Generative Clearing for Deep Tissue Imaging

Author name(s) withheld EMAIL(S) WITHHELD
Address withheld

Editors: Under Review for MIDL 2022

Abstract

Investigating cellular level tissue architecture requires the imaging of intact biological samples from 3D volumes with high-resolution imaging methods, including confocal microscopy. A significant challenge for the quantitative analysis of such volumetric data using computer vision is the degradation of image quality at increased depths due to light scattering, absorption and optical factors. Here, we introduce a generative cycle consistent adversarial network (Cycle-GAN) to mitigate these effects, which exploits the property that since the tissue is self-similar, the appearance of the shallow layers can serve as a proxy for that of depth degradation-free data. The network model obtains a bi-directional mapping between the shallow and deep layers so that the restored deep layers resemble the shallow ones. We demonstrate this approach's utility on a dataset of images obtained by confocal imaging of thick cardiac tissue sections from the mouse. Our experiments show that the restored deeper layers are qualitatively and quantitatively similar to the shallow ones, that the restored tissue images are amenable to geometric analysis and that in general such an approach outperforms other methods both qualitatively, as well as quantitatively.

Keywords: Microscopy, reconstruction, confocal imaging.

1. Introduction

In cellular and molecular biology, 3D tissue is often examined using high-resolution imaging methods such as confocal microscopy. This permits an analysis of 3D tissue architecture at the micron scale, at which cell membranes, cell nuclei, capillaries, and larger blood vessels, and extracellular material become visible. Deep tissue imaging has many applications in biology and medicine, including biopsy assessment, tissue ultra-structure analysis, and assessment of tissue health (Feuchtinger et al., 2016; Abadie et al., 2018). Nonetheless, thick biological tissue samples often yield poorly resolved imaging data in deeper layers, where light penetration can be poor. In addition, optical factors such as light scattering and aberrations caused by the lens also diminish image quality. In recent years, active tissue clearing methods have been introduced to reduce the impact of the distinct refractive indices of heterogeneous biological tissue types (Chung et al., 2013; Richardson and Lichtman, 2015). The tissue samples, combined with compatible refractive index matching materials, have reduced scattering and absorption, allowing for better penetration of the optical signal. Even so, deeper layers in actively cleared tissue can still suffer from blurring and image degradation effects, when imaging with a conventional confocal microscope, as illustrated in Fig. 3 (top row). With knowledge of the point spread function of the microscope's objective, it is possible to apply a deconvolution method to sharpen these images (Dey et al., 2006), but the depth degradation effects remain.

How can ultra-structure cellular properties be assessed from microscopy in these deeper layers in actively cleared tissue, in the face of such degradation effects? This presents

© 2022 Anonymous.

a fundamental computer vision challenge, since the ideal (ground truth) restored images are not available. In the present article we propose to use a generative model to solve this problem, where the key idea is to use the texture and intensity pattern of the lower layers as a proxy for what restored images from deeper layers should resemble, under the assumption that the tissue is self-similar. A deep tissue image is then altered so as to match the statistics of lower layers while being faithful to its original isophote geometry. We introduce a cycle consistent generative adversarial model (Cycle-GAN) to accomplish this, a model which has already shown great promise in computer vision and medical imaging tasks involving style transfer (Zhu et al., 2017; Tmenova et al., 2019), super-resolution (You et al., 2019), denoising (Kang et al., 2019; Song et al., 2020; Li et al., 2019) and image enhancement (Ma et al., 2020). The particular problem we address falls in the realm of image restoration, since unlike in style transfer applications, where the underlying source and target domains are different, in ours both are the same, i.e., microscopy images of the same tissue sample. A similar restoration problem has been addressed in (Xiao et al., 2020), where a pairing of clean and noisy images is required. This is achieved by using elaborate multiview light sheet microscopy. In order to acquire images of the same tissue under different levels of depth degredation, the authors image the sample from six different views and then apply 3D registration. They also use a step by step layering of 100μ thick slices of non-florescent tissue on top of the tissue sample. In contrast, our approach uses the simpler standard confocal microscopy set up and does not require multi-view 3D registration since we do not assume pairing.

We have compared our Resnet based noise model free generative approach to classical generative models based on dictionary learning, as well as modern deep learning based methods such as residual learning based denoising (Zhang et al., 2017) and analytically derived methods (Pronina et al., 2020; Li et al., 2017) which assume a degredation model. All these methods have shown promise in various microscopy image enhancement and denoising tasks. However, in our work we exploit the similarity of structure in a 3D tissue block to learn an implicit tissue model from degredation free images from close to surface of the tissue block, to construct restored images of the deeper tissue. We do not require a degredation model or explicit ground truth training data. We note that the Cycle-GAN architecture has shown great promise in other biomedical applications including denoising (Elad and Aharon, 2006; Li et al., 2011), segmentation (Zhang et al., 2012) and restoration (Ma et al., 2020). Our experiments using thick tissue sections of the mouse heart imaged using confocal microscopy demonstrate the superior qualitative and quantitative performance of our approach over other methods.

2. Methods

As explained earlier, the challenge in our restoration problem is that we cannot explicitly model the degradation process. Hence we exploit the power of a Cycle-GAN, demonstrating the first computer vision application to restoring deeper tissue layers in confocal microscopy images. In our experiments we demonstrate our method on tissue that has been actively cleared prior to confocal imaging (Richardson and Lichtman, 2015); this should not be confused with our figurative use of the word "clearing" in the paper, which refers to the Cyle-GAN based restoration process of the deeper layers.

2.1. Cycle-Consistent Generative Adversarial Networks

Our essential goal is to increase the effective depth for analysis of tissue images obtained in microscopy. Whereas active clearing methods such as the Clarity protocol (Chung et al., 2013; Richardson and Lichtman, 2015) make the tissue sample translucent and thus increase light penetration, the 3D microscopy images still suffer from depth dependent degradation in the optical signal. This depth degradation is caused by a combination of the scattering of light, and non-uniformities in the tissue sample and tissue preparation procedure, and the uneven penetration of stains used to highlight particular tissue components. Whereas these effects reduce the quality of images obtained from deeper layers, modelling these sources of degradation analytically is not feasible.

For a specific block of tissue, while the shape and structure of organelles can differ from region to region, the general underlying cellular structure is similar in deep and shallow layers. The microscopy images of the deeper layers suffer from an unknown non-linear degradation. Cycle-Consistent Generative Adversarial Networks (Cycle-GAN) (Zhu et al., 2017) provide an adversarial approach to learn unpaired image to image translations, and have been used successfully for applications including unsupervised image super-resolution (Yuan et al., 2018) and image denoising in medical imaging (Kang et al., 2019; Song et al., 2020; Li et al., 2019). In the present article we describe the first use of a Cycle-GAN to restore the degraded images in the deeper layers of such confocal microscopy images. We use the forward and a backward networks of a Cycle-GAN to model the degradation from shallow to deep tissue and vice-versa respectively. The network architecture is illustrated in Fig. 1, with the domain of shallow clear images being Y and that of the deeper degraded images being X. The two transformation networks learn the mapping between domains, given by $G_s : X \to Y$ and $F_d : Y \to X$. Here the subscripts refer to the parameters of the shallow mapping and deep mapping neural networks. To discriminate the translated "fake" samples in each domain we have two discriminator networks $D^X_\psi : X \to [0, 1]$ and $D^Y_\xi : Y \to [0, 1]$. Let the data distribution of domain X be $p(x)$ and that of domain Y be $p(y)$. Then the adversarial objective of the model (Zhu et al., 2017) is

$$\mathcal{L}_{adv}(s, \xi) = \mathbb{E}_{y\sim p(y)}[log D^Y_\xi(y)] + \mathbb{E}_{x\sim p(x)}[log(1 - D^Y_\xi(G_s(x))], \tag{1}$$

$$\mathcal{L}_{adv}(d, \psi) = \mathbb{E}_{x\sim p(x)}[log D^X_\psi(x)] + \mathbb{E}_{y\sim p(y)}[log(1 - D^X_\psi(F_d(y))]. \tag{2}$$

We also enforce cycle consistency so that an image x in deep domain X when translated by $G_s()$ and then $F_d()$ should be similar to x, i.e, $F_d(G_s(x)) \approx x$. For this we use an additional cycle-consistency loss given by

$$\mathcal{L}_{cycle}(\theta, \phi) = \mathbb{E}_{x\sim p(x)}[||F_\phi(G_\theta(x)) - x||_1)] + \mathbb{E}_{y\sim p(y)}[||G_\theta(F_\phi(y)) - y||_1)]. \tag{3}$$

This leads to the familiar cycle-GAN min max optimization problem given by

$$\min_{s,d} \max_{\psi,\xi} \mathcal{L}_{adv}(s, \xi) + \mathcal{L}_{adv}(s, \psi) + \lambda \mathcal{L}_{cycle}(s, d). \tag{4}$$

Here, λ is a hyper-parameter, which we set to 10 for our experiments. The value was chosen after experimentation with different values. Empirically, a value of 10 led to convergence of both the discriminator as well as the generator networks.

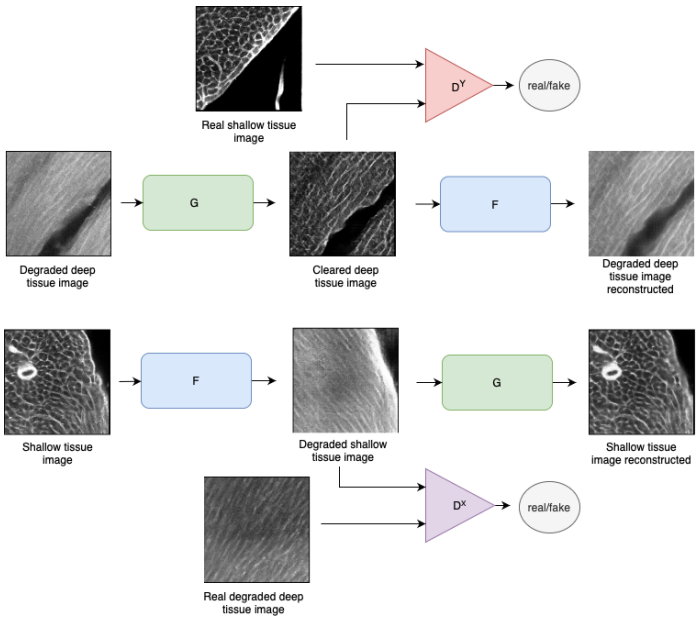

Figure 1: A schematic of the Cycle-GAN network. See Section 2.1 for a discussion.

For our translation networks, we used an architecture based on 9 Resnet blocks. The Residual network architecture with skip connections is suited to the task of image restoration and has been used for denoising as well as other image restoration tasks. For example (Zhang et al., 2017) exploits the residual framework for denoising Gaussian degraded images. While (Zhang et al., 2017) uses a single residual unit we use a stacked layer of 9 Resnet blocks in out generator to model our complex degredation process. Both the generators are identical in their architecture, consisting first of three convolutional layers with Relu activation. The first three layers consist of 64, 128 and 256 filters with a stride of 1, 2 and 2 respectively. This is followed by 9 Resnet blocks followed by two layers of transposed convolution of stride 2 and a final layer of convolution followed by *tanh* activation. We used the Adam optimizer with a learning rate of 0.002 for optimization. The optimizer as well as the learning rate were chosen empirically, and worked well.

3. Experiments

3.1. Datasets

The heart wall presents particular challenges in deep tissue imaging via confocal microscopy, due to its density and consequently high degree of tissue opacity. To test our methods, we have used cardiac tissue obtained from two different wild-type mice. The tissue has been optically cleared using the Clarity method (Chung et al., 2013), followed by sectioning in axial and coronal planes and staining with fluorescently labeled wheat germ agglutinin (WGA). WGA serves as a marker for the cell membranes of cardiac myocytes, and also capillaries and vessel boundaries. With the optical working distance of the objective lens set to resolve at $2\mu m$, we were able to image to a depth of $350\mu m$ with an Olympus FV3000

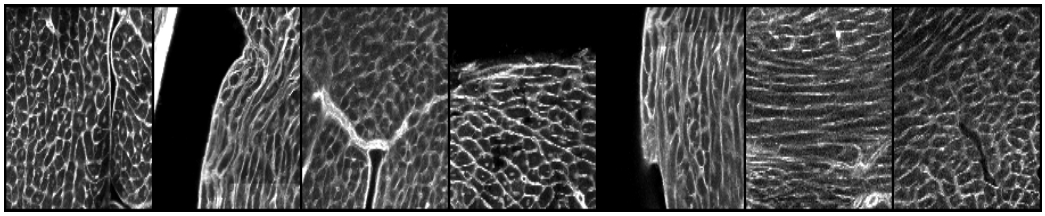

Figure 2: Sample shallow confocal microscopy images of mouse tissue.

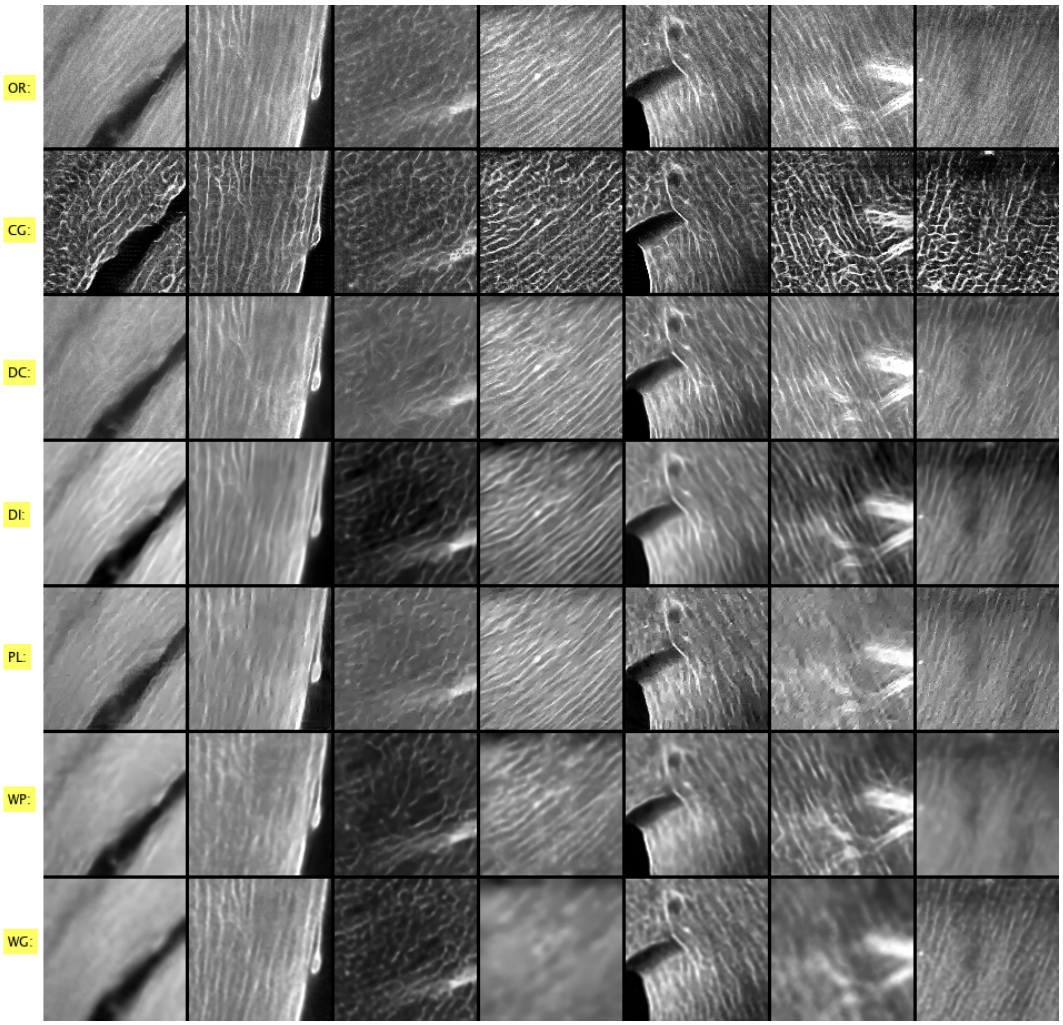

Figure 3: First Row: Sample original (OR) deep confocal microscopy images, which show degradation effects. Second - Seventh Row: The images from the top row, restored using the Cycle-GAN (CG), DnCnn (DC), Dictionary Learning (DI), Pure LET (PL), Wiener-Poisson (WP) and Wiener-Gaussian (WG) methods.

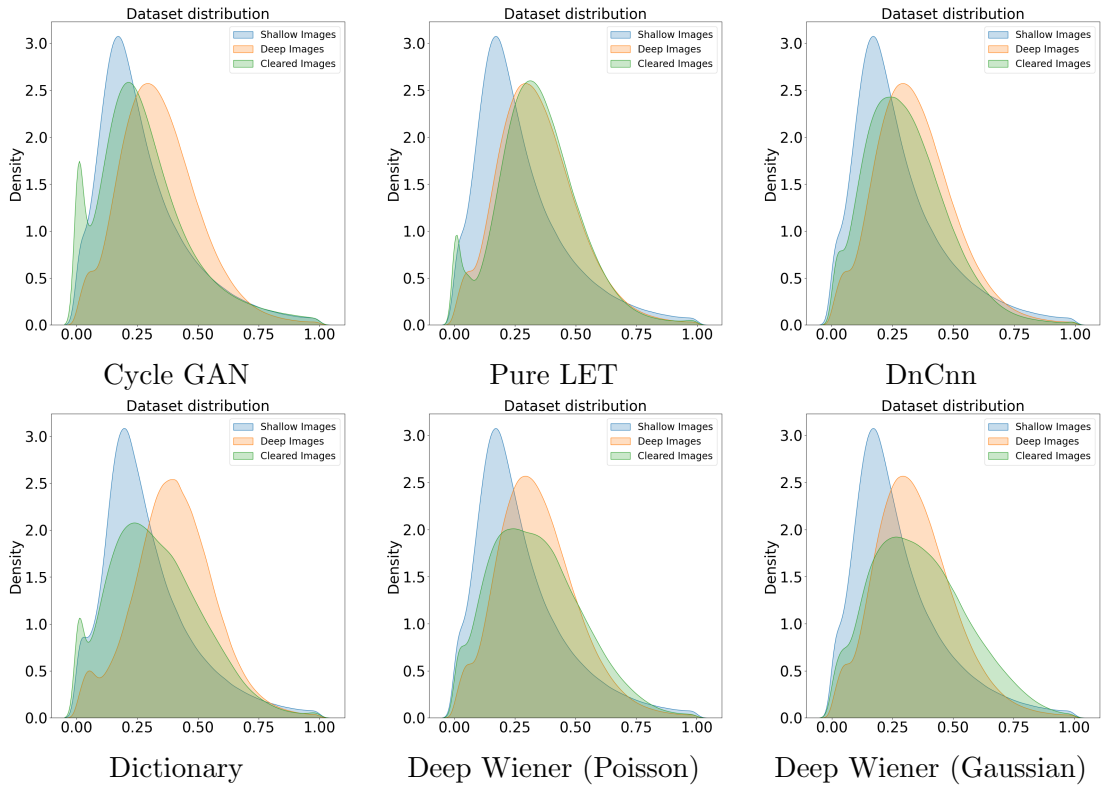

Figure 4: A comparison of histograms across methods.

confocal microscope. We divide our images into two sets: a 'deep' set consisting of 1225 images acquired at depths beyond $200\mu m$ from the surface of the tissue block closest to the objective and a 'shallow' set consisting of 1225 images acquired at depths upto $60\mu m$ from the surface. The datasets were generated from tissue samples from two different animals. We randomly sampled overlapping regions of size 128×128 pixels with 25% overlap. We further applied a manual quality control to remove any borderline images from each set. Fig. 2 shows typical samples from shallow layers, depicting sharp myocyte and capillary boundaries. Fig. 3 (top row) shows typical samples from deep layers, showing the effects of degradation.

3.2. Experiments and Results

In the dictionary learning experiments we used the full set of shallow images to learn the dictionary D and the full 'deep' set for testing. The image intensities in the 'shallow' as well as the 'deep' set are approximately distributed as a gamma distribution, with peaks at around 0.2 and 0.4 respectively, as illustrated in Fig. 4. The dictionary D was then used to clear the images from the 'deep' set. To quantify the results we used the following two similarity metrics: a) the Structural Similarity Image Metric (SSIM) (Wang et al., 2004) between the cleared and the deep image, and b) the Hellinger Similarity H_s, which we define as the Hellinger Distance between two normalized intensity histograms subtracted from 1:

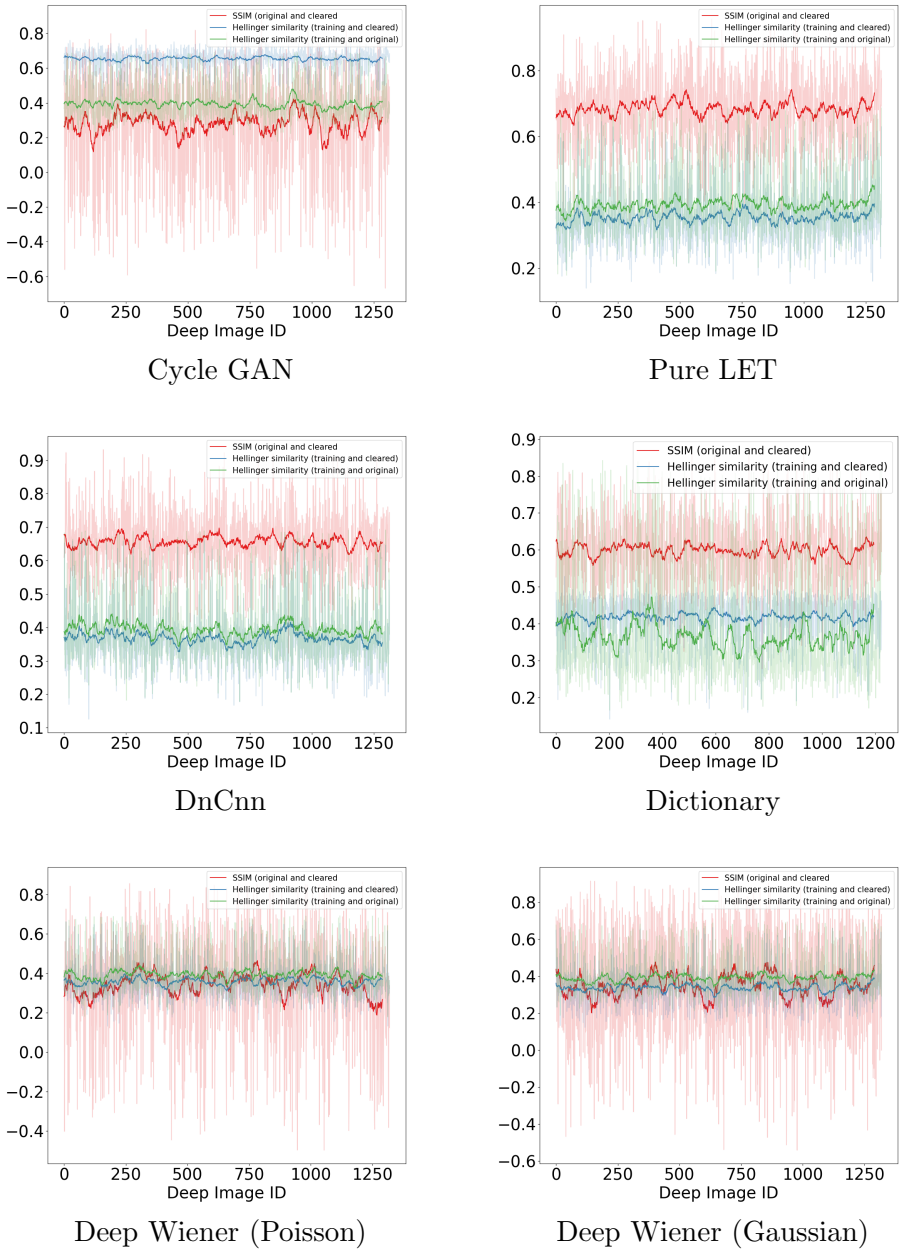

Figure 5: A comparison of SSIM and Hellinger similarity across methods.

$H_s = 1 - \sqrt{1 - \sum_{\text{bins}} \sqrt{\frac{H_a \cdot H_b}{\sum H_a \cdot \sum H_b}}}$. The SSIM is calculated between an original deep image and its cleared version, and H_s is computed between the normalized intensity distribution of a deep image and that of the 'shallow set', and also that of the corresponding cleared image and that of the 'shallow' set. Fig. 5 shows these three metrics for each deep test image - the

dark curve is a running average. Fig. 4 (bottom, left) shows the intensity distribution of the 'cleared' set, which has shifted towards that of the 'shallow' set. This shift is also apparent in Fig. 5 (middle, right), where the H_s between each cleared image and the 'shallow' set shows a slight increase from below 0.4 to above 0.4, across the dataset. The SSIM between the 'cleared' and 'deep' images also remains stable at around 0.6, indicating that the cleared image is structurally similar to the original degraded deep image. To test the learning based models including our Cycle-GAN model, the DnCnn model, the Pure LET model, and the Wiener-Kolmogorov(WK) model we carry out a similar set of experiments. Here we further divide our 'shallow' set into 'TrainS' and 'TestS' subsets and the 'deep' set into 'TrainD' and 'TestD' subsets, both containing 1169 and 56 images, respectively. We use these 4 sets to train our Cycle-GAN model for 200 epochs, with a batch size of 2. For DnCnn we used TrainS to train the model for 365000 epochs with a batch size of 64 and for the WK model we use a pre-trained UNet based method with a peak Poission noise level of 50 for WP and a Gaussian noise of variance 0.01 for WG. We tested all models on a separate validation set of 1315 images, distinct from the one used during training.

3.3. Discussion

We observe that with the Cycle-GAN, the restored image distribution better matches that of the 'shallow' set. The restored images are closer to the shallow ones in terms of their intensity histograms. The SSIM metric decreases (Fig. 4 (top left)) because each deep image has undergone more alteration when being translated to its restored version. There is also a much greater increase in H_s between the original 'deep' and 'shallow' images, and the 'cleared' and 'shallow' ones (from 0.4 to 0.6). These metrics are also supported by the qualitative results in Fig. 3. In Pure LET we observe that the histograms of the deep and restored images are quite close. This is also evident from a qualitative comparison of row 1 and 5. Similarly, a slight shift in the distribution for the case of DnCNN is apparent in row 3. While this method does restore some of the structure, the images still have the washed out look of the original degraded images in row 1, although they are less grainy. From Fig. 3 it is clear that the Cycle-GAN restorations most closely match the appearance of the shallow (sharp) images in Fig. 2. We note that Cycle-GAN model is based on the assumption that we are able to acquire good quality shallow images to train the system. A degradation such as one caused by loss in stain penetration strength, could lead to unreliable reconstruction.

4. Conclusion

The application of computer vision to the analysis of deep tissue images from microscopy is an emerging area of research, where methods based on deep learning (Weigert et al., 2018; Li et al., 2017; Pronina et al., 2020; Xiao et al., 2020; Ma et al., 2020) as well as more classical approaches (Nasser and Boudier, 2019) can be applied. Here we have shown the promise of a Cycle-GAN to clear image degradation effects at increased depths in a tissue stack from confocal microscopy. Our method can be applied to any 3D tissue sample which is self-similar so long as the ultra-structure features in the shallow layers are similar to those in the deeper ones. Our analysis shows that a Cycle-GAN produces cleared results that are closer to those in the shallow layers, both in a qualitative sense (Fig. 3) and by quantitative measures (Figs. 4 and 5) to those produced by other methods.

References

S. Abadie, C. Jardet, J. Colombelli, B. Chaput, A. David, J.-L. Grolleau, P. Bedos, V. Lobjois, P. Descargues, and J. Rouquette. 3d imaging of cleared human skin biopsies using light-sheet microscopy: A new way to visualize in-depth skin structure. *Skin Research and Technology*, 24(2):294–303, 2018.

Kwanghun Chung, Jenelle Wallace, Sung-Yon Kim, Sandhiya Kalyanasundaram, Aaron S Andalman, Thomas J Davidson, Julie J Mirzabekov, Kelly A Zalocusky, Joanna Mattis, Aleksandra K Denisin, et al. Structural and molecular interrogation of intact biological systems. *Nature*, 497(7449):332–337, 2013.

Nicolas Dey, Laure Blanc-Feraud, Christophe Zimmer, Pascal Roux, Zvi Kam, Jean-Christophe Olivo-Marin, and Josiane Zerubia. Richardson–lucy algorithm with total variation regularization for 3d confocal microscope deconvolution. *Microscopy research and technique*, 69(4):260–266, 2006.

Michael Elad and Michal Aharon. Image denoising via sparse and redundant representations over learned dictionaries. *IEEE Transactions on Image processing*, 15(12):3736–3745, 2006.

Annette Feuchtinger, Axel Walch, and Michael Dobosz. Deep tissue imaging: a review from a preclinical cancer research perspective. *Histochemistry and cell biology*, 146(6):781–806, 2016.

Eunhee Kang, Hyun Jung Koo, Dong Hyun Yang, Joon Bum Seo, and Jong Chul Ye. Cycle-consistent adversarial denoising network for multiphase coronary ct angiography. *Medical physics*, 46(2):550–562, 2019.

Jizhou Li, Florian Luisier, and Thierry Blu. Pure-let image deconvolution. *IEEE Transactions on Image Processing*, 27(1):92–105, 2017.

Shutao Li, Leyuan Fang, and Haitao Yin. An efficient dictionary learning algorithm and its application to 3-d medical image denoising. *IEEE transactions on biomedical engineering*, 59(2):417–427, 2011.

Z. Li, J. Huang, L. Yu, Y. Chi, and M. Jin. Low-dose ct image denoising using cycle-consistent adversarial networks. In *2019 IEEE Nuclear Science Symposium and Medical Imaging Conference (NSS/MIC)*, pages 1–3, 2019. doi: 10.1109/NSS/MIC42101.2019.9059965.

Yuhui Ma, Yonghuai Liu, Jun Cheng, Yalin Zheng, Morteza Ghahremani, Honghan Chen, Jiang Liu, and Yitian Zhao. Cycle structure and illumination constrained gan for medical image enhancement. In *International Conference on Medical Image Computing and Computer-Assisted Intervention*, pages 667–677. Springer, 2020.

Lamees Nasser and Thomas Boudier. A novel generic dictionary-based denoising method for improving noisy and densely packed nuclei segmentation in 3d time-lapse fluorescence microscopy images. *Scientific reports*, 9(1):1–13, 2019.

Valeriya Pronina, Filippos Kokkinos, Dmitry V Dylov, and Stamatios Lefkimmiatis. Microscopy image restoration with deep wiener-kolmogorov filters. In *Computer Vision–ECCV 2020: 16th European Conference, Glasgow, UK, August 23–28, 2020, Proceedings, Part XX 16*, pages 185–201. Springer, 2020.

Douglas S Richardson and Jeff W Lichtman. Clarifying tissue clearing. *Cell*, 162(2):246–257, 2015.

Joonyoung Song, Jae-Heon Jeong, Dae-Soon Park, Hyun-Ho Kim, Doo-Chun Seo, and Jong Chul Ye. Unsupervised denoising for satellite imagery using wavelet subband cyclegan. *arXiv preprint arXiv:2002.09847*, 2020.

Oleksandra Tmenova, Rémi Martin, and Luc Duong. Cyclegan for style transfer in x-ray angiography. *International journal of computer assisted radiology and surgery*, 14(10): 1785–1794, 2019.

Zhou Wang, Alan C Bovik, Hamid R Sheikh, and Eero P Simoncelli. Image quality assessment: from error visibility to structural similarity. *IEEE transactions on image processing*, 13(4):600–612, 2004.

Martin Weigert, Uwe Schmidt, Tobias Boothe, Andreas Müller, Alexandr Dibrov, Akanksha Jain, Benjamin Wilhelm, Deborah Schmidt, Coleman Broaddus, Siân Culley, et al. Content-aware image restoration: pushing the limits of fluorescence microscopy. *Nature methods*, 15(12):1090–1097, 2018.

Le Xiao, Chunyu Fang, Lanxin Zhu, Yarong Wang, Tingting Yu, Yuxuan Zhao, Dan Zhu, and Peng Fei. Deep learning-enabled efficient image restoration for 3d microscopy of turbid biological specimens. *Optics Express*, 28(20):30234–30247, 2020.

Chenyu You, Guang Li, Yi Zhang, Xiaoliu Zhang, Hongming Shan, Mengzhou Li, Shenghong Ju, Zhen Zhao, Zhuiyang Zhang, Wenxiang Cong, et al. Ct super-resolution gan constrained by the identical, residual, and cycle learning ensemble (gan-circle). *IEEE transactions on medical imaging*, 39(1):188–203, 2019.

Yuan Yuan, Siyuan Liu, Jiawei Zhang, Yongbing Zhang, Chao Dong, and Liang Lin. Unsupervised image super-resolution using cycle-in-cycle generative adversarial networks. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition Workshops*, pages 701–710, 2018.

Kai Zhang, Wangmeng Zuo, Yunjin Chen, Deyu Meng, and Lei Zhang. Beyond a gaussian denoiser: Residual learning of deep cnn for image denoising. *IEEE transactions on image processing*, 26(7):3142–3155, 2017.

Shaoting Zhang, Yiqiang Zhan, and Dimitris N Metaxas. Deformable segmentation via sparse representation and dictionary learning. *Medical image analysis*, 16(7):1385–1396, 2012.

Jun-Yan Zhu, Taesung Park, Phillip Isola, and Alexei A Efros. Unpaired image-to-image translation using cycle-consistent adversarial networks. In *Proceedings of the IEEE international conference on computer vision*, pages 2223–2232, 2017.