

PROCEEDINGS OF SPIE

SPIDigitalLibrary.org/conference-proceedings-of-spie

Conditional generative adversarial network for synthesizing hyperspectral images of breast cancer cells from digitized histology

Halicek, Martin, Ortega, Samuel, Fabelo, Himar, Lopez, Carlos, Lejeune, Marylene, et al.

Martin Halicek, Samuel Ortega, Himar Fabelo, Carlos Lopez, Marylene Lejeune, Gustavo M. Callico, Baowei Fei, "Conditional generative adversarial network for synthesizing hyperspectral images of breast cancer cells from digitized histology," Proc. SPIE 11320, Medical Imaging 2020: Digital Pathology, 113200U (16 March 2020); doi: 10.1117/12.2549994

SPIE.

Event: SPIE Medical Imaging, 2020, Houston, Texas, United States

Conditional generative adversarial network for synthesizing hyperspectral images of breast cancer cells from digitized histology

Martin Halicek^{a,b}, Samuel Ortega^{a,c}, Himar Fabelo^c, Carlos Lopez^{d,e},
Marylene Lejaune^{d,e}, Gustavo M. Callico^c and Baowei Fei^{a,f,g,*}

^a University of Texas at Dallas, Department of Bioengineering, Dallas, TX

^b Georgia Institute of Technology and Emory University, Department of Biomedical Engineering, Atlanta, GA

^c University of Las Palmas de Gran Canaria, Institute for Applied Microelectronics, Las Palmas, Spain

^d Hospital de Tortosa Verge de la Cinta, Department of Pathology, Tortosa, Spain

^e Universitat Rovira i Virgili, Tortosa, Spain.

^f University of Texas Southwestern Medical Center, Advanced Imaging Research Center, Dallas, TX

^g University of Texas Southwestern Medical Center, Department of Radiology, Dallas, TX

* E-mail: bfei@utdallas.edu; Web: www.feilab.org

ABSTRACT

Hyperspectral imaging (HSI), which acquires up to hundreds of bands, has been proposed as a promising imaging modality for digitized histology beyond RGB imaging to provide more quantitative information to assist pathologists with disease detection in samples. While digitized RGB histology is quite standardized and easy to acquire, histological HSI often requires custom-made equipment and longer imaging times compared to RGB. In this work, we present a dataset of corresponding RGB digitized histology and histological HSI of breast cancer, and we develop a conditional generative adversarial network (GAN) to artificially synthesize HSI from standard RGB images of normal and cancer cells. The results of the GAN synthesized HSI are promising, showing structural similarity (SSIM) of approximately 80% and mean absolute error (MAE) of 6 to 11%. Further work is needed to establish the ability of generating HSI from RGB images on larger datasets.

1. INTRODUCTION

In recent years, the use of hyperspectral imaging (HSI) for computer-aided histological sample analysis is increasing [1,2]. This trend is spurred by both increases in computational pathology and the richness of information that HSI can provide. However, the acquisition of hyperspectral (HS) images can present several challenges compared to standard RGB digitized histology. HSI technology can be combined with machine learning algorithms to retrieve useful information for diagnosis support. Training machine learning algorithms and especially deep learning algorithms requires a large amount of data for convergence and generalization. Although several researchers have shown promising results in the diagnosis of pathological slides using HSI, the challenges regarding current acquisition methods may limit histological HSI to a few selected regions of interest within a slide. Therefore, data augmentation, such as flips and rotations, can increase the quantity of HS data for sufficient training of HS computer-aided diagnostic algorithms.

To overcome these challenges, we propose a technique for faster generation of HS data from only RGB information. The synthetic HS data generated could be used for augmentation of HS machine learning or perhaps for other clinical purposes. However, the RGB to HS transformation is restricted to only the visual range from 400 to 750 nm, which is by definition the range where RGB images are acquired. On the other hand, HSI technology can work on different ranges, most commonly the visual range (400 to 750 nm) and the near infrared (NIR, 750-1000 nm) [1]. The proposed technique could help in two ways. First, the spatial resolution of the synthesized HS images could be very high, depending on the RGB camera employed. This fact is caused by the relatively low spatial resolution of HS cameras compared to conventional RGB cameras. Secondly, a NIR camera could be used to quickly complement the spectral range beyond visible light, and the synthesized HS images from RGB could be used to increase the spatial characteristics of the real HS NIR images, leading in a reduction of cost and complexity of a high spectral range HS system.

In this research, we propose the use of a generative adversarial network (GAN) to generate new HS data from cellular level histological RGB images. The main motivation is that the annotation can be easily performed in RGB digitized slides, where the tissue can be examined at different magnification levels. This is currently not feasible in HS images because the acquisition of whole-slide-images using a HS camera is currently a challenge itself. The

procedure employed to perform the experiments is as follows. First, a skilled pathologist annotated cells from selected regions within two histological RGB digitized slides from two patients with breast cancer. Next, several HS images were acquired with a HS microscope of these regions, and the images were registered with the correspondent RGB image. Finally, a conditional GAN was employed to generate synthetic HS images from conventional RGB digital images, which was compared with the original HS data.

2. METHODS

2.1 Breast Cancer Histology Hyperspectral Dataset

In this study, two histopathological slides from two patients with breast cancer were included. The tissue sample was paraffin embedded, sectioned, and stained with hematoxylin and eosin (H&E). The whole-slide was digitized using a Panoramic SCAN digital scanner, and a pathologist annotated the digitized slide image using Panoramic Viewer software (3D Histech Ltd., Budapest, Hungary). Cellular-level annotations of both normal and breast cancer cells were manually performed by a skilled pathologist. Histological sample selection, preparation, and annotation was performed by the department of pathology of the Hospital de Tortosa Verge de la Cinta, Spain. A push-broom HS camera (Hyperspec VNIR A-Series, HeadWall Photonics Inc., Fitchburg, MA, USA) was mounted to a conventional light microscope (Olympus BX-53, Tokyo, Japan), and a mechanical stage performed the movement necessary to capture the HS image. The regions of interest (ROIs) annotated by the pathologist were acquired from 400 to 1000 nm with a spectral resolution of 2.8 nm, producing 826 spectral bands. Standard calibration and a band reduction were performed to reduce the number spectral bands to 159 between 400 to 750 nm, according to a previously described method [2].

This study had two different imaging modalities, which required registration of the annotated regions from digitized RGB slides to the corresponding HS images using a synthetic RGB image obtained from the original HS image. Briefly, the previously described image registration process consists of a geometric transformation that matches the images, applied to the annotations of the HS image [2]. After image registration, patches are extracted from the digitized RGB slide and the HS image of corresponding cells that were annotated. From each cell centroid, patches are generated with a size of 39×39 spatial pixels and 159 spectral bands.

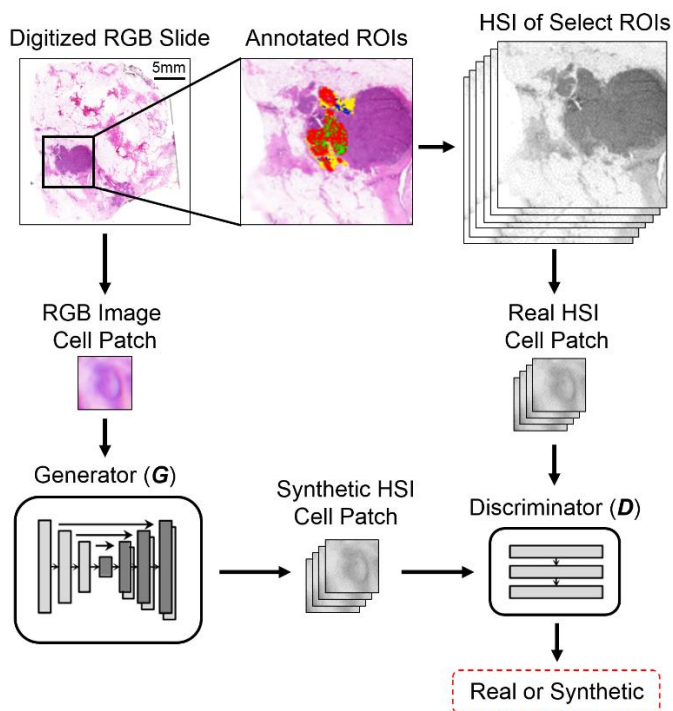


Figure 1: Schematic diagram of the experimental design. A standard RGB digitized whole-slide image is taken of the H&E slide of the breast cancer specimen. A pathologist manually annotated cell nuclei in selected ROIs (red is cancer, green is normal). HSI is acquired of the selected ROIs. Image patches are made that correspond from the RGB digitized images and the original HS images. The generator (G) is trained to make synthetic HS images from the RGB images. A discriminator (D) is trained to detect synthetic and real HS images to better train the generator.

2.2 Conditional Generative Adversarial Network

A conditional generative adversarial network (cGAN) was implemented to generate synthetic HS images from the digitized, whole-slide RGB images on a cell-by-cell level. Figure 1 shows the schematic diagram of the experimental design proposed in this work. After extracting a total of 20,800 cell patches from both RGB and HSI, the patches were split into training (65%), validation (10%), and testing (25%). The cGAN was inspired by the pix2pix design, which consists of a generator (G) to synthesize data and a discriminator (D) to detect real and synthetic pairs [3]. For better spatial reconstruction, the RGB image patches ($39 \times 39 \times 3$) were stacked to match the spectral resolution size of the HS image output patch ($39 \times 39 \times 159$), according to the visible spectrum orientation, i.e. 53 blue channels, 53 green channels, and 53 red channels. The generator was a modified U-Net consisting of 21 layers that used the stacked RGB image ($39 \times 39 \times 159$) as input and generated a synthetic HS image ($39 \times 39 \times 159$) as output. The full modified architecture of the generator is described in Table 1. The discriminator consisted of 4 strided convolutional layers to produce a final size of ($3 \times 3 \times 1$), which allows fine-level detail to be learned, referred to as patchGAN [3]. The network was developed using the TensorFlow implementation of the Keras Deep Learning API [4].

Table 1. Schematic of the Generator (G), which was a modified U-Net CNN.

Layer	Notes	Shape / Filters
Input	Real RGB (Stacked)	$39 \times 39 \times 159$
Conv2D	Strided, 3×3	$20 \times 20 \times 256$
Conv2D	Padded, 3×3	$20 \times 20 \times 512$
Conv2D	Strided, 3×3	$10 \times 10 \times 768$
Conv2D	Padded, 3×3	$10 \times 10 \times 1024$
Conv2D	Strided, 3×3	$5 \times 5 \times 1280$
Conv2D	Padded, 3×3	$5 \times 5 \times 1536$
Conv2D_Transpose	Strided, 3×3	$10 \times 10 \times 1536$
Concatenation	Skip Connection	$10 \times 10 \times 2560$
Conv2D	Padded, 3×3	$10 \times 10 \times 1280$
Conv2D_Transpose	Padded, 3×3	$20 \times 20 \times 1024$
Concatenation	Skip Connection	$20 \times 20 \times 1536$
Conv2D	Padded, 3×3	$20 \times 20 \times 768$
Conv2D_Transpose	Padded, 3×3	$40 \times 40 \times 320$
Cropping 2D	1 Pixel	$39 \times 39 \times 320$
Concatenation	Skip Connection	$39 \times 39 \times 479$
Conv2D	Padded, 3×3	$39 \times 39 \times 296$
Conv2D	Padded, 3×3	$39 \times 39 \times 284$
Conv2D	Padded, 3×3	$39 \times 39 \times 256$
Conv2D	Padded, 3×3	$39 \times 39 \times 220$
Conv2D	Padded, 3×3	$39 \times 39 \times 196$
Conv2D	Padded, 3×3	$39 \times 39 \times 159$
Output	Synthetic HSI	$39 \times 39 \times 159$

Training the cGAN consisted of first training the discriminator and then the generator. In Figure 2, we show the framework for the training of one iteration of the cGAN. First, the discriminator is trained on real and synthetic pairs using L2 loss. Next, the generator is trained using the combined model with L2 loss and alone with L1 loss to the real HS image. The loss functions for the discriminator were L2 loss. For training the generator, the G and D were combined but the weights from D were frozen, which allows calculating both L1 from the synthetic and real HS image pair and L2 loss from D . The optimizer used was stochastic gradient descent with a learning rate of 10^{-4} and a value of 0.90 for momentum. The cGAN was trained for 40 epochs, and optimization was performed using the validation set.

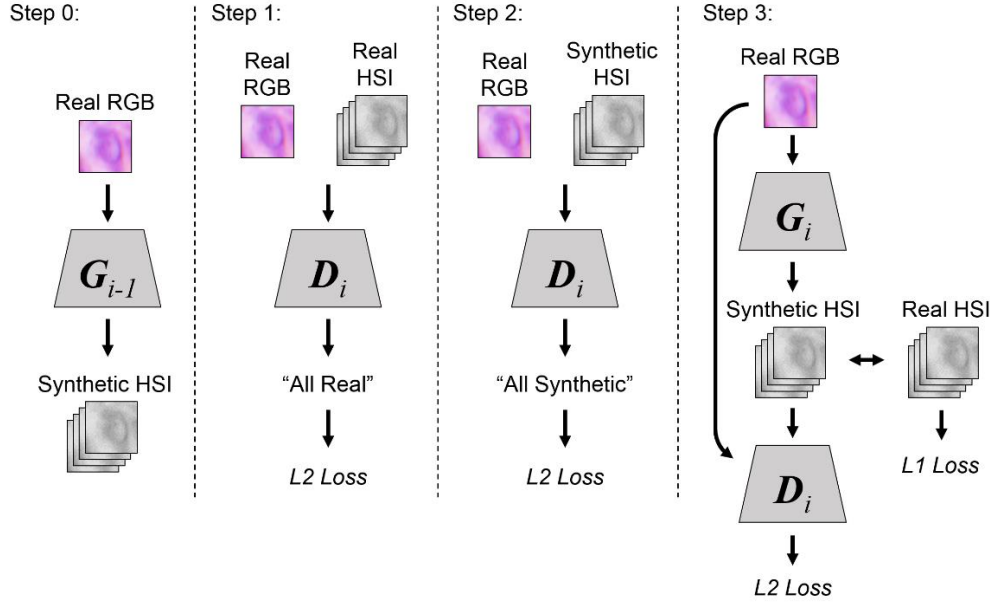


Figure 2. Framework for one iteration (i -th) of training the cGAN. The preliminary step is to generate the synthetic HS image from the previous training iteration (step 0). For training, first, the discriminator is trained on real (step 1) and synthetic (step 2) pairs using L2 loss. Next, the generator is trained using the combined model with L2 loss and alone with L1 loss to the real HS image (step 3).

2.3 Validation

To evaluate the performance of the GAN, the synthetic HS image of the cells are compared to the real HS image using well-established metrics for single-band, panchromatic images, which are then averaged across all spectral bands. The structural similarity (SSIM) index is calculated on a band-by-band basis, with one band (R) from the real HSI and the corresponding band (S) from the synthetic HS image, according to the standard equation (1), as defined in [5].

$$SSIM_{\lambda}(R, S) = \frac{(2\mu_R\mu_S + C_1)(2\sigma_{RS} + C_2)}{(\mu_R^2 + \mu_S^2 + C_1)(\sigma_R^2 + \sigma_S^2 + C_2)} \quad (1)$$

Where the constants, $C_1 = (0.01 \times L)^2 = 0.0001$ and $C_2 = (0.03 \times L)^2 = 0.0009$, are chosen as the default non-zero values dependent on the dynamic range (L) of the image values, and μ and σ are the mean and standard deviations of images. The total SSIM for a cell HSI is the average of all SSIM values from all 159 bands, according to the equation (2) below.

$$SSIM(\text{Real HSI}, \text{Synthetic HSI}) = \frac{1}{159} \sum_{\lambda=1}^{159} SSIM_{\lambda}(R, S) \quad (2)$$

Additionally, the peak signal to noise ratio (PSNR) is calculated in the same way by averaging the values obtained in a band-by-band basis across all 159 spectral bands.

$$PSNR(\text{Real HSI}, \text{Synthetic HSI}) = \frac{1}{159} \sum_{\lambda=1}^{159} PSNR_{\lambda}(R, S), \text{ where } PSNR_{\lambda}(R, S) = 10 \log_{10}(\text{peak value}^2 / \text{MSE})$$

Where MSE is the mean squared error term. The final evaluation metric was mean absolute error (MAE), which was also calculated on a band-by-band basis between the real and synthetic HS images, and the final result for an HS image was obtained by averaged the results of all spectral bands.

Parameter optimization was performed on the validation set using all the evaluation metrics above, and the final performance was calculated on the test set. The validation set was always from the same slide as the training set. Since two slides from two different patients were used in this study, the testing results were calculated on both, the same slide and the other patient's slide and reported separately. For calculating MAE, because we are interested in the HS signature of the cells, only the pixels within the annotated cell nuclei are used for calculation and background is excluded. For calculating both the SSIM and PSNR, the three pixels at each edge of the cell image-patch are not used for calculation, for the same reason. All performance calculations were performed in MATLAB (MathWorks, Inc., Natick, MA).

3. RESULTS

The quantitative experimental results obtained in this work are shown in Table 2. As stated previously, to measure the quality of synthetic HSI generation, the metrics utilized were SSIM, PSNR, and MAE. The SSIM and PSNR measures are more related to the spatial component of the HS image, while the MAE is more sensitive to errors in the spectral component. The results from the testing group were separated by normal cells and breast cancer cells. The average SSIM was around 80% for both cell types when testing on the same patient (intra-slide test). When testing on the other patient slide (inter-slide test), the average SSIM was slightly lower at 78%. However, the PSNR (18 dB) and the MAE (11%) were equivalent testing intra-slide or inter-slide patients. Additionally, the MAE of the normal cells was lower than breast cancer cells.

Table 2: Quantitative results of the cGAN-synthesized HSI of normal and cancer cells from both patients.

Group	Patient	Cell Type	No. of Cells	SSIM (%)	PSNR (dB)	MAE (%)	
Validation	<i>P1</i>	Combined	1281	80 ± 5	18 ± 3	10 ± 4	
Intra-slide	<i>P2</i>	Combined	798	81 ± 5	20 ± 5	10 ± 5	
Test	<i>P1</i>	Cancer	3106	80 ± 4	17 ± 2	11 ± 4	
		Normal	75	81 ± 6	21 ± 3	6 ± 2	
	<i>P2</i>	Cancer	1900	79 ± 6	18 ± 4	13 ± 7	
		Normal	91	84 ± 2	23 ± 2	6 ± 2	
	Average				80 ± 4	18 ± 3	11 ± 5
	Test	<i>P1</i>	Cancer	3106	76 ± 4	19 ± 2	10 ± 2
Normal			75	78 ± 7	20 ± 1	9 ± 1	
<i>P2</i>		Cancer	1900	81 ± 3	17 ± 3	13 ± 5	
		Normal	91	83 ± 2	14 ± 2	18 ± 5	
Average				78 ± 4	18 ± 2	11 ± 3	

The outcomes of the HS images generated using the proposed cGAN were investigated both qualitatively and quantitatively. The spatial components are shown in Figure 3 with single-band images and synthetic RGBs for both real and synthesized HS images, which shows a strong relationship between the real images and the synthetic ones. Nevertheless, although the most relevant spatial features of the cells are generated successfully by the cGAN, the spatial information is quite blurry compared to the sharp edges of the real HS image. In Figure 4 and 5, the spectral signatures from real and synthetic HS images are plotted for both normal and breast cancer cells, respectively. The synthetic HS image generated from intra-slide cGAN training is shown in blue, and inter-slide cGAN training is shown in red. It can be observed that our method provides a good approximation of the spectral signature of each type of cell. However, the normal cells (Figure 4) have more realistic signatures than cancer cells (Figure 5). Additionally, the intra-patient training yields more realistic spectral signatures, compared to inter-patient cGAN training.




Figure 3. Single-band images from HSI and RGB composite from HSI showing the spatial features of a breast cancer cell (left) and normal cell (right).

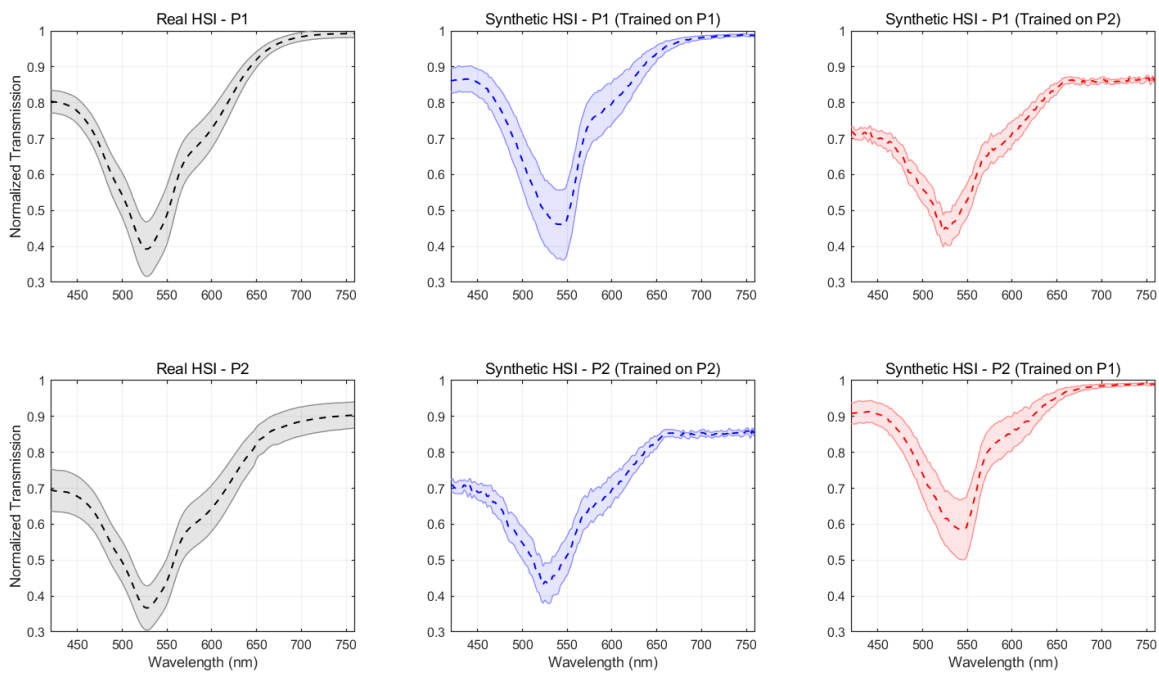


Figure 4. Plots of the spectral signatures of normal cells from both patients ($P1$ and $P2$). Real HS spectral signatures are shown in black (left). Synthetic HS spectral signatures using cGAN trained on same patient's data (intra-slide) are shown in blue (middle). Synthetic HS spectral signatures using cGAN trained on the other patient's data (inter-slide) are shown in red (right).

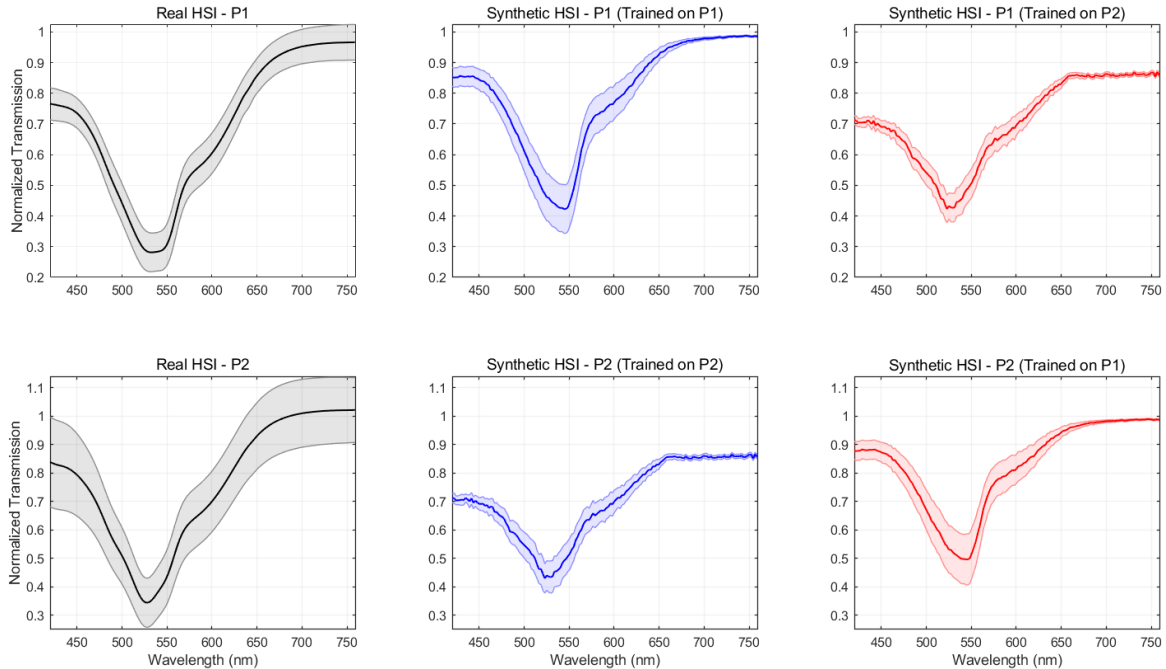


Figure 5. Plots of the spectral signatures of breast cancer cells from both patients (*P1* and *P2*). Real HS spectral signatures are shown in black (left). Synthetic HS spectral signatures using cGAN trained on same patient's data (intra-slide) are shown in blue (middle). Synthetic HS spectral signatures using cGAN trained on the other patient's data (inter-slide) are shown in red (right).

4. DISCUSSION

In this work, we present the use of conditional GAN to use digitized RGB information to construct artificial HS images of normal and breast cancer cells from two histology slides. The novelty of this approach is the dataset of an annotated slide with both HSI and RGB components, which allows training a GAN. In the literature, GANs have been utilized for artificial staining of histological samples with HSI [6], but this is the first application for direct synthesis of HS data from RGB data for breast cancer histology on a cellular level. Future work is necessary to expand the dataset to more patient slides for training to allow better generalization to new slides. Additionally, other organ systems, cell types, and even other staining techniques need to be explored using the proposed method.

From the preliminary results shown in this study, it was observed that normal cells have better spectral representations synthesized from RGB than cancer cells. The reason for this could be that there is more consistency in spectral signatures of normal cells and larger heterogeneity in the signatures of breast cancer cells (Figure 5). Both the shape and standard deviation of the signatures are more realistic for normal cells (Figure 4). The relative shapes and peaks appear worse for cancer cells. Most notably, the standard deviation for real spectral signatures of cancer cells from the second patient is quite substantial, but the synthetic curves do not have the same heterogeneity.

The results from intra-patient compared with inter-patient cGAN training appear to be equivalent quantitatively, with both around 11% error. However, inspecting the spectral signatures in Figures 4 and 5, more appropriate curves are generated when the GAN is trained on the same patient's data. This result indicates that more patient data is required for training and generalization. Additionally, the optimization of the modified U-net generator component of the GAN was done on the *P1* experiment. As a result of this, the synthetic spectral representations look better for this patient, compared to *P2*. The GAN architecture and hyper parameters should be optimized on a large dataset with cells from multiple patients.

In conclusion, histological HSI may be able to provide more information for computer-aided diagnostic algorithms compared to standard RGB digitized histology, but HSI is currently more difficult to acquire, especially in whole-slide fashion. In this work, we present a dataset of corresponding RGB digitized histology and histological HSI of breast cancer and develop a conditional GAN to artificially synthesize HSI from standard RGB images of normal and cancer cells. The results of the GAN synthesized HS images are promising, showing SSIM of approximately 80% and mean absolute error of 11%. Further work is needed to establish the ability of generating HSI from RGB images on larger datasets with more diverse cancer morphologies, different organ systems, and different staining techniques. Additionally, it is necessary to demonstrate the ability of GAN-synthesized HS data from RGB to improve training and classification ability of CNNs when the GAN is utilized for data augmentation proposes. Alternatively, the method of GAN-synthesized data combined with real HS data for extending spectral or spatial information and reducing costs needs to be investigated.

Disclosures

The authors have no relevant financial interests in this article and no potential conflicts of interest to disclose.

ACKNOWLEDGEMENTS

This research was supported in part by the U.S. National Institutes of Health (NIH) grants (R01CA156775, R01CA204254, R01HL140325, and R21CA231911) and by the Cancer Prevention and Research Institute of Texas (CPRIT) grant RP190588. Additionally, this research was supported in part by the Canary Islands Government through the ACIISI (Canarian Agency for Research, Innovation and the Information Society), ITHACA project under Grant Agreement ProID2017010164. This work was completed while Samuel Ortega was beneficiary of a pre-doctoral grant given by the “*Agencia Canaria de Investigacion, Innovacion y Sociedad de la Informacion (ACIISI)*” of the “*Conserjería de Economía, Industria, Comercio y Conocimiento*” of the “*Gobierno de Canarias*”, which is part-financed by the European Social Fund (FSE) (*POC 2014-2020, Eje 3 Tema Prioritario 74 (85%)*).

REFERENCES

- [1] Halicek, M., Fabelo, H., Ortega, S., et al. “In-Vivo and Ex-Vivo Tissue Analysis through Hyperspectral Imaging Techniques: Revealing the Invisible Features of Cancer,” *Cancers*, vol. 11, no. 6, p. 756 (2019).
- [2] Ortega, S., Halicek, M., Fabelo, H., et al. “Hyperspectral imaging and deep learning for the detection of breast cancer cells in digitized histological images,” *Proc. SPIE. Medical Imaging 2020: Digital Pathology*. 11320-7 (2020).
- [3] Isola, P., Zhu, J.Y., Zhou, T., and Efros, A.A. “Image-to-Image Translation with Conditional Adversarial Networks,” *arXiv: Computer Vision and Pattern Recognition*. arXiv/1611.07004 (2016).
- [4] Abadi, M. et al., “TensorFlow: Large-Scale Machine Learning on Heterogeneous Distributed Systems,” Mar. 2016.
- [5] Zhou, W., A. C. Bovik, H. R. Sheikh, and E. P. Simoncelli. “Image Quality Assessment: From Error Visibility to Structural Similarity,” *IEEE Transactions on Image Processing*. Vol. 13, Issue 4, pp. 600–612 (2004).
- [6] Bayramoglu, N., Kaakinen, M., Eklund, L., and Heikkilä, J. “Towards Virtual H&E Staining of Hyperspectral Lung Histology Images Using Conditional Generative Adversarial Networks,” *2017 IEEE International Conference on Computer Vision Workshops (ICCVW)*, Venice, Italy, pp. 64-71. (2017).