

UC Davis

UC Davis Previously Published Works

Title

Generative modeling for label-free glomerular modeling and classification.

Permalink

<https://escholarship.org/uc/item/4m83j6qx>

Journal

Proceedings of SPIE--the International Society for Optical Engineering, 11320

ISSN

0277-786X

Authors

Lutnick, Brendon

Ginley, Brandon

Jen, Kuang-Yu

et al.

Publication Date

2020-02-01

DOI

10.1117/12.2548757

Peer reviewed



Published in final edited form as:

Proc SPIE Int Soc Opt Eng. 2020 February ; 11320: .

Generative modeling for label-free glomerular modeling and classification

Brendon Lutnick¹, Brandon Ginley¹, Kuang-Yu Jen², Wen Dong³, Pinaki Sarder^{1,*}

¹Department of Pathology and Anatomical Sciences, SUNY Buffalo

²Department of Pathology, University of California at Davis

³Department of Computer Science and Engineering, SUNY Buffalo

Abstract

Generative modeling using GANs has gained traction in machine learning literature, as training does not require labeled datasets. This is perfect for applications in biological datasets, where large labeled datasets are often difficult and expensive to acquire. However, generative models offer no easy way to encode real images into feature-sets, something that is desirable for network explainability and may yield potentially informative image features. For this reason, we test a VAE-GAN architecture for label-free modeling of glomerular structural features. We show that this network can generate realistic looking synthetic images, and be used to interpolate between images. To prove the biological relevance of the network encodings, we classify small-labeled sets of encoded glomeruli by biopsy Tervaert class and for the presence of sclerosis, obtaining a Cohen's kappa values of 0.87 and 0.78 respectfully.

Keywords

Unsupervised data-mining; variational autoencoder; generative adversarial network; glomeruli

1. INTRODUCTION

Recent advancements in machine learning techniques have created exciting technologies for synthesis of synthetic natural images (generative modeling). In particular, GANs¹ (generative adversarial networks) have seen an explosion of popularity in machine learning literature due to the high quality synthetic images they produce as well as the ability to be trained using unlabeled datasets. This latter quality makes these networks particularly intriguing for adaption in pathological image datasets, where explicitly labeled images suited to machine learning tasks are limited. GANs map a latent code (generally a noise vector) to realistic natural images but lack the ability to perform the inverse transformation. The ability to map an image to a latent code is desirable not only to explain the output of GANs, but also to predict potentially informative features from input images. To accomplish this we use of a VAE-GAN² (the combination of a VAE³-Variational Autoencoder and GAN). Using this architecture, real images are encoded to produce a multivariate latent

*Address all correspondence to: Pinaki Sarder, Tel: 716-829-2265; pinakisa@buffalo.edu.

code, and realistic synthetic images are generated from a latent code sampled from a multivariate Gaussian distribution. While image generation was the original intention of these networks, we propose encoding of images to a latent code has more uses that are impactful to the field of pathology.

At core of the VAE-GAN is an encoder-decoder architecture, which acts as a funnel for input image information. Simply, the encoder learns a series of transformations that condense the information in input images. This produces a latent code, which can be understood as a vector of image features. This code is passed to the decoder, which learns to transform the code back into a real looking image. It is useful to consider this network from an information theory perspective, where this can be described as a lossy data compression of input images. Here the encoder reduces the redundancy of information in the highly correlated input images, where each image pixel can be considered a data dimension. The latent code acts as an information channel, which has a limited bandwidth. To reconstruct the images with high fidelity, the network must encode image features that have a higher density of information, and are therefore descriptive of the original image. This approach is powerful because the latent code is learned without the need for data labels. The network is self-supervised, using a discriminator network to approximate the data distribution of input images and penalize the encoder/decoder for poor performance. In this work, we explore the robustness of these automatically determined features, validating their usefulness for the description of pathologically derived priors.

2. RESULTS

We trained a VAE-GAN using histopathological glomerular images segmented from whole slide images (WSIs) of murine and human kidney sections. Despite never seeing data labels, the VAE-GAN architecture (Figure 4) is able to produce realistic looking synthetic glomeruli images as well as smooth interpolations between real input images as seen in Figure 1. We note that while generated images from encoded real images look similar to the input images, there are significant micro-anatomical structural differences. We hypothesize this is due to the adversarial loss used, which uses a discriminator network to penalize the network for synthetic looking images. While the output from one of the middle convolutional layers is compared between corresponding real and synthetic images the network never directly compares the input and generated images directly.

To verify the ability of our network to capture pathologically important features we classify smaller sets of labeled human glomeruli images based on their latent code predicted by the VAE-GAN. On a set of 1193 images labeled as normal vs sclerotic, we use a simple multilayer perceptron classifier⁴ (MLP) to predict sclerosis. This classifier, which predicts sclerosis from the latent code of the labeled images, achieved 0.964 sensitivity and 0.821 specificity, with a Cohen's kappa of 0.78 using fivefold cross validation.

On a separate set of glomeruli extracted from 121 biopsies labeled by diabetic nephropathy Tervaert class⁵ we use a recurrent neural network⁶ (RNN) to predict biopsy level Tervaert class. Specifically this is done by considering the sequence of latent codes of glomeruli within a biopsy, which is fed into two-stacked long short-term memory (LSTM) units. This

network models Tervaert biopsy class from the sequence of glomerular features using a regression analysis. Using 10 fold cross validation we achieve a linear weighted Cohen's kappa of 0.87. Examples of these images are shown in Figure 2, and classifier confusion matrices in Figure 3.

3. METHODS

We have developed a VAE-GAN architecture² (Figure 4) in Tensorflow, which utilizes 2D convolutional layers to encode color images of size 256×256 pixels. This architecture is roughly based in the DCGAN architecture⁷ but utilizes modern machine learning approaches to produce high quality images, including: subpixel upsampling⁸, and instance normalization⁹. This architecture uses convolutional up and down sampling, followed by fully connected layers to predict the mean and standard deviation for a multivariate Gaussian distribution with 256 dimensions (a reduction of 256 times the original input dimensionality). This Gaussian distribution is fed to the generator of a GAN, which reconstructs synthetic images, attempting to fool the discriminator network. The specifics of this network are described in Figure 4, which shows an overview of the network architecture.

2.1. Network training

The network was trained for 87 epochs with a batch size of 16 using an RMSProp optimizer with a learning rate of 0.0003. The training set consisted of 87K PAS, H&E, and Trichrome stained glomeruli images: 59930 came from human biopsies and 27508 from murine kidney sections. During training, images were randomly cropped, rotated, and flipped to prevent against overfitting. The training images were segmented from WSIs using H-AI-L¹⁰, a neural network based segmentation technique developed by our lab. The glomeruli were extracted from WSIs at 10X(1.008μm/pixel). The background tissue was excluded from the images (set to white) for the sake of simplicity.

2.2. Murine model

We used streptozotocin (STZ) treated diabetic mouse model¹¹, and an STZ treated nephrin knockdown mouse model¹² for generating mouse data. All animal studies were performed in accordance with protocols approved by the University at Buffalo Animal Studies Committee.

2.3. Human data

Biopsy samples from human diabetic nephropathy patients were obtained from Kidney Translational Research Center at Washington University School of Medicine directed by Dr. Sanjay Jain and from Vanderbilt University Medical Center via collaborator Dr. Agnes Fogo. Description of these samples are available in our recent publication¹³. The glomerular structural changes in these biopsies suggest DN related changes spanning different DN stages as discussed in Tervaert *et al*⁶. As control, renal tissue samples of non-diabetic patients with renal cell carcinoma were considered. We used sections with completely normal renal tissues as verified by Dr. Sanjay Jain and Dr. John E. Tomaszewski. We also used human transplant biopsy samples provided by co-author Dr. Jen from University of

California at Davis. Human data collection procedure followed a protocol approved by the Institutional Review Boards at University at Buffalo, Vanderbilt University Medical Center, and University of California at Davis. Ground-truth annotations of DN structural disease state were performed by the co-author Dr. Kuang-Yu Jen.

2.4. Imaging and data preparation

Tissue slices of 2 μm (for murine data) and 2–5 μm (for human data) were stained using diverse histological stains, and imaged using a whole-slide imaging scanner (Aperio Versa, Leica, Buffalo Grove, IL). We followed similar imaging protocol as described in our recent works¹⁴.

2.5. MLP sclerotic classification

The latent encoding of the 1193 labeled glomeruli was used as input to a MLP classifier. We used the implementation of a MLP from the python package sklearn. We specify 400 hidden nodes, and two output nodes (for normal and sclerotic probability). This network was trained for a maximum of 300 iterations, and stopped when the convergence criteria of sklearn were met. Figure 3 presents the results of five-fold cross validation.

2.6. RNN Tervaert classification

All the glomeruli from a biopsy are extracted using H-AI-L as described above. These are encoded using the VAE- GAN and the latent codes are concatenated to form a sequence. We use two stacked LSTM layers with 50 and 25 hidden features respectively. This was followed by a fully connected layer to predict the Tervaert DN stage per biopsy and modeling it as a continuous number with a single output node. The network performance is computed using the Euclidean distance of the predicted stage from the true stage. More information about the particulars of this RNN biopsy classifier is outlined in our recent journal article.¹³ Figure 3 presents the results of ten-fold cross validation per biopsy.

4. DISCUSSION

The work presented in this write-up utilizes a custom VAE-GAN architecture. We present an extension of our earlier work on automatic unsupervised feature determination of glomeruli using a VAE¹¹, adding the ability to generate convincing synthetic glomeruli images. The motivation for this work is primarily as a feature finding technique. While purely GAN architectures are capable of generating realistic looking images, they offer no way to encode real images to the latent distribution. The addition of an encoder network is essential for the prediction of a latent code from an input image. This is the motivation for using a VAE architecture, which is well known for its compact data modeling. However, using images as input data, VAE's are notoriously poor at generating realistic results. This is remedied by the addition of the discriminator network.

We expect the ability of this network to generate realistic looking images to bring some explainability to the network. We hypothesize the ability to walk the latent space (see Figure 1) will be very useful for feature exploration. Decoding latent codes back to the image space allows effect of perturbation of the latent code to be visualized on the output image.

However, the discriminator cost function used by the network does not directly compare input images to the corresponding generated images. This means that structures, which are encoded from real images, need not be generated by the decoder network to fool the discriminator. Therefore, the synthetic results, while proving that the network is indeed learning to encode tissue structure, may not be directly useable to understand the latent code in the current form.

In summary, we have developed a pipeline for mapping images to latent codes without labels. We show this mapping is predictive of common pathological classifications using smaller labeled datasets. Particularly exiting is the results of Tervaert DN classing using the RNN model. This is the same classification method that we recently published in JASN¹³, where and derived biologically motivated features were used as input to the RNN. Using these carefully hand crafted features, we report a Cohen's Kappa of 0.9 in the work¹³. Using the same holdout technique, and inputting the latent codes of our VAE-GAN to the RNN classifier, we report a Kappa of 0.87. This performance is comparable to the highly motivated hand derived feature set from the JASN work, despite having no supervision and no data labels.

Finally, this architecture is capable of generating realistic looking images, and producing smooth interpolations between input images. We use this as another validation of the networks ability to encode relevant biological features.

5. FUTURE WORK

While our application of a VAE-GAN architecture has been customized, there is not currently an adequate network penalty for dissimilar encoded and generated images. In future work we intend to modify the loss function to mitigate this issue. This can be done by adapting a cycle-consistent loss function such as the one used in the newly popular cycleGAN¹⁵. This would enable the decoder to be used as a latent exploration tool, helping the explainability of our network. Additionally we plan to explore the addition of sparse labels to the VAE-GAN to encourage the network to learn biologically relevant features in a semi-supervised fashion.

ACKNOWLEDGEMENT

The project was supported by the faculty startup funds from the Jacobs School of Medicine and Biomedical Sciences, University at Buffalo; Buffalo Blue Sky grant, University at Buffalo; NIDDK Diabetic Complications Consortium grant DK076169; NIDDK grant R01 DK114485 & DK114485 02S1; and NIDDK CKD Biomarker Consortium grant U01 DK103225. We thank NVIDIA Corporation for the donation of the Titan X Pascal GPU used for this research (NVIDIA, Santa Clara, CA).

REFERENCES

1. Goodfellow I et al. in Advances in neural information processing systems. 2672–2680.
2. Larsen ABL, Sønderby SK, Larochelle H & Winther O Autoencoding beyond pixels using a learned similarity metric. arXiv preprint arXiv:1512.09300 (2015).
3. Kingma DP & Welling M Auto-Encoding Variational Bayes. ArXiv e-prints 1312, arXiv:1312.6114 (2013).

4. Gardner MW & Dorling S Artificial neural networks (the multilayer perceptron)—a review of applications in the atmospheric sciences. *Atmos. Environ* 32, 2627–2636 (1998).
5. Tervaert TW et al. Pathologic classification of diabetic nephropathy. *J Am Soc Nephrol* 21, 556–563, doi:10.1681/ASN.2010010010 (2010). [PubMed: 20167701]
6. Hochreiter S & Schmidhuber J Long short-term memory. *Neural computation* 9, 1735–1780 (1997). [PubMed: 9377276]
7. Radford A, Metz L & Chintala S Unsupervised representation learning with deep convolutional generative adversarial networks. arXiv preprint arXiv:1511.06434 (2015).
8. Shi W et al. in *Proceedings of the IEEE conference on computer vision and pattern recognition*. 1874–1883.
9. Ulyanov D, Vedaldi A & Lempitsky V Instance normalization: The missing ingredient for fast stylization. arXiv preprint arXiv:1607.08022 (2016).
10. Lutnick B et al. An integrated iterative annotation technique for easing neural network training in medical image analysis. *Nature Machine Intelligence* 1, 112–119, doi:10.1038/s42256-019-0018-3 (2019).
11. Lutnick B, Tomaszewski JE & Sarder P Leveraging unsupervised training sets for multi-scale compartmentalization in renal pathology. *Proceedings of SPIE (SPIE Medical Imaging 2017: Digital Pathology)* 10140, 101400I: 101401–101407, doi:10.1117/12.2254750 (2017).
12. Li X et al. Nephric Preserves Podocyte Viability and Glomerular Structure and Function in Adult Kidneys. *J Am Soc Nephrol* 26, 2361–2377, doi:10.1681/ASN.2014040405 (2015). [PubMed: 25644109]
13. Ginley B et al. Computational segmentation and classification of diabetic glomerulosclerosis. *J. Am. Soc. Nephrol* 30, 1953–1967 (2019). [PubMed: 31488606]
14. Ginley B, Tomaszewski JE, Yacoub R, Chen F & Sarder P Unsupervised labeling of glomerular boundaries using Gabor filters and statistical testing in renal histology. *Journal of Medical Imaging* 4, 021102: 021101–021112, doi:10.1117/1.JMI.4.2.021102 (2017). [PubMed: 28680908]
15. Zhu J-Y, Park T, Isola P & Efros AA in *Proceedings of the IEEE international conference on computer vision*. 2223–2232.




Figure 1. Interpolation between validation human glomeruli images with different stain. Input images are encoded into a code of 256 numbers, which is representative of the input image. These codes are interpolated via spherical linear interpolation and decoded back into intermediate images. The VAE-GAN network was trained using the input images alone—without data labels. It is interesting to note the difference between the input images and generated images, this may indicate a disconnect between the encoder and generator networks due to the adversarial loss.

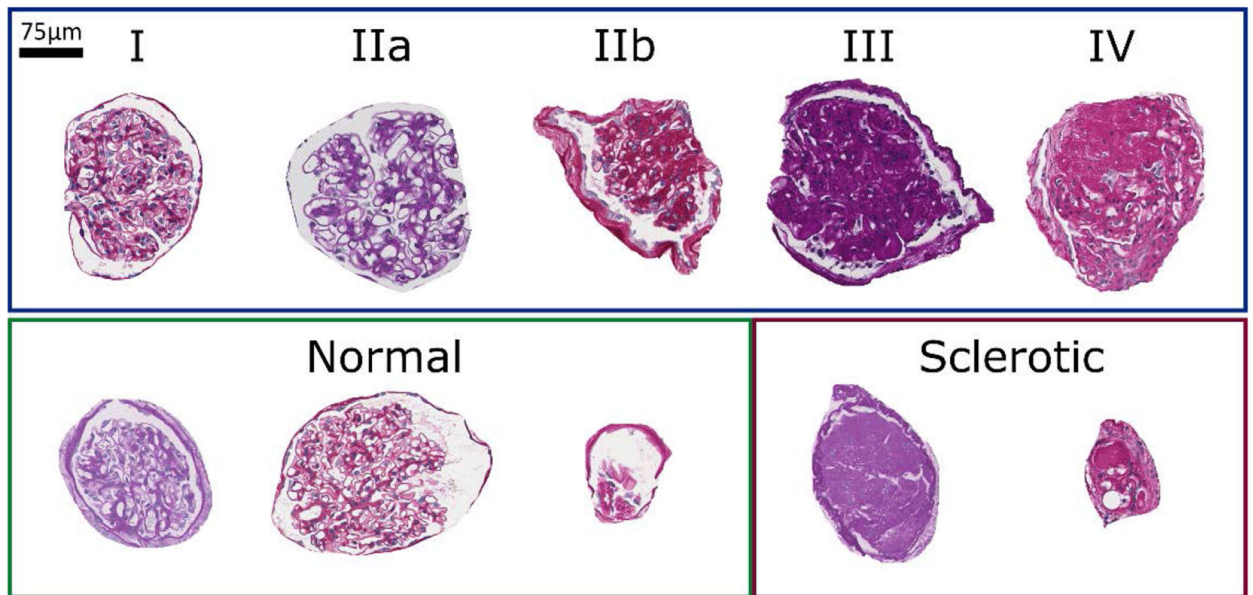


Figure 2. Examples of the labeled training images.

Selected example images used to produce the classification results presented in Figure 4. These sets contain glomeruli from 121 Tervaert staged biopsies and 1193 individual glomeruli labeled for sclerosis respectively. The Tervaert biopsies are PAS stained while the sclerotic glomeruli set contains a combination of H&E and PAS stains. The glomeruli are cropped from segmented WSIs and extracted into 256×256 pixel patches with a white background.




Figure 3. Confusion matrices from classification using VAE-GAN.

Classification results using the VAE-GAN latent code. The VAE-GAN network was trained without using labeled data, and the encoder was used to predict 256 features for each input image, examples are given in Figure 2. **(A)** Latent codes produced by the encoder network of the VAE-GAN were input to a multilayer perceptron (MLP) to predict glomerular sclerosis. These latent codes (image features) were randomly shuffled and input to a MLP with 400 hidden nodes, which was trained using 5-fold cross validation. The performance of the aggregated folds are presented above. We observe a Cohen's Kappa of 0.78 for Sclerotic classification. **(B)** Latent codes were assembled into sequences using all the segmented glomeruli from a Tervaert staged biopsy. These sequences were fed into a recurrent neural network (RNN) which was trained to predict the Tervaert stage from the input glomerular feature sequence using 5-fold cross validation. The performance of the aggregated folds are presented above. We observe a Cohen's Kappa of 0.87 for Tervaert classification.

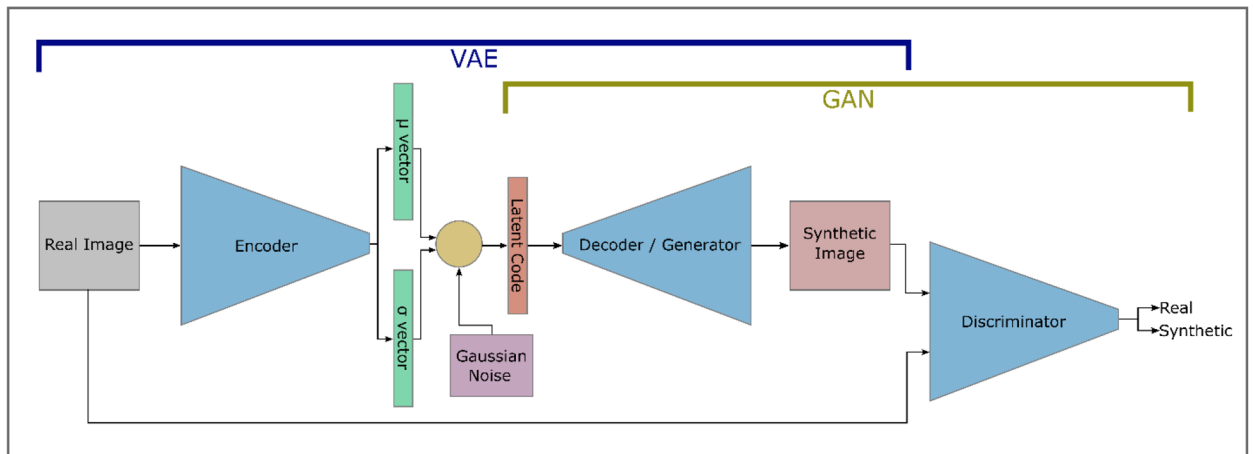


Figure 4. Example of the VAE-GAN architecture used.

Basic layout of the VAE-GAN used in this work. Images are fed into a convolutional encoder network, which uses a series of fully connected layers to predict mean and standard deviation vectors of length 256. These vectors are used in addition to a Gaussian noise vector to predict the latent code (length 256) via variational inference. A convolutional decoder/generator network is used to convert this latent code to a synthetic image. Synthetic and real images are fed into a discriminator network, which learn to classify them. The discriminator and generator network compete during training leading to high quality synthetic images. After training, the encoder network can be used also to predict a latent code for an input image. The decoder network can be fed Gaussian noise and used alone to produce synthetic images.