstract

- Automated diagnosis of diabetic retinopathy is challenging due to the lack of available datasets.
- **Can we generate high-quality synthetic data to improve performance on downstream tasks?**
- We propose a two-step process for generating photo-realistic fundus images with accompanying semantic labels, conditioned on diabetic retinopathy severity grade.
- We demonstrate the potential for synthetic retinal fundus images to improve performance in tasks such as lesion segmentation.

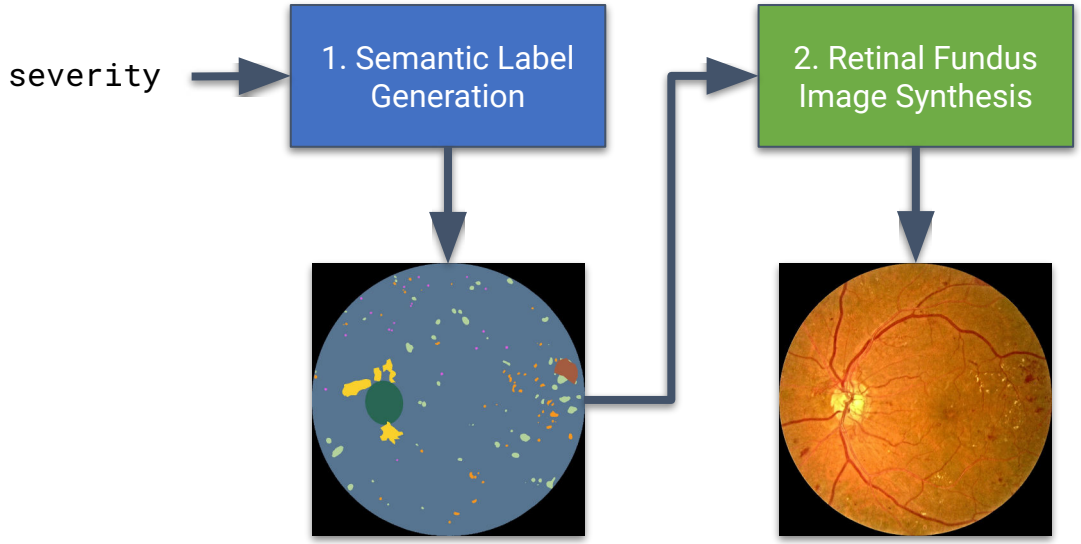

Note

- This work presents preliminary results of a master's thesis which can be found in full at sonj.me/semantic-retina.

Method

- We first generate semantic labels conditioned on DR severity.
- Standard GAN architectures are prone to discriminator collapse when applied directly to semantic labels.
- We use a GAN with **imbalanced network capacities**, **adaptive discriminator augmentation [2]**, and run **two generator iterations** for each discriminator iteration to mitigate this.
- Semantic labels are generated at 256×256, upsampled to 512×512.
- These are then translated to realistic retinal fundus images, conditioned on the generated semantic labels.

1. Semantic Label Generation (ACGAN [1])

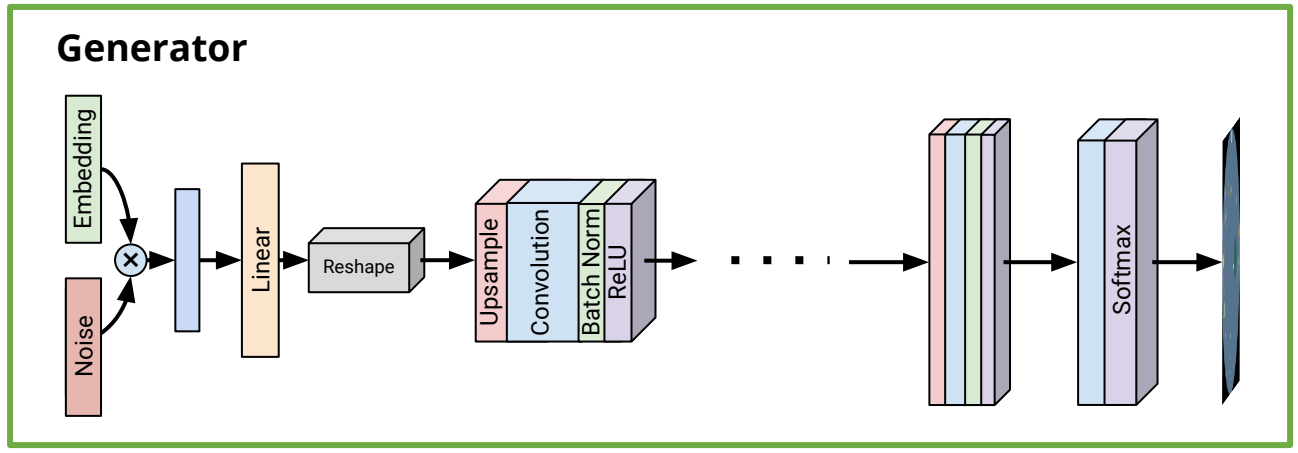

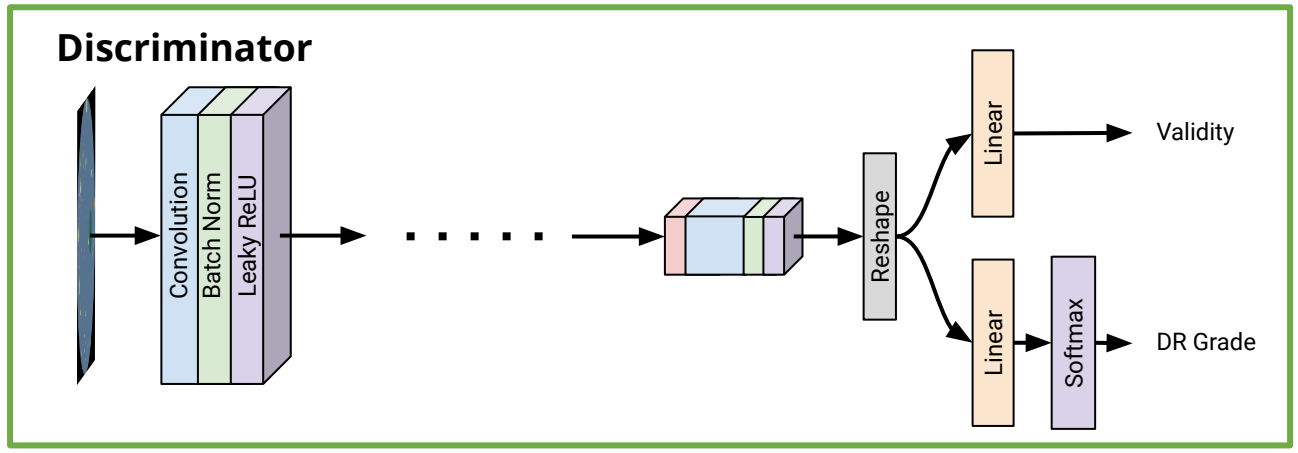

2. Retinal Fundus Image Synthesis (SPADE [3])

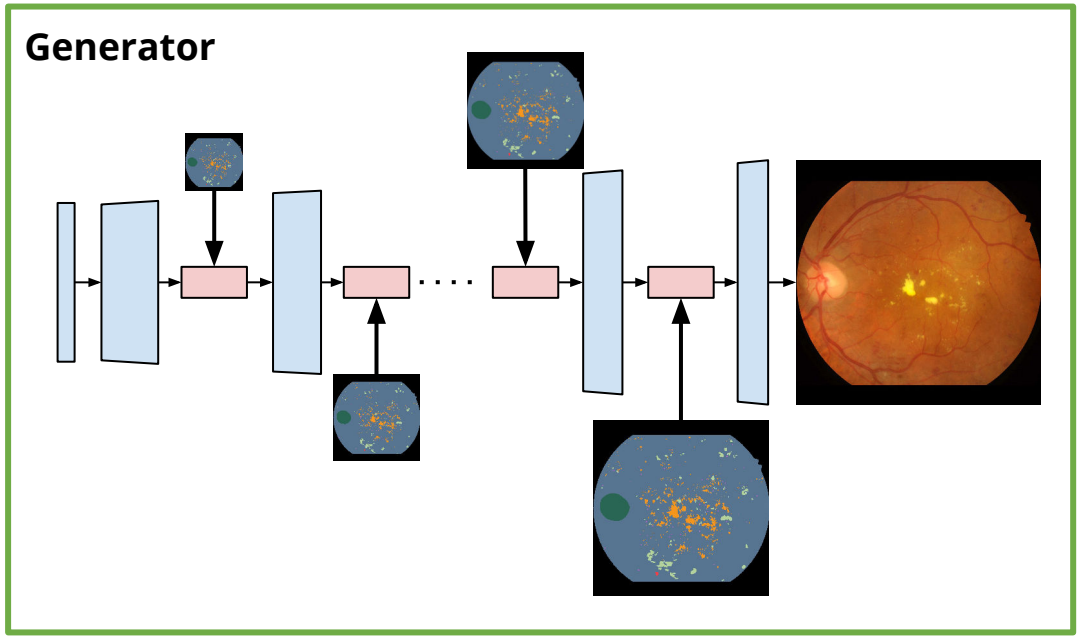

Experiments and Results

- Trained a U-Net to segment hard exudate lesions on different proportions of real and synthetic data.
- Evaluated performance on the held-out set.
- Used a combination of the **FGADR [4]** and **IDRiD [5]** datasets.

Synthetic Data		Precision	Recall	F_1
% Synthetic	Count			
0	0	0.5777	0.4692	0.4694
25	246	0.5860	**0.4938**	**0.4965**
50	738	0.5925	0.4727	0.4824
75	2214	**0.6569**	0.3990	0.4522

Contributions

- A method for generating synthetic retinal fundus images with severity classes and semantic labels, which can be used as segmentation maps.
- We demonstrate the viability of training segmentation models on synthetic data.
- Extensions of this work have been have shown this method to be effective on the classification task also.

Samples

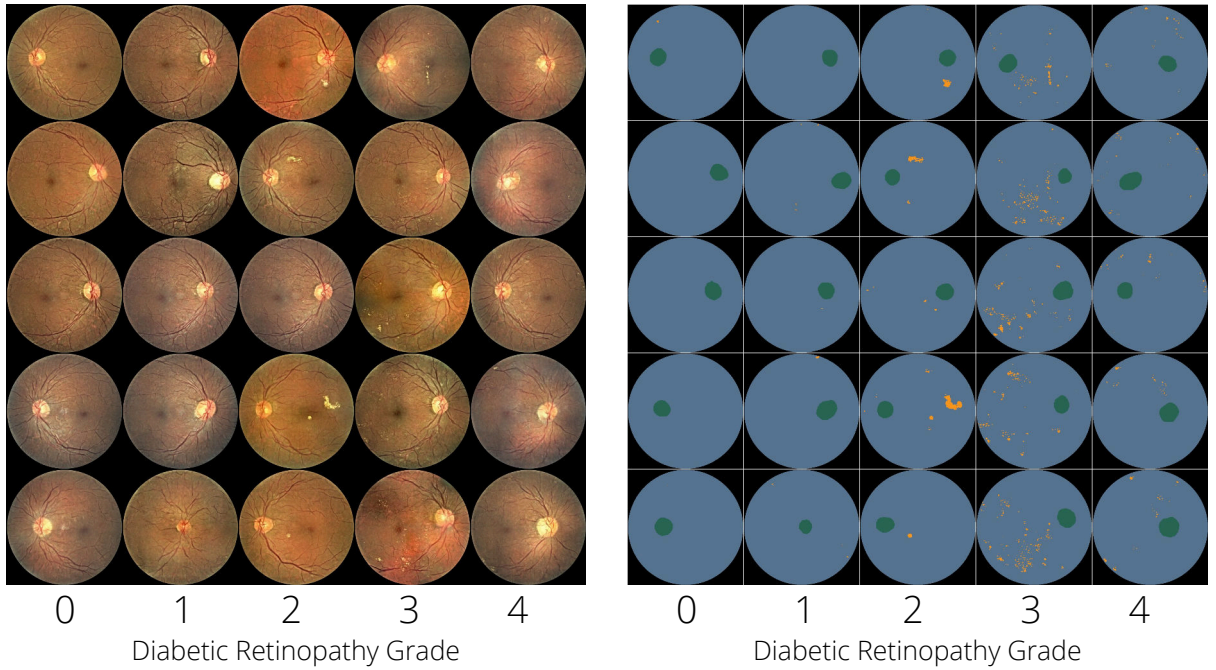

References

[1] Augustus Odena, Christopher Olah, and Jonathon Shlens. Conditional Image Synthesis With Auxiliary Classifier GANs, 2017

[2] Tero Karras, Miika Aittala, Janne Hellsten, Samuli Laine, Jaakko Lehtinen, and Timo Aila. Training Generative Adversarial Networks with Limited Data, 2020

[3] Taesung Park, Ming-Yu Liu, Ting-Chun Wang, and Jun-Yan Zhu. Semantic Image Synthesis with Spatially-Adaptive Normalization. CoRR, abs/1903.07291, 2019

[4] Y. Zhou, B. Wang, L. Huang, S. Cui, and L. Shao. A Benchmark for Studying Diabetic Retinopathy: Segmentation, Grading, and Transferability. IEEE Transactions on Medical Imaging 2020. doi: 10.1109/TMI.2020.3037771.

[5] Prasanna Porwal, Samiksha Pachade, Ravi Kamble, Manesh Kokare, Girish Deshmukh, Vivek Sahasrabuddhe, and Fabrice Meriaudeau. Indian Diabetic Retinopathy Image Dataset (IDRiD): A Database for Diabetic Retinopathy Screening Research. Data, 3(3): 25, 7 2018. ISSN 2306-5729. doi: 10.3390/data3030025